# FLIR Systems Infrared Imaging Radiometry Handbook

May 5, 2021

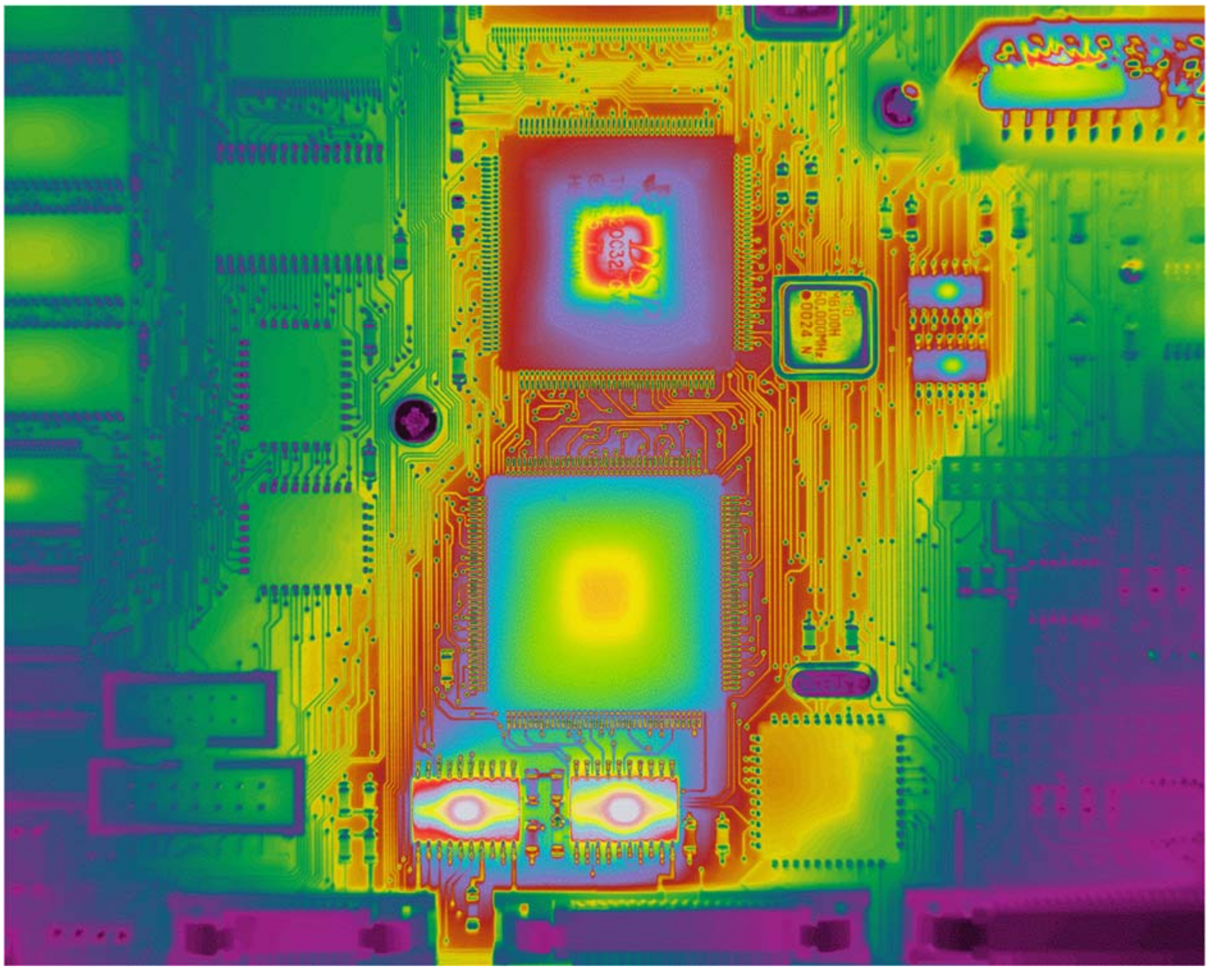


Austin Richards, Ph.D.

FLIR Systems, Inc.

This book was written to help FLIR customers understand the theory and practice of infrared imaging radiometry, and how to make radiometric measurements with FLIR's science cameras working in concert with FLIR ResearchIR software. Please send comments and feedback to austin.richards@flir.com. I would be particularly grateful if you point out any errors in the text or figures, and I would love to learn about novel applications for infrared imaging radiometry.

Austin Richards, Ph.D.
Senior Principal Research Scientist, FLIR Systems, Inc., Goleta, CA
SPIE Fellow

# Table of Contents:

**Chapter 1:** Imaging Radiometry

**Chapter 2:** How to Make NUCs

**Chapter 3:** Radiometric User Calibration

**Chapter 4:** Radiometric Measurement Accuracy

**Chapter 5:** Optical Filters in Radiometric Cameras

**Chapter 6:** Image Noise in Infrared Camera Systems

**Chapter 7:** MTF and Optics in Infrared Camera Systems

**Chapter 8:** Displaying Infrared Image Data - AGC and Color Palettes

**Chapter 9:** Step by Step User Calibration

**Chapter 10:** User and Factory Calibration Comparison

**Chapter 11:** Changing Include Files

**Chapter 12:** Generating NUC Tables

# Chapter 1 : Imaging Radiometry

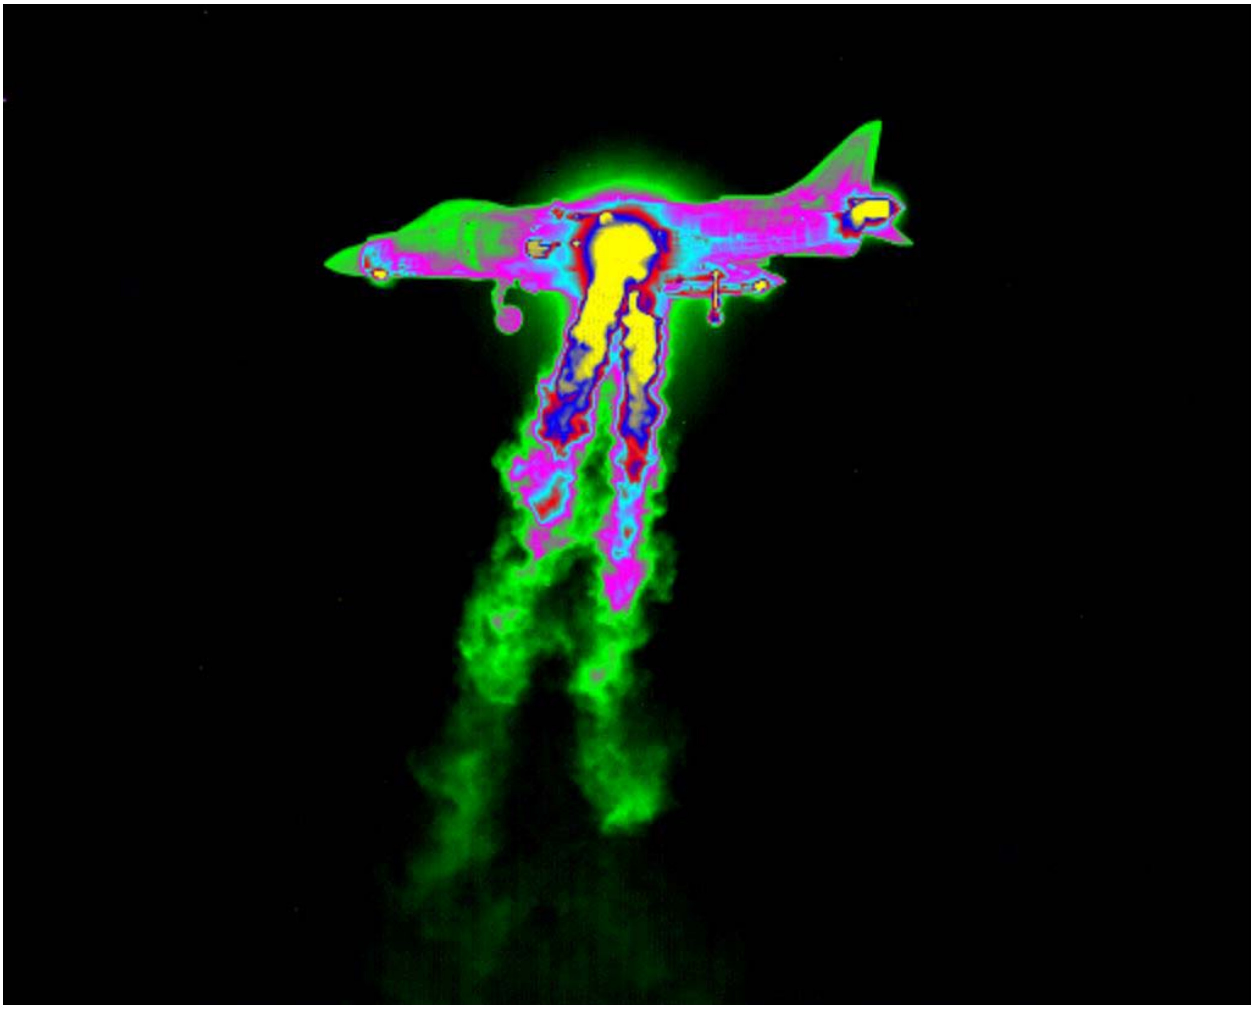

**Midwave infrared image of a Harrier jet hovering. The color palette is Color Wheel 6.**

Introduction ........ 1—5
Radiometric Calibration ........ 1—8
Radiometric Cameras ........ 1—8
Optics ........ 1—14
Standard Focal Lengths ........ 1—14
InGaAs and VisGaAs Optics: ........ 1—16
Non-Standard Optics ........ 1—16
Radiometric Calibration ........ 1—16
Calibration Ranges ........ 1—17
Radiance Ranges ........ 1—18
Exposure Control ........ 1—20
$T_{Drift}$ Correction ........ 1—21
Integration Times and Subranges ........ 1—21
Coldshield f/number ........ 1—22
Neutral Density Filters ........ 1—24
Spectral Ranges of Different Detectors ........ 1—24
Spectral Filters ........ 1—24
Dewars and Coolers ........ 1—25
Infrared Imaging Radiometry ........ 1—27
How does infrared imaging radiometry work? ........ 1—31
Radiometric Calibration Process ........ 1—32
Optics and the optical path ........ 1—35
Camera Radiance Linearity ........ 1—38
Planck's Law ........ 1—39
User Calibration: ........ 1—41
Calculating Radiance ........ 1—43
Calculating Temperature ........ 1—47
Stability ........ 1—49
Repeatability ........ 1—49
Short Integration Time and $T_{drift}$ ........ 1—50

Radiant Intensity Measurement ........ 1—50
Radiant Intensity of Point Targets ........ 1—56
Superframing ........ 1—57
Summary ........ 1—64

## Introduction

Infrared imaging radiometry is the science of using an infrared camera to measure the amount of IR radiation from a target in a scene. Every pixel in the image formed by a radiometrically-calibrated camera carries quantitative information. Radiometric images are more than just pictures; they are also two-dimensional grids of measurements of the infrared radiation imaged onto the sensor from the scene. The camera lens "maps" the IR radiation emitted from objects in the scene to corresponding positions on the sensor. The applications for infrared radiometry are plentiful, and well understood by scientists and engineers that work with this equipment. The most common industrial application for radiometry is non-contact surface temperature measurement, but military test ranges often measure radiometric quantities of targets and scenes such as radiance and radiant intensity[1], rather than temperature. It is worth noting right from the start that IR camera measurements of temperature are really measurements of *apparent* temperature based on emitted radiation. The camera makes a measurement of radiation from a target or scene and converts the measurement into temperature units. The actual physical or kinetic surface temperature of the target can be very different from the apparent temperature, and the process of correcting an apparent temperature back to the physical temperature can often be very challenging. Going forward, when temperature measurement is mentioned, the "apparent" nature of the measurement is implied.

In the same vein, both radiance and radiant intensity measurements made with a calibrated infrared camera are really measurements made in a spectral sub-band of the infrared band, because infrared cameras do NOT measure all wavelengths of electromagnetic radiation. It is therefore more correct to call them "in-band radiance" and "in-band radiant intensity". As with apparent temperature, the "in-band" part is always implied in this book. Radiance and radiant intensity units will be discussed in more detail later, but Figures 1a and 2a below are a preview of what radiance image data look like. Radiance and temperature measurements are shown in these two radiometric images of a Beechcraft King Air with running engines. The original image data is in units of digital counts – in this case a 14-bit binary number that ranges up to a maximum value of 16,383. The FLIR software application ResearchIR was used to capture the digital image data, and with a radiometric calibration in the camera itself or on the host computer, the software can process these images into radiance or temperature units of measure.

[1] These parameters will be defined in this chapter.

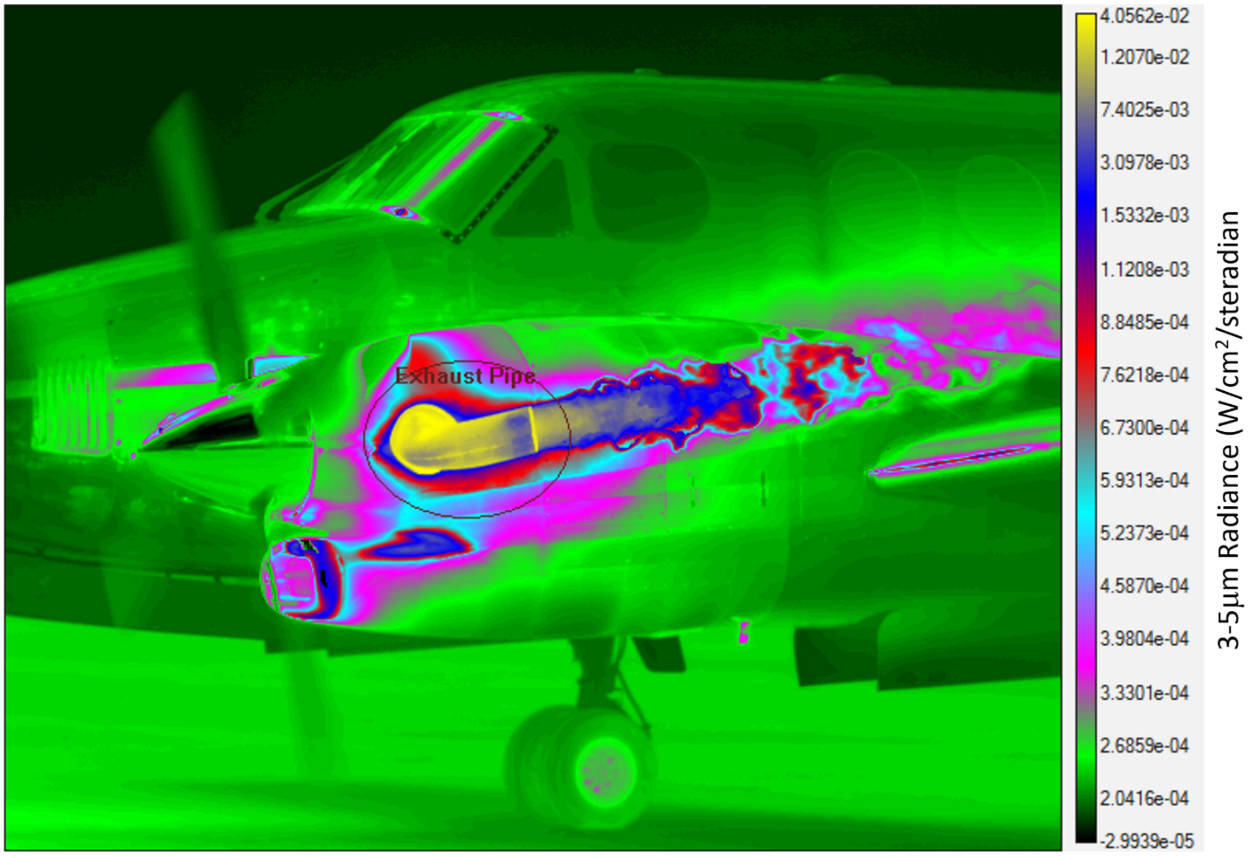


| Statistic [units] | Exhaust Pipe |
|---|---|
| Mean [W/(sr-cm²)] | 5.9770e-03 |
| Std. Dev. [W/(sr-cm²)] | 8.7230e-03 |
| Intensity [W/sr] | 1.4807e+01 |
| Center [W/(sr-cm²)] | (279.5, 262.5) 1.1715e-02 |
| Maximum [W/(sr-cm²)] | (257, 282) 4.0562e-02 |
| Minimum [W/(sr-cm²)] | (308, 221) 3.1948e-04 |
| Number of Pixels | 9162 |
| Single Pixel Area [cm²] | 0.2704 |
| Area [cm²] | 2477.4 |
| Length [cm] | N/A |
| u Emissivity ☐ | 1 |
| u Distance [m] ☑ | 20.8 |

**Figure 1a-b. Radiometric image of Beechcraft King Air during engine warmup (a) and radiance measurement statistics on an elliptical region of interest that is labeled "Exhaust Pipe" (b)**

Figure 1a shows the measurement area in the image - an elliptical region of interest (ROI) centered over the hottest part of the engine's exhaust pipe. The first measurement in the statistics window (Figure 1b) is the mean radiance within the region of interest in units of watts/square cm/steradians. The second measurement is the standard deviation of the radiance within the ROI. The Intensity field says 14.80 Watts/steradian – it is actually radiant intensity, but it's called Intensity for short. The derivation of radiant intensity from radiance will be discussed later in this chapter. Radiance values for the center, min and max of the ROI and the number of pixels in the ROI come next. The ROI distance is set to 20.8 meters – this is the distance from the camera to the objects in the ROI. A combination of this distance and a spatial calibration that defines the angular subtense of a single pixel yields the single pixel area of 0.25 sq. cm, and the ROI area which is 2477 square centimeters. The color bar on the right side of the image allows the user to estimate radiance in any part of the image based on the false color palette applied. This color palette is called Color Wheel 6 and was designed to produce a sensation of high contrast to the human visual system.

The same raw image data can also be processed into temperature units, as shown in Figure 2a below. The software indicates that the mean apparent temperature in the ROI is 120.2 °C, as shown in Figure 2b. The other statistics are the same as in the radiance case, except in temperature units. The colors in the color bar indicate that the pipe temperature is over 200 °C, and the wheel hub is around 40 °C. The assumption with the color bar values is that everything is a perfect blackbody emitter with an emissivity of 1. The emissivity of the ROI was also set to 1, but one can use different values for ROI emissivity. The choice of the value 1 returns apparent temperatures as if everything in the ROI was a perfect blackbody. This is called the **apparent blackbody temperature.** It may not match the kinetic surface temperature unless the surface has an in-band emissivity that is close to 1. Emissivity can vary greatly with surface finish and material type.

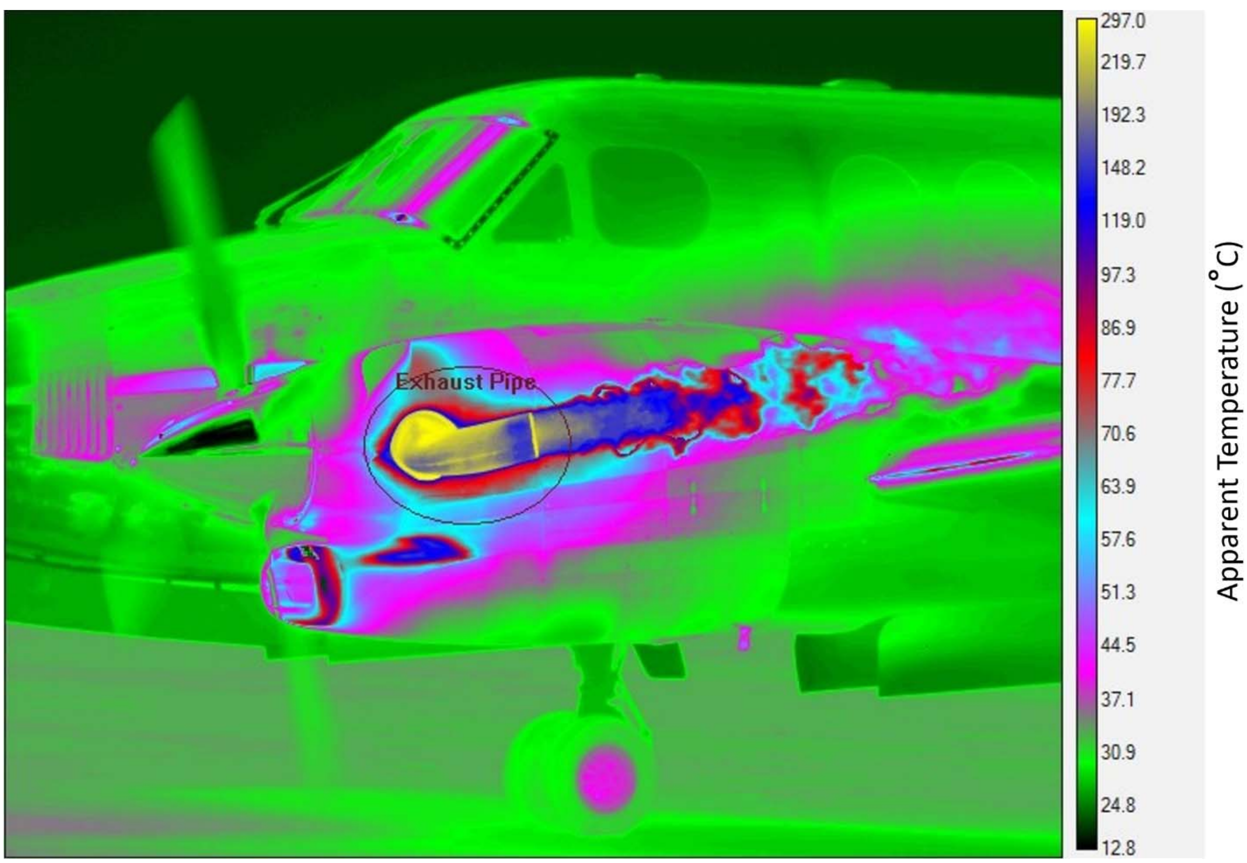

| Statistic [units] | Exhaust Pipe |
|---|---|
| Mean [°C] | 120.2 |
| Std. Dev. [°C] | 74.1 |
| Center [°C] | (279.5, 262.5) 202.5 |
| Maximum [°C] | (253, 248) 297.0 |
| Minimum [°C] | (308, 221) 42.4 |
| Number of Pixels | 9162 |
| Single Pixel Area [cm²] | 0.2704 |
| Area [cm²] | 2477.4 |
| Length [cm] | N/A |
| u Emissivity | ☐ 1 |
| u Distance [m] | ☑ 20.8 |

**Figure 2a-b. Radiometric image of Beechcraft King Air during engine warmup (a) and temperature measurement statistics (b) on the same elliptical region of interest as in Figure 1a.**

## Radiometric Calibration

Radiometric cameras are calibrated in the factory calibration laboratory in order to convert the signals coming out of the imaging sensor into these radiometric quantities. The camera is pointed at a series of laboratory blackbody radiators that are set to precisely controlled temperatures, and the camera signals (in digital count units) are recorded at each temperature. An algorithm described later is used to compute the calibration coefficients which are used to convert from digital counts to radiance. The factory calibrations are stored inside the camera, but it is also possible to create a so-called User Calibration, which generates calibration coefficients. These calibration coefficients are saved in a calibration file which is subsequently made available to apply to the uncalibrated images when viewed in ResearchIR software. Image data can then be expressed in either digital count units (so-called “raw” images) or in radiometric units, but the radiometric calibration is not “baked in” to the stored image data, making it possible to go back later and recalibrate the camera and apply a new calibration if needed.

Imaging radiometry generates a large amount of image data in a short time, especially for high frame rate science-grade cameras, so radiometric infrared camera systems are very much a marriage of the camera to a high-performance PC-based data acquisition system. For the FLIR science-grade cameras, FLIR ResearchIR and the new FLIR Research Studio are the software products used to control the cameras, acquire image data, and perform simple image data analysis. For very high frame rate acquisition, there are special data recorders that use solid-state hard drives, or high-speed RAM buffers on the Niceville cameras themselves.

## Radiometric Cameras

FLIR Systems makes a variety of cameras that are designed for high-performance radiometry. The center for US-made FLIR radiometric science camera design and manufacturing is located in Niceville, Florida. The camera production line includes high frame rate models, high-definition sensor formats, and integration with a variety of optics, including standard lenses, microscopes and long-range IR zoom lenses. The spectral range goes from the visible to ~12µm in wavelength. The performance of these IR cameras is very impressive. One type of camera, the X6900 series, can operate at frame rates of 1000 frames/sec at 640x512 image sizes. A standard sensor array size for a commercial HD midwave or longwave infrared camera is 1280x1024 pixels. High-performance zoom lenses with 1200mm focal length are available, making it possible to image less than 1° fields of view onto a high-definition IR imaging sensor. Microscope optics are also available with 1X and 4X magnifications. A variety of digital data transmission interfaces are supported, including Camera Link and CoaXpress, as well as triggering of recording, synchronization with other cameras or infrared sources, and incorporation of external time signals into the image metadata header.

An example of a FLIR X6900 series camera made in Niceville is shown below in Figure 3. This camera has a cooled midwave indium antimonide sensor that is 640x512 pixels in size. These cameras are typically about the size of a large loaf of bread, weigh between 5-10 pounds, and consume ~20 watts of power. The optics are made of exotic optical materials, which include germanium, silicon and zinc selenide. A SWIR camera is shown in Figure 4. It has a 640x512 pixel InGaAs sensor and can run at up to 180 frames per second. Figure 5 shows the back of the X6900 series camera with all its interfaces. Figure 6 shows the FLIR RS8303, a high-definition midwave range camera with a 120-1200mm zoom lens that is metric, i.e. the focal length and focus position are measured and encoded into the header of the image data output. The lens is thermally compensated, so that focus is preserved even over changes in the lens's ambient temperature.


**Figure 3. FLIR X8500 high-speed, high definition midwave IR science camera with InSb sensor**


**Figure 4. Shortwave infrared camera with InGaAs sensor and InGaAs-band lens**

1. 24 VDC at 2 amps
2. GigE for digital video, control
3. SDI, Composite, HDMI video outs, USB
4. Camera Link Full
5. CoaXpress
6. Aux connector for triggers, analog inputs
7. Sync In and Out
8. IRIG and Genlock


**Figure 5. The back of a FLIR X6900-series camera**


**Figure 6. FLIR RS8300-Series Range Camera on sturdy tripod and adjustable gearhead mount. The InSb sensor is 1344x784 pixels with a 14µm pixel pitch and can run at 124 frames per second.**

The sensors in these cameras are hybrid “sandwiches” composed of a silicon readout IC and an array of IR detectors on a wafer die. The readout ICs, or ROICs, are very complex mixed-signal devices, meaning that they contain analog and digital circuitry working together. Detector technologies include Indium Antimonide (InSb), Indium Gallium Arsenide (InGaAs), and Strained-Layer Superlattice (SLS). Taken together, these three detector types span the wavelength range from 0.6µm to 12µm. InSb is the most commonly used material for high-performance thermal imaging, and InSb cameras are typically set up to image in the 3-5µm band, the so-called midwave IR band. This is accomplished with a "cold filter", a bandpass filter located inside the refrigerated compartment that holds the imaging sensor. Standard InGaAs detector material[2] is responsive between 0.9 and 1.7µm, with special processing available to move the lower wavelength limit down to 0.6µm. SLS is responsive over a wide range of wavelengths, but it is typically used in a camera with a spectral filter operating in the longwave IR band (7.5-10.5µm is typical).

Both indium antimonide and SLS sensors must be cooled to cryogenic temperatures to function properly. This is done with a closed-cycle Stirling cryocooler, a refrigerator that uses helium gas as the "working fluid". The coolers can get the sensors down to 77K or colder and stabilize the sensors at these cold temperatures. In the case of InGaAs, it is not necessary to cool the sensor to cryogenic temperatures. InGaAs can be run at ambient temperature values (around 30 °C) and still operate with low dark current, especially at short integration times. It is desirable to stabilize the operating temperature, though, so that dark current is stable, otherwise the non-uniformity correction will not work well at longer integration times as the sensor warms or cools in response to the camera's ambient operating temperature.

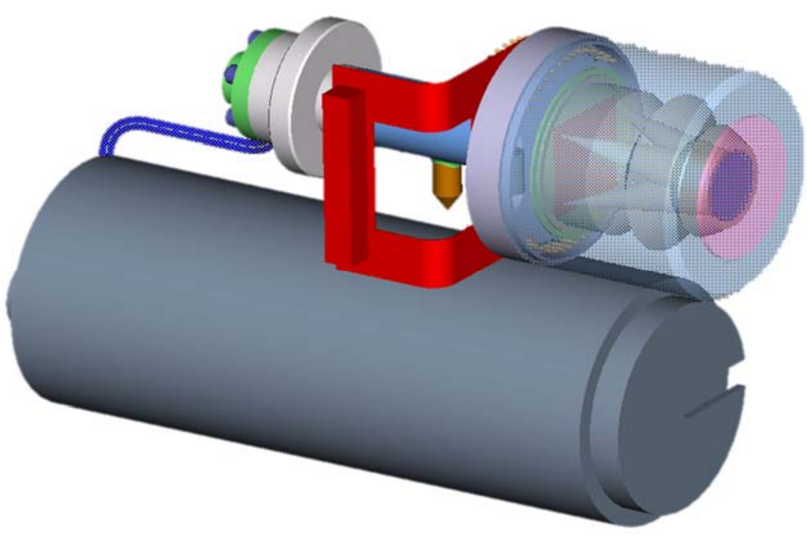

**Figure 7. Linear Stirling cycle cryocooler and transparent view of dewar assembly**

These three common IR camera operational wavebands were chosen to correspond to terrestrial atmospheric transmission windows in the infrared band, as shown in Figure 8 below. These are

[2] Standard InGaAs alloy has a lattice constant that is matched to the indium phosphide substrate upon which the InGaAs material is grown.

regions of the spectrum where atmospheric gases do not absorb so strongly that imaging at appreciable distances is impossible. Most of the absorption features are due to water vapor, with the strongest absorption feature located between 5 and 7.5µm. Water vapor molecules are extremely well coupled to infrared radiation! An infrared camera operating in this band may only be able to see ambient-temperature terrain out to a range of 100m or less.

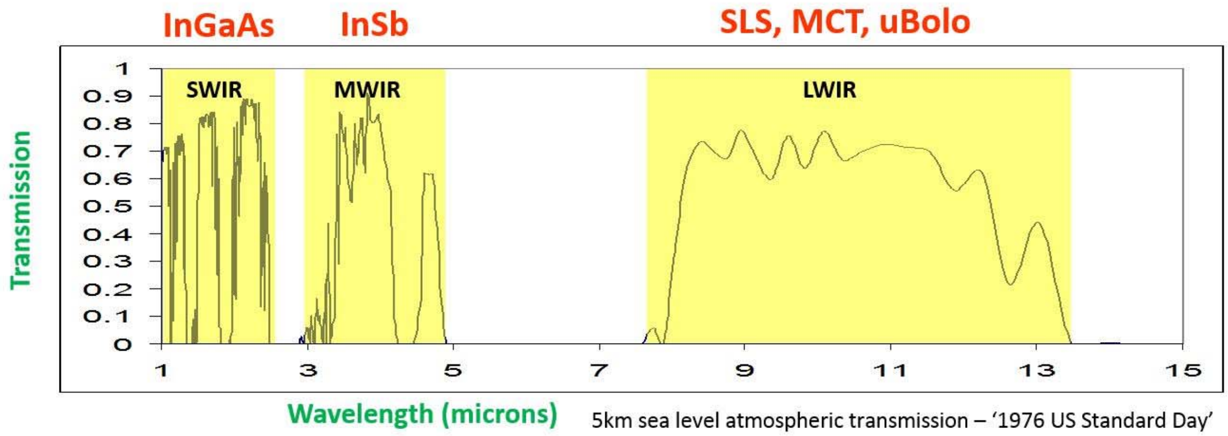


**Figure 8. Atmospheric transmission of typical horizontal air path with wavebands and typical detector types for each band**

Figure 9a-c shows three views of an outdoor scene. The right-hand image was taken with a very special QWIP camera that imaged right in the middle of the 5-8µm absorption band. The air looks very murky!

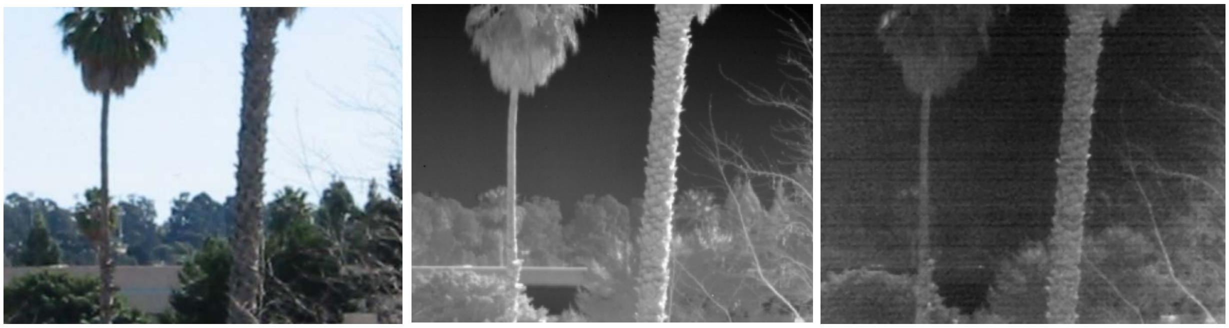

**Figure 9a-c. Visible light (a), 8-9µm (b) and 5.7-6.7µm (c) images of an outdoor scene. The atmosphere is very opaque around 6.2µm due to absorption by water vapor**

Another very strongly absorbing atmospheric gas is carbon dioxide, which has its primary absorption feature at 4.3µm. The 3-5µm band is a very commonly used band for general-purpose thermal imaging. For long-range imaging, special cold filters are used in midwave cameras to notch out the 4.2-4.4µm band which can sometimes improve contrast at multi-km ranges.

## Optics

FLIR offers a large selection of optics for the Niceville cameras.  The standard interchangeable lenses come in focal lengths ranging from 13mm to 200mm and are manually adjusted for focus. There are interchangeable microscope lenses for midwave and longwave imaging of small objects.  There is also a family of "range cameras" with very long variable field of view lenses which are motorized and instrumented to record their settings.  The long-range optics are integrated with several different camera types into rugged environmental housings with sunshades and optional motorized lens covers for outdoor use, as pictured above in Figure 6. Infrared camera optics are made of exotic IR-transmitting optical materials in low volumes, and as such they are much more expensive than visible-light lenses.

## Standard Focal Lengths

The standard focal lengths for both 3-5µm  midwave IR and longwave IR Niceville cameras are:

**17mm, 25mm, 50mm, 100mm, 200mm, 1X microscope**

Two examples of these lenses are shown in Figure 10.  Additionally, there is a 4X microscope lens for midwave cameras, but not for longwave.  A 4X magnification makes little sense in the longwave IR band, as there will be far too much diffraction blurring to make use of the magnification.  The fields of view for these standard lenses on the standard FLIR A6701 camera (640x512 15µm pixels) are as follows:

| **Focal length** | A6701 HFoV (deg.) | A6701 VFoV (deg.) | X6901sc HFoV (deg.) | X6901sc VFoV (deg.) | X8501sc HFoV (deg.) | X8501sc VFoV (deg.) |
|---|---|---|---|---|---|---|
| 17mm | 32 | 26 | 50 | 41 | 49 | 40 |
| 25mm | 22 | 18 | 36 | 29 | 34 | 28 |
| 50mm | 11 | 8.8 | 18 | 15 | 18 | 14 |
| 100mm | 5.5 | 4.4 | 9.1 | 7.3 | 8.8 | 7.0 |
| 200mm | 2.7 | 2.2 | 4.6 | 3.7 | 4.4 | 3.5 |

**Table 1.  Fields of view for different Niceville camera and lens combinations**

The microscope fields of view are given in millimeters in object space, since the working distance is fixed at 50.8mm for the 1X and 38.1mm for the 4X.

| Magnification | A6701 HFoV (mm) | A6701 VFoV (mm) | X6901sc HFoV (mm) | X6901sc VFoV (mm) | X8501sc HFoV (mm) | X8501sc VFoV (mm) |
|---|---|---|---|---|---|---|
| 1X | 9.6 | 7.7 | 16 | 13 | 15 | 12 |
| 4X | 2.4 | 1.9 | 4.0 | 3.2 | 3.8 | 3.1 |

**Table 2. Fields of view for different Niceville camera and microscope optics combinations**

Broadband optics are designed for midwave cameras that have response outside of the 3-5µm standard midwave band. The focal lengths currently available are 25mm, 50mm and 100mm. These optics are designed to work in the 1-5µm band. There are no broadband microscope lenses currently available.

These standard lenses listed above are designed to accept filter holders that thread into the back of the lens. FLIR sells ND1, ND2 and ND3 filters already mounted in these holders. ND stands for neutral density, which means that the filter has a constant transmission as a function of wavelength over the bandpass of the camera. An ND1 filter has a transmission of 0.1, or 10%. The transmission for an ND2 filter is 1%, and 0.1% for an ND3. ND filters are called gray filters in Europe. There are also blank filter holders available that can hold a 1.5-inch diameter filter. This is not a standard filter diameter, and bandpass filters of this size are much more expensive than 1-inch diameter filters. A better solution is to use the models of Niceville science cameras that are equipped with filter wheels. The filters in the wheel are held very close to the warm window on the camera's sensor housing, and thus they can be 1 inch in diameter without vignetting the scene. These filters can be user-supplied, or they can be one of various types sold by FLIR for different applications.

**Figure 10a-b. Renderings of FLIR midwave 25mm (a) and 200mm (b) science camera lenses**

## InGaAs and VisGaAs Optics:

There are currently 5 different focal lengths of lenses available in FLIR inventory for Niceville InGaAs cameras, which have spectral response in the 0.9-1.7µm band. The focal lengths are:

**16mm f/1.4, 25mm f/1.4, 35mm f/1.4, 50mm f/1.8, 100mm f/1.5**

All these lenses have adjustable apertures that can be set from f/1.4 or whatever the lowest value for the lens is, to f/16 and then to a closed position. There are two lenses currently available for Niceville VisGaAs cameras; they are a 25mm, f/2 unit and a 100mm, f/1.5 unit. FLIR VisGaAs sensors have a spectral response in the 0.4-1.7µm range. The "Vis" stands for visible light. InGaAs is processed to thin down the indium phosphide substrate, which allows more visible light to reach the InGaAs detectors.

**Figure 11a-b. Renderings of 100mm InGaAs (a) and 25mm VisGaAs (b) lenses available from FLIR**

## Non-Standard Optics

It is possible to purchase non-standard lenses for Niceville cameras. Vendors such as StingRay Optics have designed lenses for midwave cameras that can be configured with the correct relay optics and lens mount to attach to Niceville science cameras. They have 500mm and 750mm catadioptric units, both at f/2.5. As a rule, FLIR does not release the opto-mechanical interface drawings to anyone except a select group of US-based optics vendors, so customers wanting custom lens designs are strongly encouraged to work with these vendors to obtain the best performance match between the optics and the FLIR science camera on which they will be used. Contact FLIR customer support for more information.

StingRay also offers InGaAs and VisGaAs lenses in various focal lengths such as 200mm and 500mm, and several semi-custom midwave and InGaAs-band zoom lens options.

# Radiometric Calibration

There are four different types of waveband range for standard-product Niceville cameras that are calibrated as part of a standard factory calibration service product. The bands are 3-5µm InSb, 1.5-5.7µm broadband InSb, longwave SLS and 0.9-1.7µm InGaAs. The factory calibrations each cover their own predetermined scene temperature range. The temperatures specified in the range are the minimum and maximum blackbody temperatures that can be measured to the stated accuracy specifications, which are typically +/- 2 °C or 2% of the temperature value in Celsius units, whichever is greater. The accuracy of every calibration is evaluated with laboratory blackbodies, which are a close approximation to theoretically perfect blackbodies. Real-world scenes and objects can depart significantly from theoretical blackbody performance, and it is imperative that the so-called “object parameters” in ResearchIR are set correctly by the camera system user to get the maximum temperature measurement accuracy.

The calibration ranges are each comprised of multiple sub-ranges, each with a particular sensor integration time setting. It is not possible to image from -20 °C to 350 °C with a single integration time with a 3-5µm camera, because the change in the in-band radiance for objects at those temperatures is far too high. A 3-5µm camera operating at an integration time short enough to reach 350 °C will have an extremely noisy image at scene temperatures below 100 °C. A 3-5µm camera operating with an integration time long enough image down to -20 °C will saturate at temperatures above 60 °C. Longwave cameras can (usually) measure temperatures across a much wider scene temperature range at a given integration time, due to blackbody physics which results in a lower change in radiance for the same change in temperature, at least in the most common temperature measurement ranges.

The FLIR Niceville facility is set up to perform standard calibrations from -20 °C to 3000 °C. Lower temperature calibrations are possible, but they are done on a case-by-case basis only and are not trivial to create or to use in practice. When a target is at -50 °C, the parasitic radiation in the camera optics themselves can introduce significant measurement errors. Parasitic radiation (also known as self-radiation) is infrared radiation that is emitted by surfaces inside the camera optical path. This unwanted IR radiation can reach the detectors in the sensor and adds to the image noise. Another issue is that a calibration source at -50 °C will immediately get covered in water frost in ambient air, requiring that the target-to-camera air path be immersed in dry nitrogen or be placed in a vacuum chamber.

## Calibration Ranges

The following calibrations are the standard ranges which do not require an ND (neutral density) filter to limit scene energy:

- InSb 3-5µm : -20 to 350 °C with no ND filter, -10 to 350 °C for microscope lenses
- InSb Broadband : -20 to 300 °C with no ND filter
- Longwave SLS: -20 to 650 °C with no ND filter

Calibrations for higher scene temperatures require the use of an ND filter. Here are those ranges:

**ND1 filters:**

- InGaAs 0.9-1.7µm : 400 °C to 1200 °C with ND1 filter
- InSb 3-5µm : 45 to 600 °C with ND1 filter, 45 to 600 °C for microscopes with ND1 filter

- Longwave SLS: 250 to 1500 °C with ND1 filter and 250 to 2000 °C with ND1 filter (extended range)

**ND2 filters:**

- InGaAs 0.9-1.7µm : 600 to 1500 °C with ND2 filter
- InSb 3-5µm : 250-1500 °C or 2000 °C (extended range) with ND2 filter
- InSb Broadband : 250-1500 °C or 2000 °C (extended range) with ND2 filter
- Longwave SLS: 500 to 3000 °C with ND2 filter

**ND3 Filters:**

- InGaAs 0.9-1.7µm : 1100 to 2200 °C with ND3 filter
- InSb 3-5µm : 500-3000 °C with ND3 filter
- InSb Broadband: 500-3000 °C with ND3 filter

There is no Longwave SLS calibration with an ND3, because we don't have blackbody sources that can go higher than 3000 °C. In principal, an ND3 filter would allow one to measure up to ~25,000 °C with a Longwave SLS camera. That is even hotter than Niceville in late July.

## Radiance Ranges

For radiometric IR camera users that are measuring radiance and radiant intensity, it is useful to know what equivalent radiance values correspond to the calibration temperature limits. We calibrate the cameras using cavity and area blackbodies that are controlled by precision temperature controllers, which is why we use temperature units to specify the limits and the accuracy of the factory calibration, but the calibrated cameras can output either temperature or radiance values. We do not certify the accuracy of radiance measurements, only the accuracy of temperature measurements on laboratory-grade blackbody targets.

Radiant intensity measurement values are derived from radiance measurements by multiplying the radiance of each pixel by its emitting area. In ResearchIR software, the units of radiance are watts/square cm/steradian, and the units of radiant intensity are watts/steradian. Radiant intensity is referred to as Intensity in the statistics window in the software. The derivation of these units and their meaning is discussed later in this document.

Here are the temperature and the corresponding in-band radiance values for the FLIR Niceville camera standard calibrations:

| Description | Temperature (°C) | 3-5µm Radiance (W/sq. cm/sr) |
|---|---|---|
| No Filter | -20 | 2.5e-5 |
| | 350 | 6.8e-2 |

| | | |
|---|---|---|
| Microscope Lenses | -10 | 4.1e-5 |
| | 350 | 6.8e-2 |
| ND1 filter | 45 | 3.5e-4 |
| | 600 | 3.7e-1 |
| ND2 filter | 250 | 2.3e-2 |
| | 1500 | 3.8 |
| ND2 filter Extended Range | 250 | 2.3e-2 |
| | 2000 | 6.7 |
| ND3 filter | 500 | 2.1e-1 |
| | 3000 | 14 |

**Table 3. MWIR (3-5µm) InSb camera calibrations**

| **Description** | **Temperature (°C)** | **7.5-10.5µm SLS Radiance (W/sq. cm/sr)** |
|---|---|---|
| No Filter | -20 | 1.1e-3 |
| | 650 | 1.4e-1 |
| ND1 filter | 250 | 3.0e-2 |
| | 1500 | 4.4e-1 |
| ND1 filter Extended | 250 | 3.0e-2 |
| | 2000 | 6.4e-1 |
| ND2 filter | 500 | 9.1e-2 |
| | 3000 | 1.0 |

**Table 4. LWIR (7.5-10.5µm) SLS camera calibration**

| Description | Temperature (°C) | 0.95-1.66µm InGaAs Radiance (W/sq. cm/sr) |
|---|---|---|
| ND1 Filter | 400 | 3.9e-4 |
| | 1200 | 1.2 |
| ND2 filter | 600 | 1.0e-2 |
| | 1500 | 4.2 |
| ND3 Filter | 1100 | 7.1e-1 |
| | 2200 | 25 |

**Table 5. SWIR (0.95-1.66µm) InGaAs camera calibration**

## Exposure Control

These calibration ranges are determined by the practical range of integration times for a particular camera. Take the case of the standard InSb calibration for a 3-5µm camera that ranges from -20 to 350 °C. The high-end temperature of 350 °C is achieved by running the sensor with an integration time of about 10-15 microseconds for an f/2.5 camera. Typical FLIR focal plane arrays are not designed to operate below 10 microseconds without suffering from degraded performance, notably a reduced digital count range where the sensor is linear. This means that the minimum integration time is limited to about 10 microseconds or longer. This limit sets the upper end of the standard 3-5µm InSb camera range[3]. If the camera is set to a very short integration time, for example 1 microsecond, then the maximum scene temperature will be higher, but the linear range of the digital counts may max out at about 8000 counts, for example.

At the other end of the calibration temperature range, the -20 °C minimum measurable temperature was determined through experimentation. Lower scene temperatures will result in very noisy images. The integration time cannot be increased beyond the limit where the integration capacitors fill up just from parasitic radiation from the ambient temperature optics. At -20 °C temperature, the 3-5µm target radiance is only 2.5e-5 watts/sq. cm/sr. The sensor is "seeing" both this small radiance as well as the parasitic radiation from the optics and camera interior, which may be 35 °C or hotter. The parasitic radiation depends on the temperature of the optics and the camera interior, both of which can and will change during operation, even in a laboratory with controlled air temperature. If a camera is turned on and left running for several hours, the optics temperature can change substantially because of the camera electronics power dissipation. If the camera is operated outdoors in bright sunlight with no sunshade, then both the optics and camera interior temperatures can get quite high, and the converse is true if the camera

[3] For an f/4 camera, the maximum temperature can be higher, but this is considered a custom calibration.

is operated in cold weather conditions. When the target temperature is down around -20 °C, the measurement can be quite susceptible to changes in parasitic radiation in the optics and camera body. We deal with this with a technique known as $T_{drift}$ correction.

## $T_{Drift}$ Correction

Science cameras are often used in the field, and the cameras and their optics can be subjected to heating and cooling from solar loading and ambient air temperature changes. Niceville science cameras incorporate a technique known as $T_{drift}$ correction to deal with the uncontrolled levels of parasitic radiation that result. Each camera is placed in a thermal chamber, and it looks out of the chamber at a fixed-temperature blackbody radiator. The chamber interior is then cycled through a series of high and low temperatures within the operation range of the camera, and the response of the camera to the blackbody is measured to form a set of permanently stored correction parameters that are continuously applied to the digital image data. The proprietary technique helps the camera system maintain radiometric accuracy during the occasionally large ambient temperature excursions that are encountered in field operations. For a target at -20 °C, the $T_{drift}$ compensation is pushed to its limit, since the parasitic radiation is higher than the target radiation! If we turn up the integration time more to get better sensitivity, we end up filling the sensor with more and more parasitic radiation until it is full – so there is a very real sensitivity limit imposed by the presence of optics and camera interior self-radiation. Measuring very cold targets with all this excess radiation present in the system is like trying to measure the weight of a walnut by weighing a person holding the walnut, then the person by themselves, then subtracting the two values to find the weight of the walnut! That method of weighing a walnut is very prone to error. But there is no good way to reduce the parasitic radiation without cooling the whole camera and lens combination, which carries its own problems, notably moisture condensation. The best practice is to do the calibration at a known camera and optics temperature, then do the measurements with the camera and optics at that same temperature. This would obviate the need for $T_{drift}$ correction. Another technique is to image the target of interest while simultaneously imaging a laboratory blackbody source at a known temperature. A radiance offset correction can be derived from the apparent radiance of the source. If two blackbodies are used at two different temperatures, the atmospheric transmission effects can be calculated and accounted for.

## Integration Times and Subranges

All the standard FLIR Niceville radiometric calibrations use a set of sub-ranges to cover the total temperature span of the calibration. These subranges correspond to preset integration time values and their associated non-uniformity corrections[4]. The cameras can have up to four "presets" or sub-ranges loaded at any one time. Each preset consists of a non-uniformity correction with all its associated sensor settings, such as integration time, as well as the radiometric calibration file. The cameras can be repeatedly cycled through the presets in a mode

[4] Non-uniformity corrections, or NUCs are described in detail in the next chapter.

known as “superframing” mode.  The images taken at each preset can be combined in the software to produce high dynamic range images in either radiance or temperature units.  The details of superframing will be covered later in this book.  The standard ROICs used with InSb or SLS detectors are typically only able to give good radiometric performance down to 10 microseconds or greater integration time.  This drives the requirement to control the exposure with neutral density filters in order to measure hotter temperatures.

For a typical 3-5µm, f/2.5 InSb camera like the A6701sc, the integration times and subranges that go into spanning the standard factory calibration range of -20 to 350 °C are as follows:

- -20 to 55 °C: 2.4ms
- 35-150 °C:    0.29ms
- 80C-200 °C: 0.12ms
- 150-350 °C: 0.021ms

We also add a 10-90 °C subrange in the set of standard calibration sub-ranges – it is a very useful range for many applications.

These values were determined by averaging the values for a set of almost fifty as-built A6701sc production cameras with 50mm lenses that were calibrated in the Niceville factory. Different lenses tend to give quite similar values, depending on the lens transmission which varies between lens types because of the different number of optical elements in each design.

Note the integration time changes by a factor of ~100 between the hottest and coldest subranges. The camera system will have to be in preset cycling mode (superframing) in order to span the full range of the calibration during a measurement.  One can then image a hot target against a cold background and get meaningful radiometry on both the target and background simultaneously.  As always, there are caveats, since a very hot target will tend to “paint” radiation all over the image due to scattering and reflections in the optics.  That will make cooler parts of the scene look hotter than they would look if the hot target were not present.

### Coldshield f/number

The coldshield is a radiation shield used in cryogenically cooled cameras to prevent stray radiation from striking the focal plane array.  The coldshield is intended to only allow IR radiation that comes through the lens from the scene to strike the sensor.  It does this by means of an aperture in the end of the coldshield called the coldstop.  In Figure 12 below, the coldshield looks like a Christmas tree in profile. This shape and the black painted interior are designed to reduce stray light from reflecting off the sides of the coldshield into the sensor.  The red cone shows the solid angle of the coldstop as viewed by a pixel at the center of the FPA.

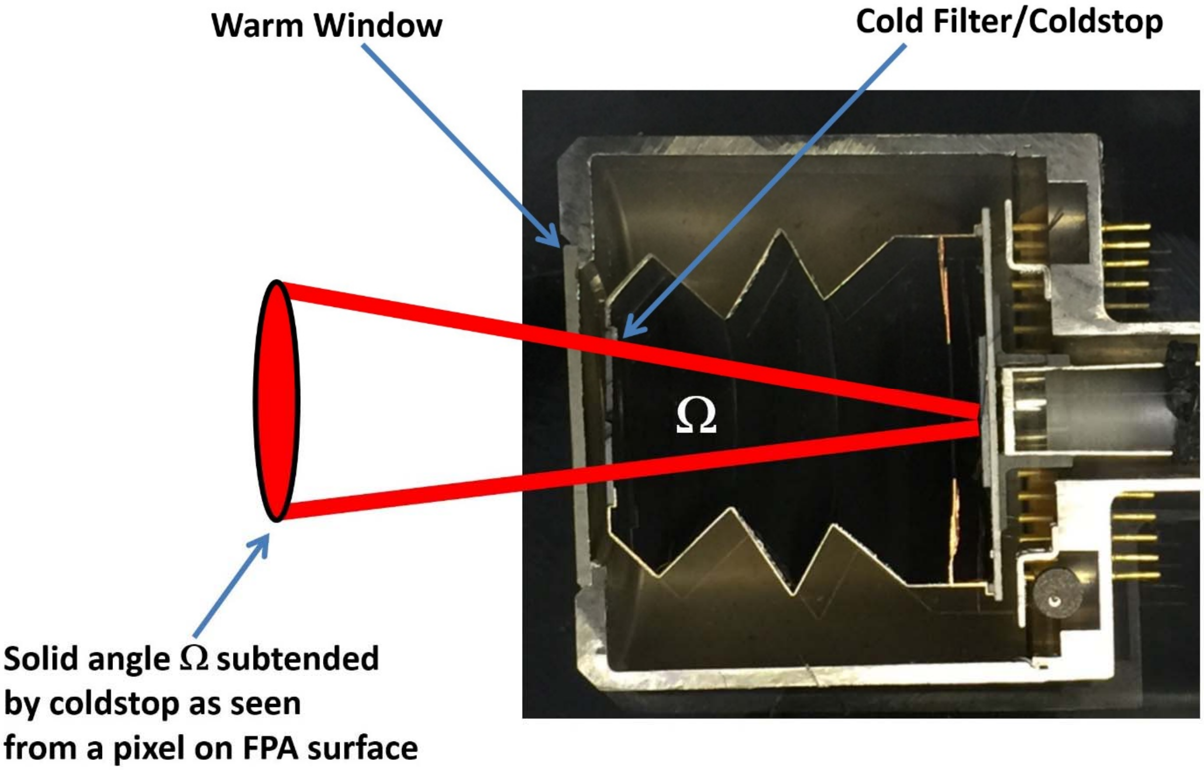


**Figure 12. Cutaway of dewar with coldshield and focal plane array**

The exposure is controlled by the f/number of the coldshield. The higher the f/number, the smaller the solid angle of the coldstop as seen from the FPA center, and the less IR radiation will strike the sensor for a given scene radiance. The f/number is defined as the distance of the focal plane array from the coldstop divided by the diameter of the circular coldstop aperture. The standard Niceville cooled cameras have f/2.5 coldshields. Some of the cameras are built with f/4.1 coldshields. These are used with the microscope lenses, which have historically been f/4. The other reason to use higher f/number coldshields is to be compatible with long focal length lenses. In the case of the RS8303 camera, the 120-1200mm lens is f/5, which requires a special f/5 coldshield to match the optics.

For a given integration time, the f/4 camera can view a higher temperature target without saturation as compared to an f/2.5 camera. It is possible to calibrate f/4 coldshield 3-5µm InSb cameras to higher temperatures than 350 °C, but it is not standard practice. For a 150-350 °C calibration on an f/4, 3-5µm InSb camera, the average integration time (for a sample of 45 as-built cameras) is 52µs with a 1µs standard deviation. This time is 2.5 times longer than that used for the standard f/2.5 camera 150-350 °C calibration sub-range described earlier. That suggests that the f/4 camera, if calibrated at 20µs, should be able to go up to a peak radiance 2.5 times higher than the 350 °C equivalent in-band radiance. This new temperature value will be around 460 °C.

## Neutral Density Filters

The calibration ranges can also be controlled with neutral density (ND) filters. If an ND1 filter is added to the system, then the radiance range high and low limits are increased by a factor of 10. An ND2 filter will increase the radiance limits by a factor of 100 and so on. Note that these factor of 10 and 100 increases are in radiance units only, NOT temperature units. In other words, an ND1 filter put on a camera that measures to 100 °C does NOT now measure to 1000 °C. The relationship between radiance and temperature is not linear, except in the limit of very small temperature ranges.

## Spectral Ranges of Different Detectors

The standard calibration ranges are for cameras with 3-5µm InSb sensors, 1.5-5.5µm Broadband InSb, 7.5-11µm Longwave SLS and 0.9-1.7µm InGaAs sensors. The temperature ranges for these standard calibrations are quite different because the relationship between in-band radiance and scene temperature vary dramatically between shortwave, midwave and longwave infrared. This difference is due to blackbody physics: at longer wavelengths, there is a smaller increase in radiance for a given temperature increase. The standard calibration for a 3-5µm camera runs from -20 °C to 350 °C. This is achieved with four subranges spanning integration times from 20 to 2400µs, the maximum practical span for these cameras. The same standard calibration for the longwave IR SLS camera runs from -20 °C to 650 °C . This is accomplished with only three subranges, with typical integration times of 67µs, 21µs, and 8µs. Here the total temperature measurement range is much larger than for the 3-5µm standard calibration. The entire range can be spanned with only three subranges, instead of four, and the integration time varies over a factor of ~10 versus a factor of ~100 for the 3-5µm InSb camera calibration. InGaAs camera calibrations start at 400 °C and go up from there. The available InGaAs-band radiance below 400 °C drops off steeply and makes it problematic to make accurate measurements below 400 °C without resorting to lengthy integration times.

## Spectral Filters

For a given detector type, like InSb, the calibration ranges will be affected by the choice of spectral filter used in the camera. For example, a custom InSb camera equipped with a SWIR band cold filter from 1.5-2.5µm will have very, very different dynamic range performance characteristics relative to a standard 3-5µm camera. For scene temperatures ranging from 25 °C to 550 °C , the in-band radiance of the 3-5µm camera changes by a factor of ~1300. For an InSb camera with a 1.5-2.5µm cold filter, however, the in-band radiance changes by a factor of **560,000** over that same temperature span! Covering this whole range would require a large number of preset subranges spanning very large changes in integration time, plus the use of neutral density filters.

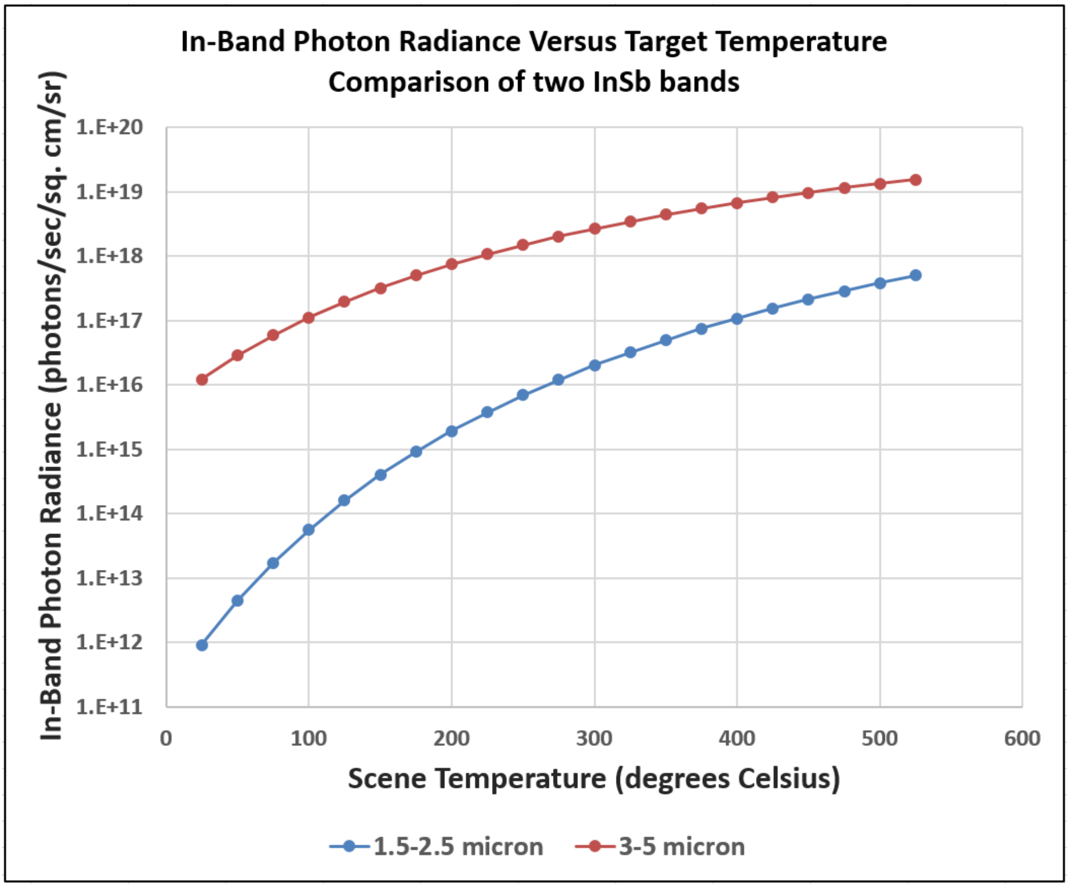


**Figure 13. Comparison of photon radiance versus scene temperature between two InSb spectral bands**

## Dewars and Coolers

InGaAs sensors in the Niceville A62xxsc series are typically operated at a temperature of 30 °C , and they are temperature stabilized at 30 °C using a small thermoelectric cooler. InSb and SLS sensors must be cooled to cryogenic temperatures in order to work properly. The sensor arrays are mounted inside a double-hulled vacuum chamber called a dewar. The sensor itself is surrounded by the coldshield, with an aperture on the end called the coldstop, which is where the cold filter is mounted (if one is used). The coldshield and sensor are cooled using a closed-cycle Stirling cryocooler, as mentioned earlier. The detector, dewar and cooler assembly are called DDCAs, as shown in Figure 14a. FLIR no longer manufactures science cameras with liquid nitrogen pour-filed dewars. A section view of a dewar in Figure 14b shows the ray cone extending out of the sensor center through the coldstop.

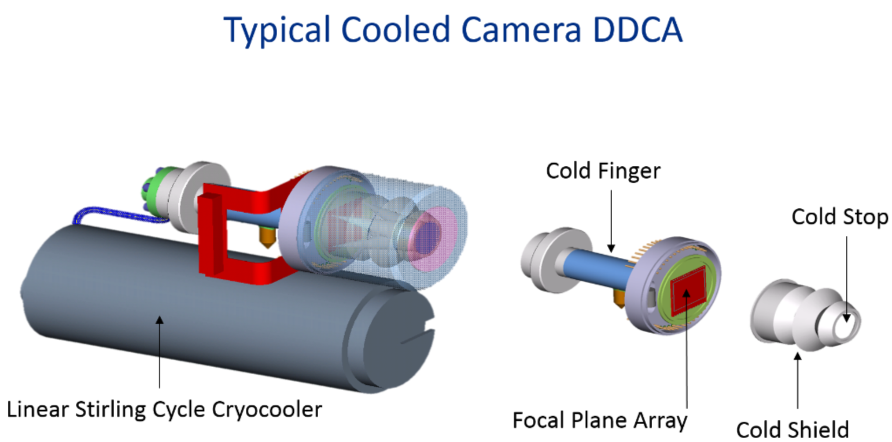


**Figure 14a. DDCA and exploded views**

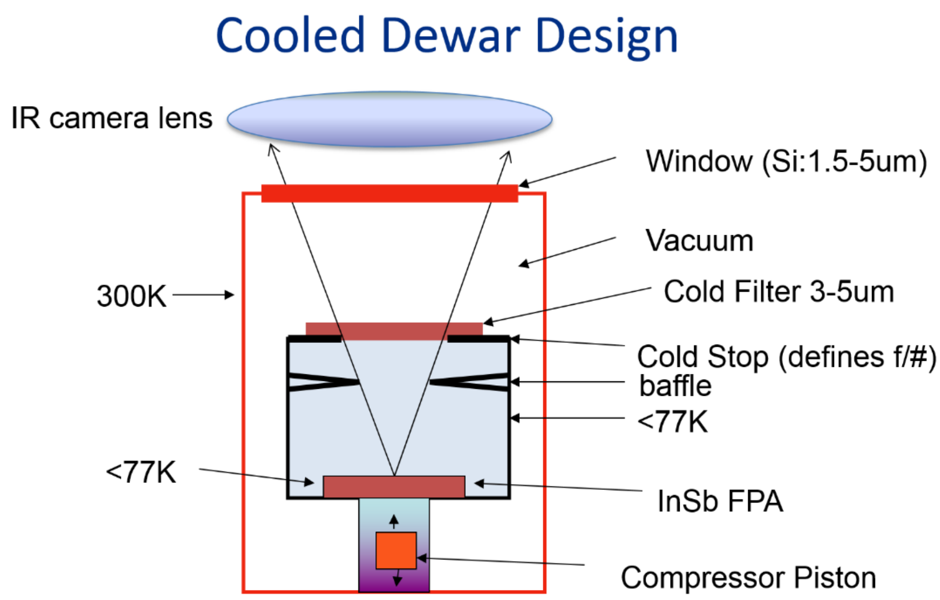


**Figure 14b. Dewar cross section cartoon**

## Infrared Imaging Radiometry

Some examples of infrared imaging radiometry measurements include:

1. Measurement of the radiance profile of a jet aircraft exhaust plume
2. Measurement of the total radiant intensity of an aircraft engine
3. Measuring the temperature of an animal's legs to look for inflammation

Figure 15a is an image of a jet's exhaust plume on takeoff. The color palette is black hot, which was used to enhance details in the plume. The plume radiance as a function of longitudinal position (along the red line profile) is shown in Figure 15b. This plot was made using the Profile tool in ResearchIR, and the red line in Figure 15a is the line used for the profile. The green dot indicates the starting pixel for the profile plot.

**Figure 15a. Jet exhaust plume during a full power takeoff**

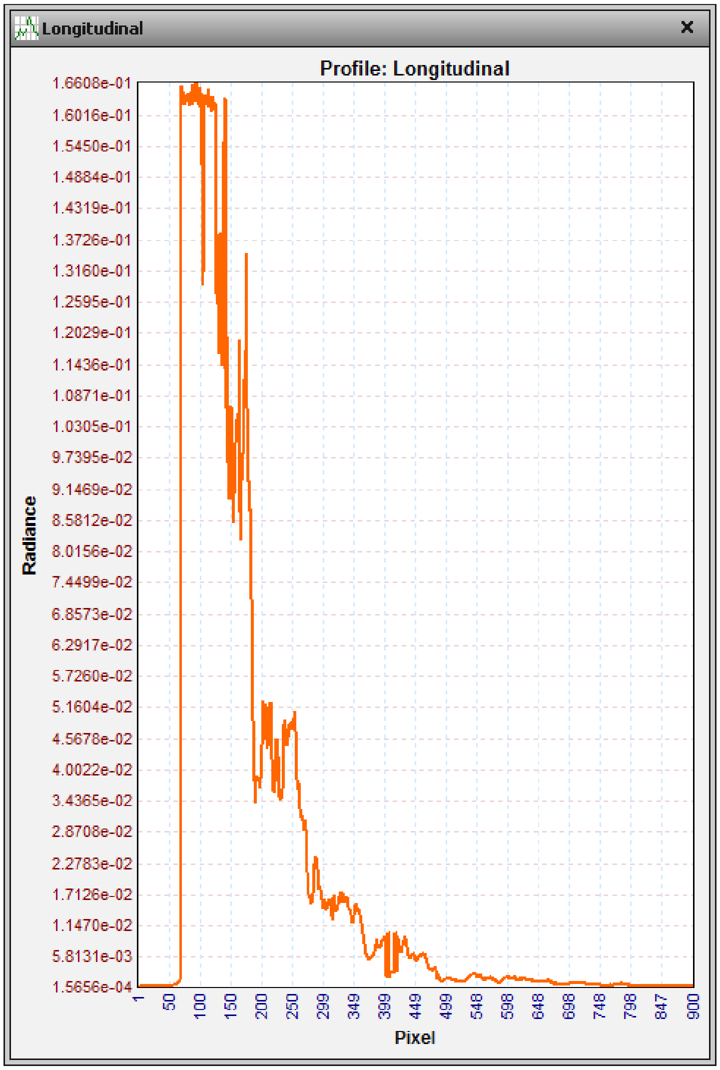


**Figure 15b. Radiance as a function of position along the jet exhaust plume**

Figure 16 shows the radiant intensity of a region of interest around a propeller plane's engine exhaust at various times during the engine start process, along with the radiant intensity versus time curve extracted from a 263-frame image sequence. The red arrows indicate the approximate times that these four images were acquired.

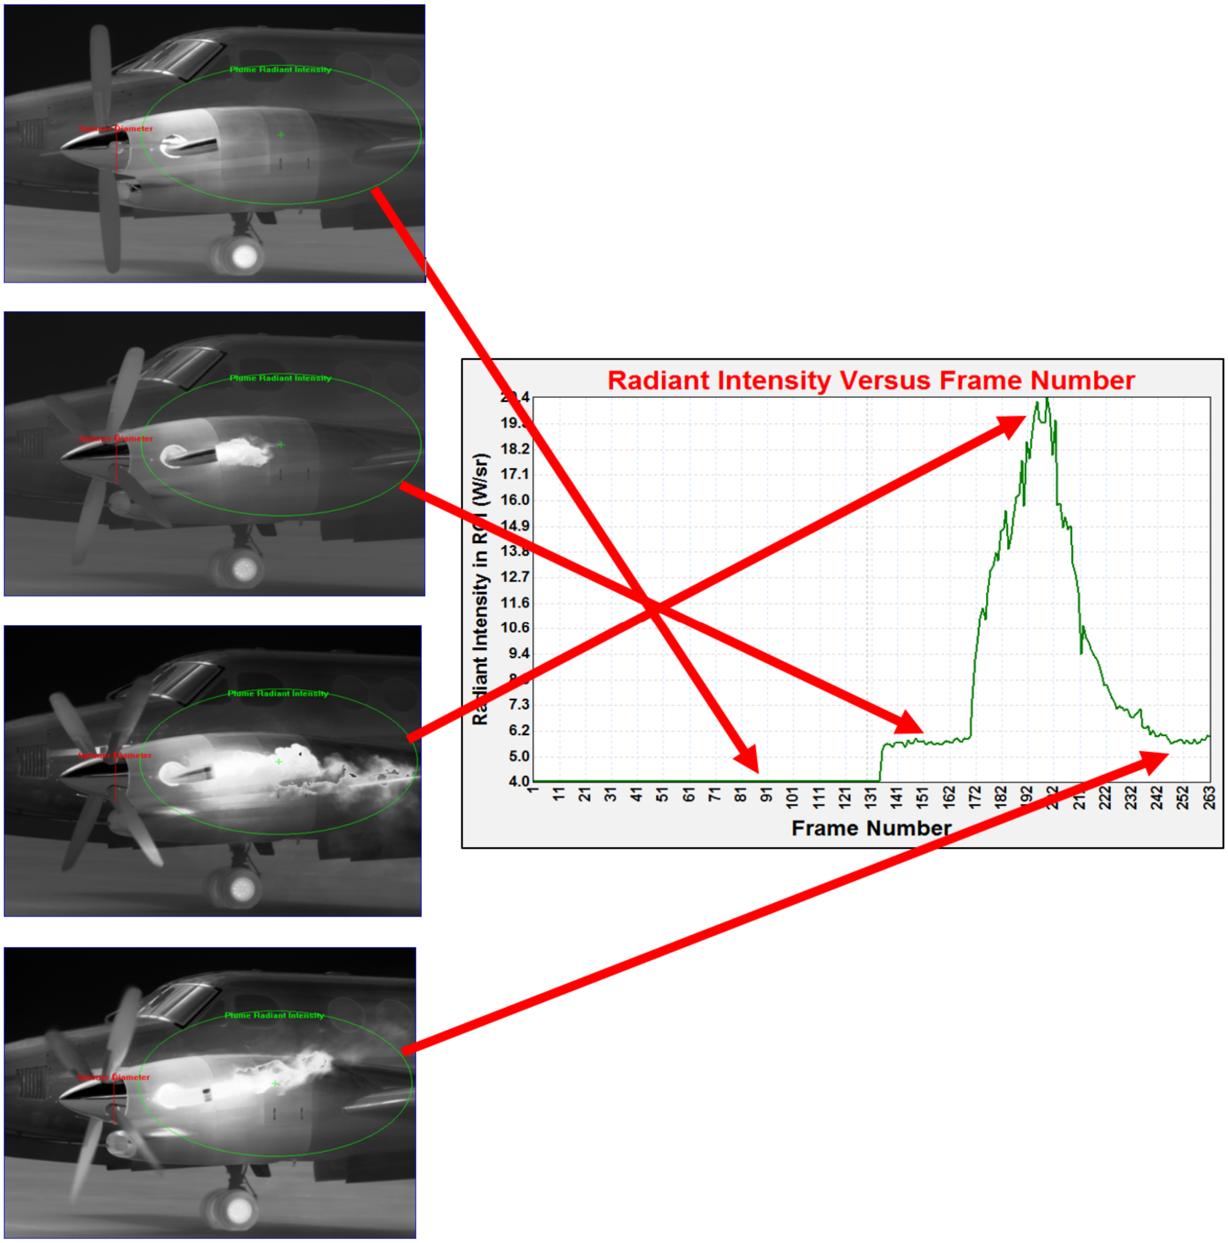


**Figure 16. Four images of a Beechcraft King Air engine start with radiant intensity within a region of interest.**

Figure 17 shows an image of an elderly elephant with an inflamed and arthritic knee joint. There is a noticeable temperature rise on the hide.

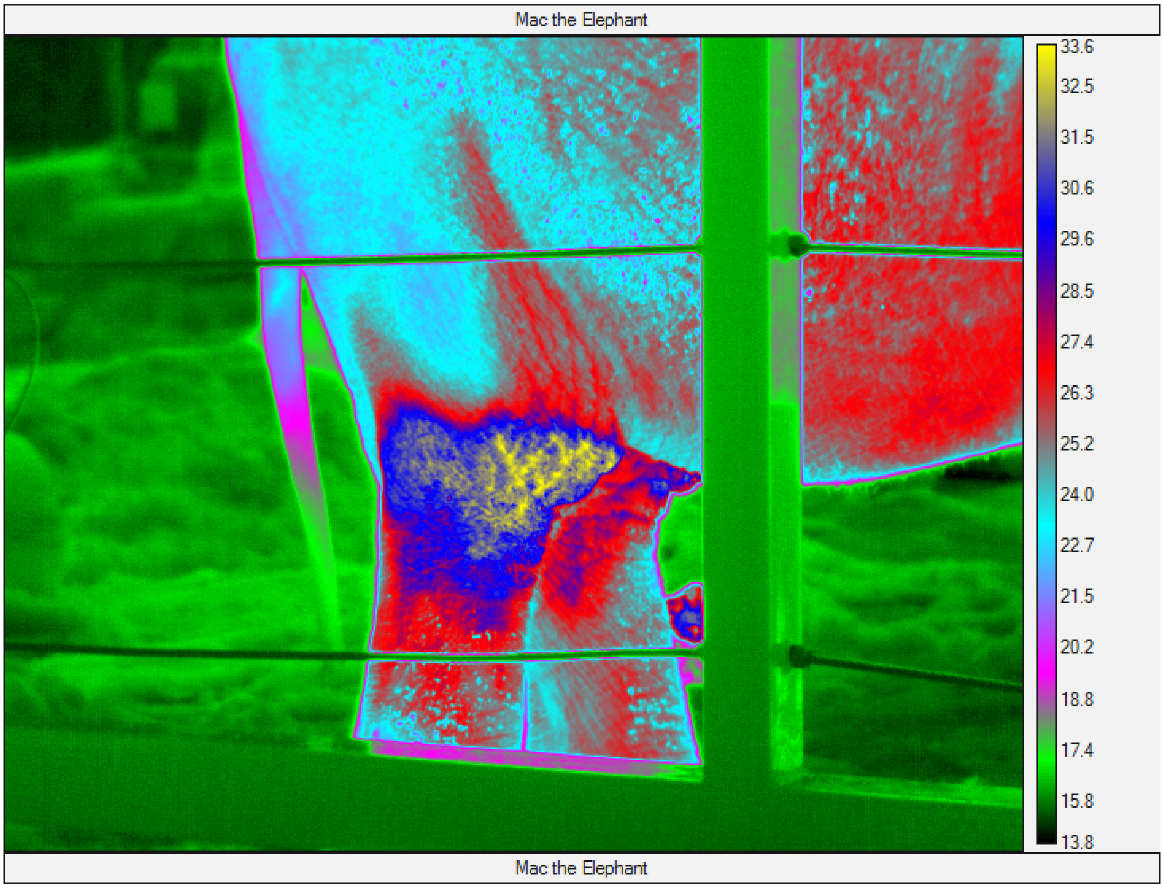


**Figure 17. Temperature map of an elephant's right rear knee joint. Unit are in degrees Celsius.**

Figure 18 shows a midwave IR image of the Oxnard Airport about 1 hour after sunset. There are thermal shadows (also known as "thermal scars") where aircraft had been positioned earlier on the tarmac, shading it from the Southern California sun high overhead. Repeated measurements of the shadow's contrast with the tarmac can be used to determine when the plane moved away.

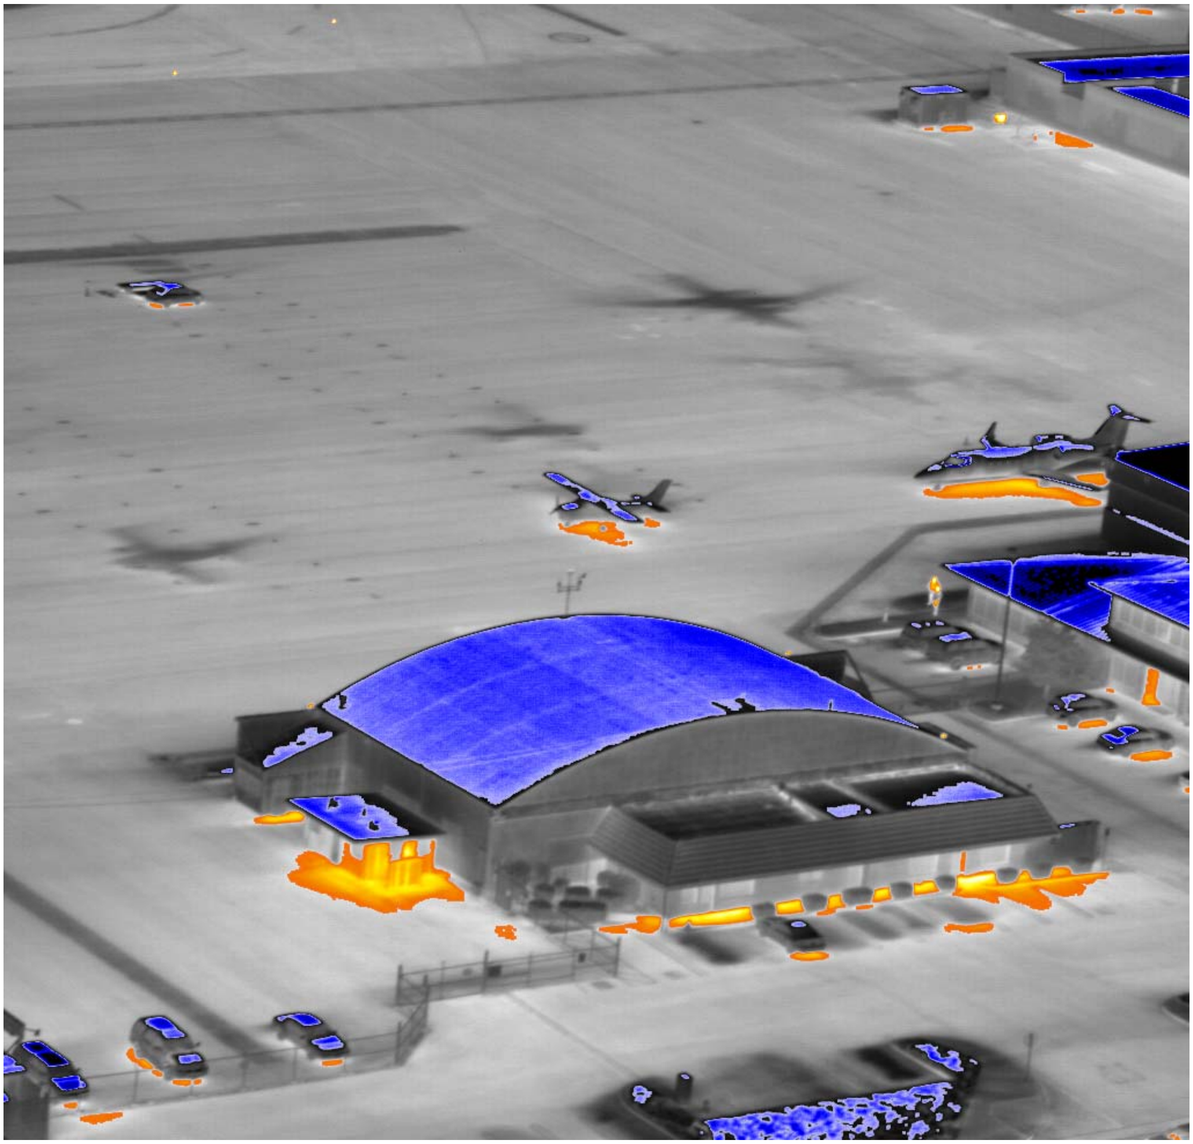

**Figure 18. Aerial image of the Oxnard Airport about 1 hour after sunset. The orange areas are the hottest parts of the image, the blue the "coldest" parts. In this case, the slightly reflective metal roofs are reflecting the clear sky, making them look much colder than their actual surface temperature.**

## How does infrared imaging radiometry work?

The radiometric camera is designed to generate raw digital counts data from each pixel in the array. These digital count data are converted to radiometric units, either in the camera itself, or in the back-end software that controls the camera and captures image data. Each pixel is represented by a photodiode "grown" into the detector material. The photodiodes generate photocurrent in response to radiation imaged onto the sensor by the optics. The photodiodes are each connected to a unit cell, a small circuit network for each pixel in the array. The pixel unit cell circuits on the ROIC generate voltages that vary with the photocurrent in a deterministic and very linear way. The voltages are streamed out of the sensor array into a high-speed analog to digital converter that converts the voltages into digital counts. From there, the image can be reconstructed and corrected for non-uniformity and then converted into radiometric image data. The electronics are carefully designed so that stable and repeatable conversions from photons to digital counts out are performed by the camera. In this manner, a calibration done in the laboratory will apply to image data taken some time later in the field.

## Radiometric Calibration Process

The camera must first be calibrated in the laboratory to be able to convert digital counts (the raw data from each pixel in the camera sensor) into in-band radiance values, and from there other radiometric units of interest. FLIR offers factory calibrations on most of the science cameras from the Niceville facility in Florida. The calibration laboratory can perform calibrations from a target temperature ranging from -50 °C up to 3000 °C on science-grade infrared cameras. Advanced users can also perform their own user calibrations using their own set of blackbody emitters and the CalibratIR routine within the ResearchIR software suite. This technique is described in detail later on in this book. The user calibration method is very useful for users that have very specific and/or ever-changing requirements that do not lend themselves well to factory calibrations.

There are some notable disadvantages to user calibrations, notably the lack of $T_{drift}$ correction. As the optics change temperature in a thermal IR camera, the amount of radiance striking the sensor changes, since the self-radiation of the lens and camera interior is detected by the sensor just the same as IR radiation from the scene. For high temperature calibrations, the effect of the optics parasitic radiation is negligible, but it can be a noticeable effect on radiometric accuracy when the target or scene is cold relative to the optics temperature. Another disadvantage of user calibrations is the requirement for "Bookkeeping" – the user has to make sure the correct calibration is associated with the image data. Finally, creating good user calibrations starts with currently calibrated blackbody sources. The cost of recalibration of blackbodies is high, and the lead times can be many weeks in the case of cavity blackbodies.

In a radiometric calibration, the calibrated infrared camera is pointed at a scene in the laboratory or in the field, and then image data is acquired in raw digital counts units. These digital counts data are converted to radiance and give the user an in-band radiance value for each pixel. These radiance values for each pixel in the scene can then be used to calculate the radiant intensity, as well as the apparent temperature of an object in the scene, since the in-band radiance emitted by a blackbody has a well understood relationship to its surface temperature.

Figure 19 shows a cartoon of this process. The laser rangefinder or radar is used to get the distance to the target in order to calculate the radiant intensity, and also to try to determine an in-band atmospheric transmission parameter. The atmosphere is a very important source of systematic error in infrared radiometry. The transmission is highly variable with wavelength, and the transmission will track changes in partial pressures of various IR absorbing gases, most notably water vapor and carbon dioxide.

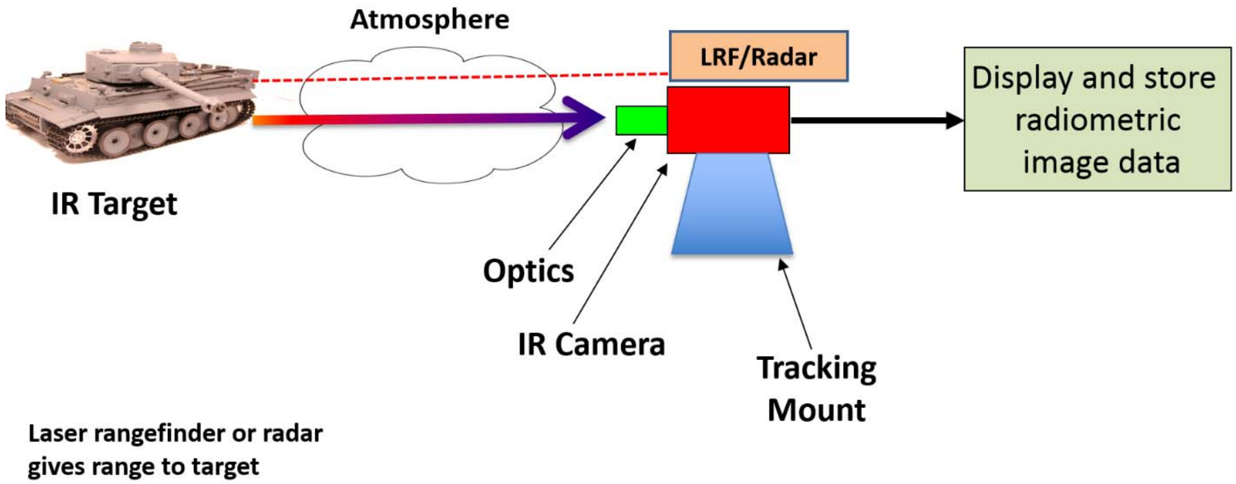


**Figure 19. Radiometry measurement scenario cartoon. The cameras are typically part of a larger system that includes range finding to the target and camera pointing motion control. The tracking mount may be equipped with a wide field of view thermal camera for automated target tracking.**

Figure 20a-b shows a radiance image of a Beechcraft King Air prop plane during engine startup, along with statistics for the image within the blue ROI around the exhaust pipe.

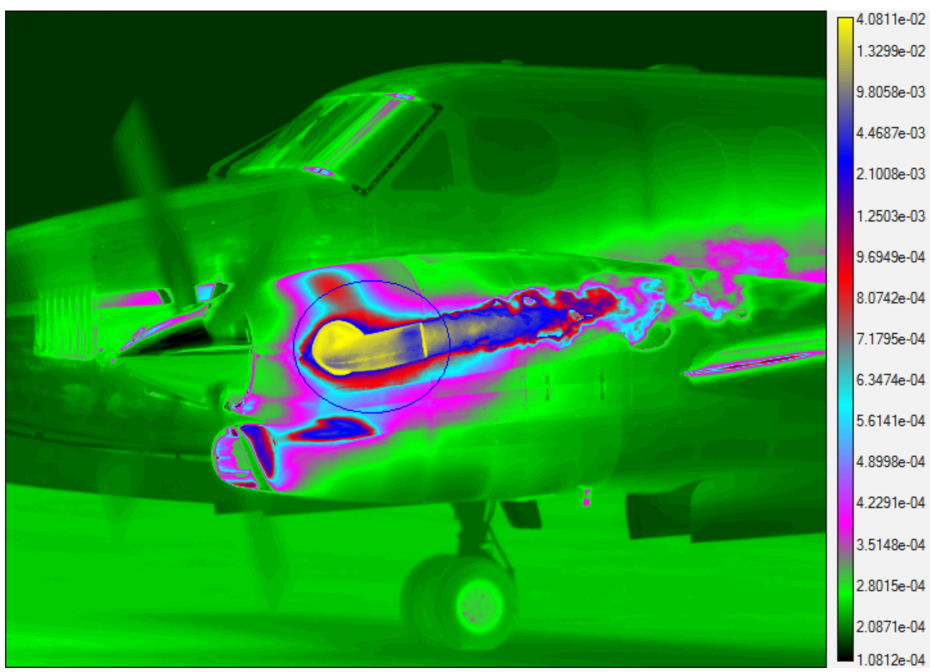


**Figure 20a. Beechcraft King Air imaged with a midwave IR camera (3-5µm) that measures radiance**

The base unit of imaging radiometry is the in-band radiance. This is the radiance measured by a camera with a particular spectral response shape, typically 3-5µm for midwave infrared cameras, but the measurement band can also be in the SWIR or LWIR bands. The in-band radiance is measured in units based either on watts or photons/sec. The dimensions of the power radiance unit are watts/sq. cm/steradian. The radiance value can be integrated over the emitter area to give the radiant intensity, or intensity for short. The radiant intensity unit has dimensions of watts/steradian. It is often used for point source characterization. In the case of the Beechcraft exhaust pipe, the pipe was 20.8m from the camera. The camera lens was 100mm focal length, and the pixels on the camera sensor had a pixel pitch of 25µm. The IFOV (also called the detector angular subtense) was 250 microradians. That information, combined with the 20.8m range, gives the physical size of each pixel at the range of the exhaust pipe (the ground sample distance or GSD). The square of the ground sample distance is the Single Pixel Area, expressed in square cm. The radiant intensity of the ROI is 14.9 watts/steradian as shown in Figure 20a. This value is the product of the mean radiance in the ROI multiplied by the ROI area.

Stats

| Statistic [units] | Exhaust Pipe |
|---|---|
| Mean [W/(sr-cm²)] | 5.7023e-03 |
| Std. Dev. [W/(sr-cm²)] | 8.6738e-03 |
| Intensity [W/sr] | 1.4969e+01 |
| Center [W/(sr-cm²)] | (285.0, 261.5) 1.2517e-02 |
| Maximum [W/(sr-cm²)] | (263, 282) 4.0811e-02 |
| Minimum [W/(sr-cm²)] | (307, 214) 2.9809e-04 |
| Number of Pixels | 9708 |
| Single Pixel Area [cm²] | 0.2704 |
| Area [cm²] | 2625.0 |
| Length [cm] | N/A |
| u Emissivity | ☐ 1 |
| u Distance [m] | ☑ 20.8 |

**Figure 20b. Radiance images statistics showing the Intensity of the Exhaust Pipe ROI.**

To summarize, infrared imaging radiometry is the measurement of an object or scene in any of these three units:

1. **Radiance**
2. **Radiant Intensity**
3. **Apparent Temperature**

Radiance is a range-invariant quantity. In the case where there is no atmospheric transmission loss (or it is corrected out), the apparent radiance of a target will not change with range. This effect works because the inverse square law is inverted for the area of a camera pixel projected back to the target. The irradiance on the camera lens aperture for any given area element on the target will drop off as the square of the distance. But the number of those area elements that are subtended by a camera pixel goes up as the square of the distance. The distance units cancel out.

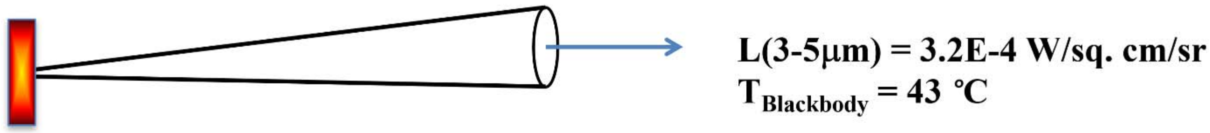


**Figure 21. Radiance cartoon. Radiance is the intensity per unit area of emitter.**

Each square centimeter on the orange glowing target emits 3.2e-4 watts/sr of radiant intensity. If we sum up the radiance values of every square centimeter of the target that is emitting, we get the radiant intensity. Radiant intensity is a radiometric unit of measure that is closely associated with target signature work. For example, a user may want to measure the radiant intensity of a helicopter as a function of view angle. The radiant intensity can be used to derive the signal at an infrared sensor like the seeker on a surface-to-air missile. Radiant intensity is the best way to characterize a small, faraway target – you can't resolve the parts of the target to measure radiance, so you just lump the whole thing into a single radiant intensity value.

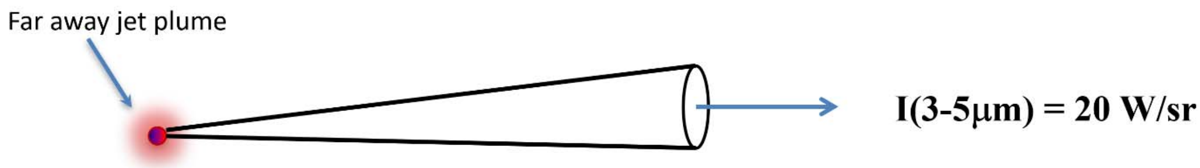


**Figure 22. Radiant intensity cartoon. Radiance can be thought of as the intensity per unit area of emitter.**

Apparent Temperature is called that because it is a temperature value that is derived from the radiance measured by the camera. The apparent temperature may not be well correlated with the true physical temperature of the object. An extreme example of this phenomenon is a sun reflection off a car windshield. A midwave infrared camera will regard this reflection as having an apparent temperature of thousands of degrees, when it is obviously not the case that the windshield has that surface temperature. In the case where the object is highly absorbent to IR radiation and closely approximates a perfect blackbody radiator, the correlation can be quite good, especially at close range where the atmospheric transmission is close to 100%. As the range gets longer and longer, the air path will absorb more of the target's emitted radiation and the apparent target temperature will more closely approximate the air temperature itself.

## Optics and the optical path

The optical path for infrared radiation is shown in the next series of images.

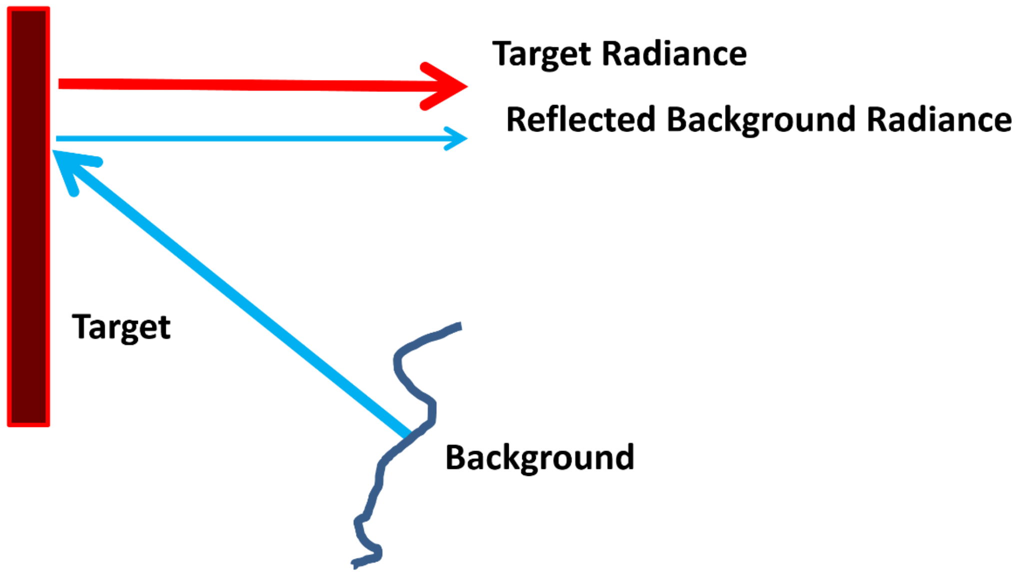


**Figure 23a. Target radiance. A graybody is a material that has a constant spectral emissivity over a particular band.**

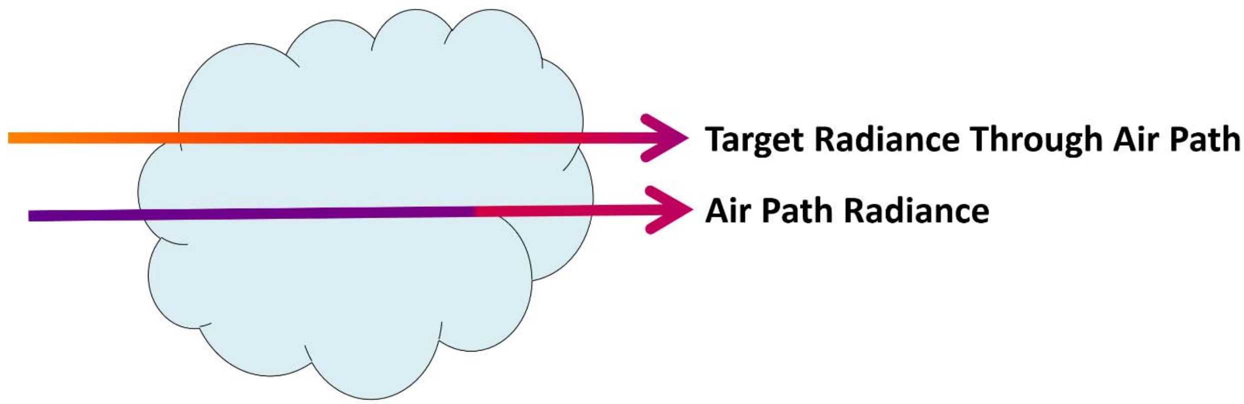


**Figure 23b. Target radiance propagation and air path radiance**

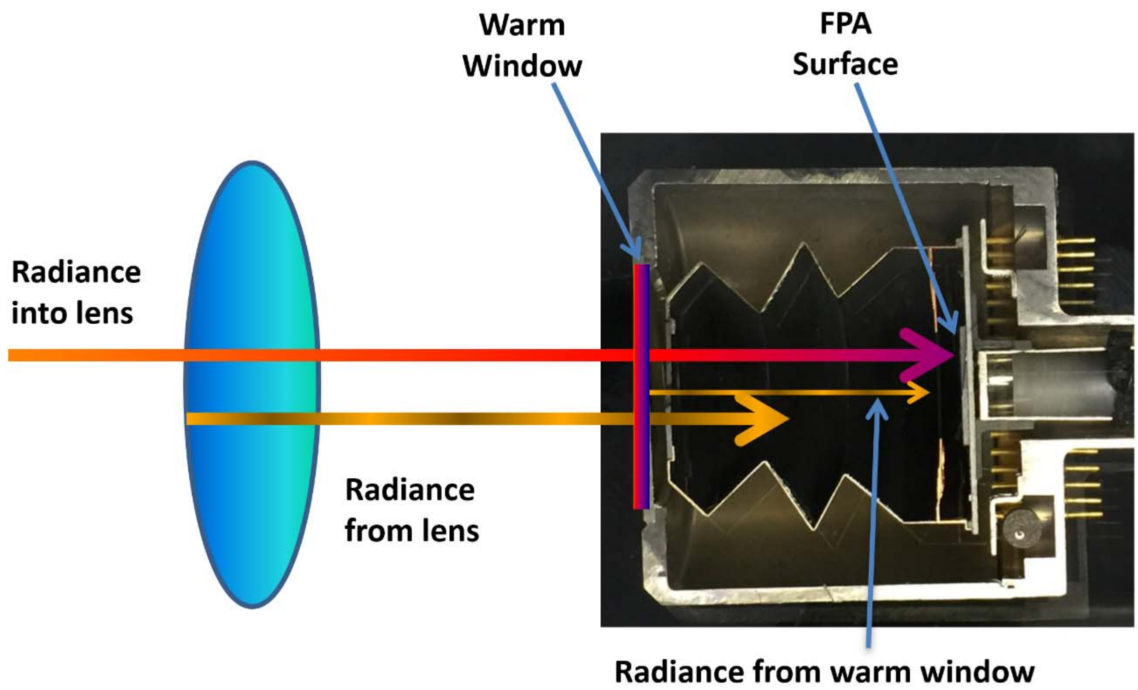


**Figure 23c. Three radiance components going into camera. Radiance from the warm window is almost always the smallest component.**

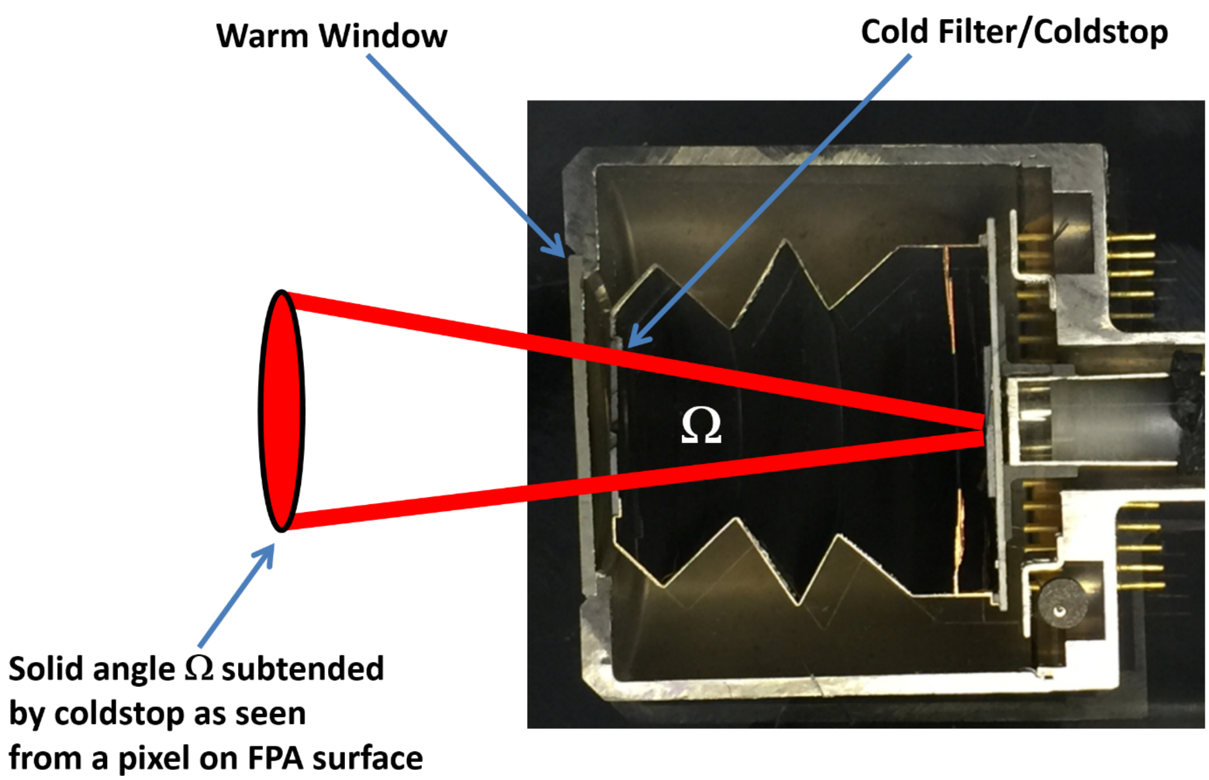


**Figure 23d. Geometry of coldstop and FPA, with solid angle of coldstop as seen from FPA surface**

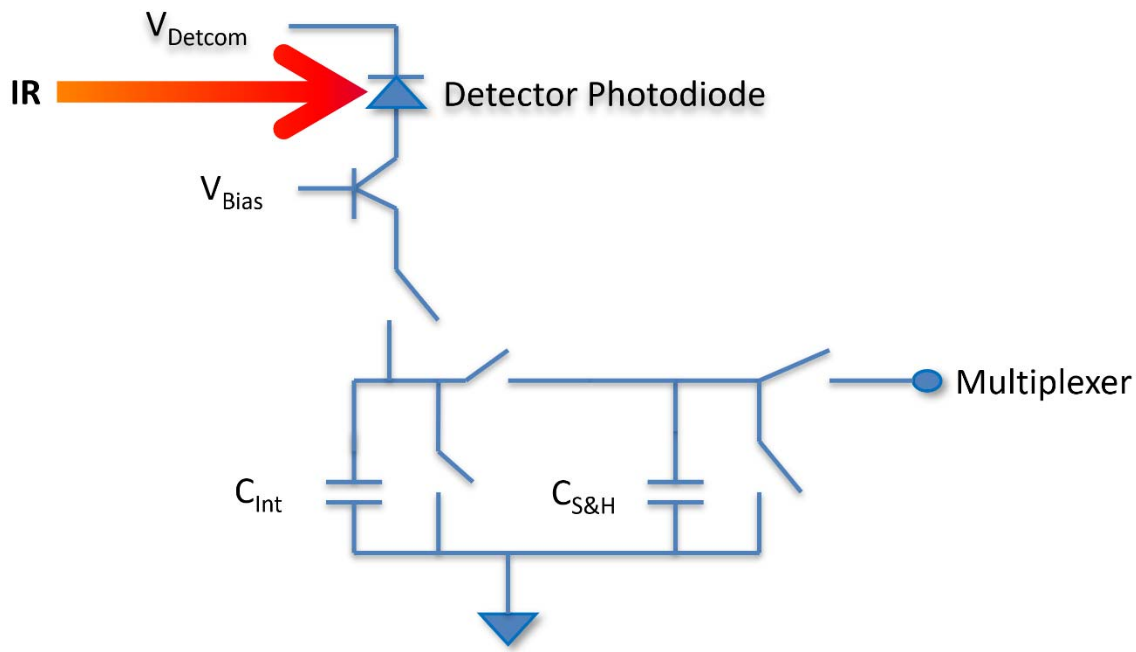


**Figure 23e. Infrared radiation and FPA unit cell circuitry**

## Camera Radiance Linearity

Both power and photons/sec incident on a camera pixel are linear in the in-band radiance of the object subtended by the pixel's instantaneous field of view (IFOV). If the camera is designed to have an output that is linear in the incident flux on a pixel, then the camera system will be linear in scene radiance. A linear transfer function between scene flux and camera output makes calibration straightforward, since a simple slope and intercept is enough to get a good fit. A linear calibration lends itself to working alongside standard infrared camera non-uniformity corrections (NUCs), which are themselves linear transfer functions. Put another way, infrared cameras use linear NUCs, and thus will only output radiometrically accurate data within the linear range of the camera.

The detectors used in most common photodiode-based cameras are very linear in incident radiant flux. The readout ICs used to convert photocurrent from these detectors into voltages are also linear over a certain range of their performance. Finally, the analog to digital converters in the cameras that convert the voltages from each pixel into digital data (digital counts, or LSBs, as they are sometimes called) are linear. Thus, there is a linear relationship between the incident radiant flux on a camera pixel and the digital counts out for that pixel.

This is what makes it possible to calibrate cameras to measure radiance of scenes and targets in a linear manner, which in turn makes it possible to correct for effects such as emissivity, atmospheric transmission and temperature, and the transmission and temperature of a protective enclosure window in front of the lens. The camera detects radiance that is emitted from a target, but also reflected off a target with an emissivity less than 1 (which describes real materials), emitted by the atmosphere and emitted by a window, if one is used. The linear relationship between radiance and digital counts makes it mathematically simple to subtract off these other radiance contributions and correct for transmission losses, all to get to the "true" radiance emitted by the object. Once that true radiance value is known, it can be scaled by the in-band emissivity of the object's surface and converted to an apparent or radiation temperature, as discussed later on in this document.

If the radiance corrections are done properly, the apparent temperature measurement can be highly correlated with the kinetic surface temperature and give very accurate temperature results.

## Planck's Law

The relationship between in-band radiance and emitter temperature is determined by Planck's Law, which is a theoretically determined and empirically tested expression that calculates the spectral exitance of a blackbody radiator as a function of wavelength and temperature. This spectral exitance is integrated over wavelength to get the total in-band exitance. Exitance is measured in units of watts/sq. cm and shown in a cartoon form below in Figure 24:

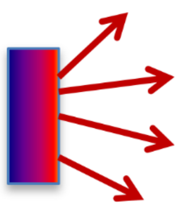

**M(3-5µm) = 3.75E-4 W/sq.cm (15 °C BB)**

**Figure 24. MWIR (3-5µm) exitance of a 15 °C blackbody**

The radiance of a blackbody emitter is simply the exitance divided by pi steradians. A flat emitter emits into a hemisphere which has a solid angle of 2pi steradians, but because of the cos(θ) functional shape of the projected area of the flat emitter, the *effective* solid angle is pi steradians. The other important assumption that gets you pi steradians of effective solid angle is that the target is a so-called Lambertian emitter, which means that light rays are emitted at completely random angles from the surface into the full hemisphere. This condition is met by cavity blackbodies, and to a large extent, rough, non-metallic surfaces.

The in-band radiance for a given temperature is determined by numerical integration of the Planck function for a given temperature over the waveband of the camera system. We often use a simple square-band approximation for the spectral response of a camera. This assumes the response is zero outside of the band limits and 1 inside. ResearchIR software can also make use of a spectral response file with the prn extension: *.prn. This file contains two columns of data: wavelengths and relative power spectral response. This is the spectral response to power with the peak normalized to 1. An example of this file and a graph of it is shown in Figure 25 below. These data were measured by a spectral response system in Niceville. The sampling interval is 0.01µm. The file is called SC8201.prn and it is used for factory calibrations of this camera type.

| | |
|---|---|
| 2.8 | 0 |
| 2.85 | 0 |
| 2.9 | 0.01 |
| 2.95 | 0.17 |
| 3 | 0.61 |
| 3.05 | 0.65 |
| 3.1 | 0.76 |
| 3.15 | 0.8 |
| 3.2 | 0.76 |
| 3.25 | 0.7 |
| 3.3 | 0.67 |
| 3.35 | 0.69 |
| 3.4 | 0.66 |
| 3.45 | 0.69 |
| 3.5 | 0.71 |
| 3.55 | 0.74 |
| 3.6 | 0.75 |
| 3.65 | 0.75 |
| 3.7 | 0.76 |
| 3.75 | 0.77 |
| 3.8 | 0.78 |
| 3.85 | 0.8 |
| 3.9 | 0.81 |
| 3.95 | 0.82 |
| 4 | 0.83 |
| 4.05 | 0.83 |
| 4.1 | 0.84 |
| 4.15 | 0.86 |
| 4.2 | 0.86 |
| 4.25 | 0.86 |
| 4.3 | 0.88 |
| 4.35 | 0.89 |
| 4.4 | 0.9 |
| 4.45 | 0.91 |
| 4.5 | 0.92 |
| 4.55 | 0.93 |
| 4.6 | 0.93 |
| 4.65 | 0.94 |
| 4.7 | 0.96 |
| 4.75 | 0.99 |
| 4.8 | 0.98 |
| 4.85 | 1 |
| 4.9 | 0.97 |
| 4.95 | 0.47 |
| 5 | 0.11 |
| 5.05 | 0.03 |
| 5.1 | 0.02 |
| 5.15 | 0.02 |
| 5.2 | 0.02 |
| 5.25 | 0.02 |
| 5.3 | 0.02 |

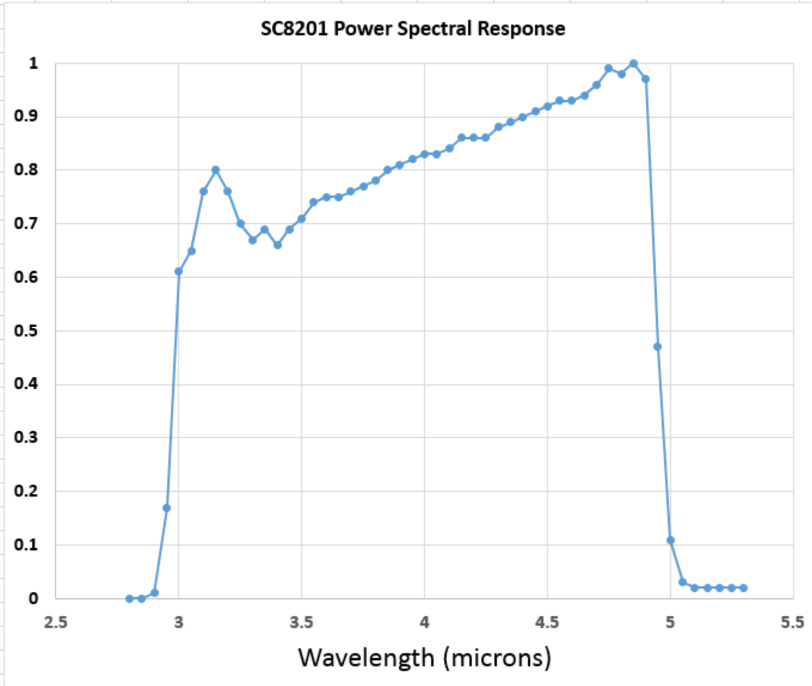


**Figure 25. Spectral response file for ResearchIR. The file is called a *.prn file which is readable in a text editor like Notepad. The graph shows the rising spectral response with wavelength, because the data are in power units, not photon units.**

The shape of the curve does not match the 3 to 5µm square band approximation shape. There is a rise in the response with wavelength up to 5µm, which is expected[5]. The radiance values that are calculated by ResearchIR will be different depending on whether the square band or a *.prn file is used. However, the apparent temperature measurements will be very similar with either spectral response function choice. This makes sense, since the radiance/temperature lookup table

[5] A detector like InSb is flat in photon response. In a spectral response curve with power units, the assumption is that the variable wavelength beam used to measure sensor response has the same power at all wavelengths. But it takes more photons to give the same power at longer wavelengths, so the response rises.

is calculated using the same spectral response used to derive the radiance calibration points. In the case of an apparent temperature measurement, it really does not matter what the intermediate radiance value is. However, it is a very good idea to use the best spectral response data available for the camera, since large deviations from the true spectral response will introduce significant curvature in the radiance versus counts plot. Yet these curves should always be very close approximations to perfect straight lines. If they aren't, then there is something wrong with the calibration, and it is probably the choice of the spectral response curve or band limits.

## User Calibration:

A user calibration is performed by pointing the camera at a calibration target blackbody and acquiring digital image data of the target. The temperature of the blackbody is changed to a new value and more data are taken. In the end, a table is created that contains the blackbody temperatures, emissivity values, in-band radiance values and digital counts. The radiance values are derived from the temperature and emissivity of the blackbody, with an optional reflected radiance component correction. A curve is then fit to the scatterplot of the in-band radiance and the digital counts data.

Figure 26 below shows a user calibration done with the correct band limits of 7.5 and 10.5µm. The goodness-of-fit $R^2$ value of the least-squared fit is 1.0000 and the data are highly linear.

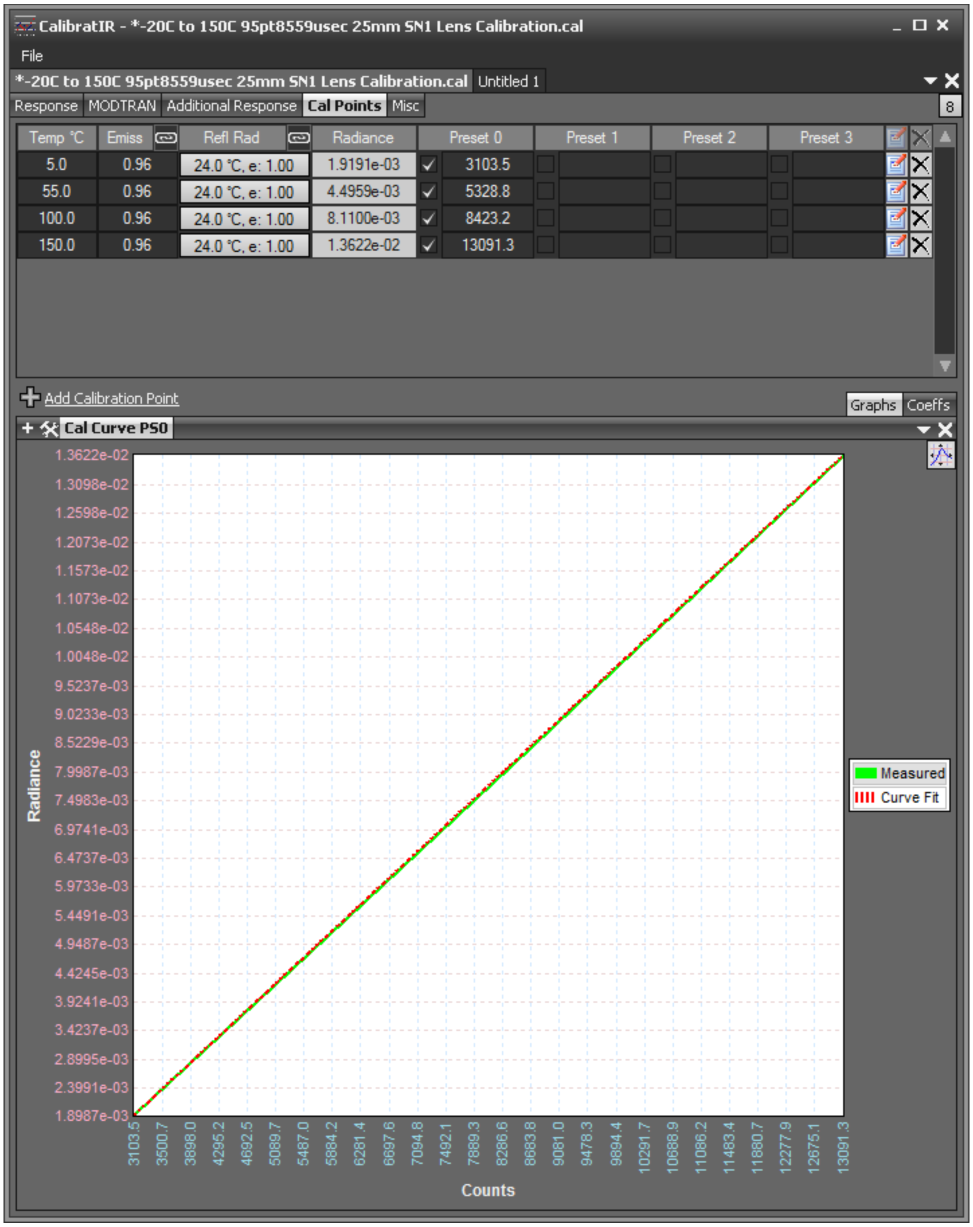


**Figure 26. Calibration graph in CalibratIR software for an SLS camera with band limits correctly set to 7.5 and 10.5µm.**

Figure 27 below is the same user calibration with the default values of 3 and 5µm used for the spectral response. Note the severe curvature of the line, and the goodness-of-fit $R^2$ value is now 0.9499, which is a poor fit indeed. This is the hallmark of an incorrect spectral response limits or the spectral response file choice.

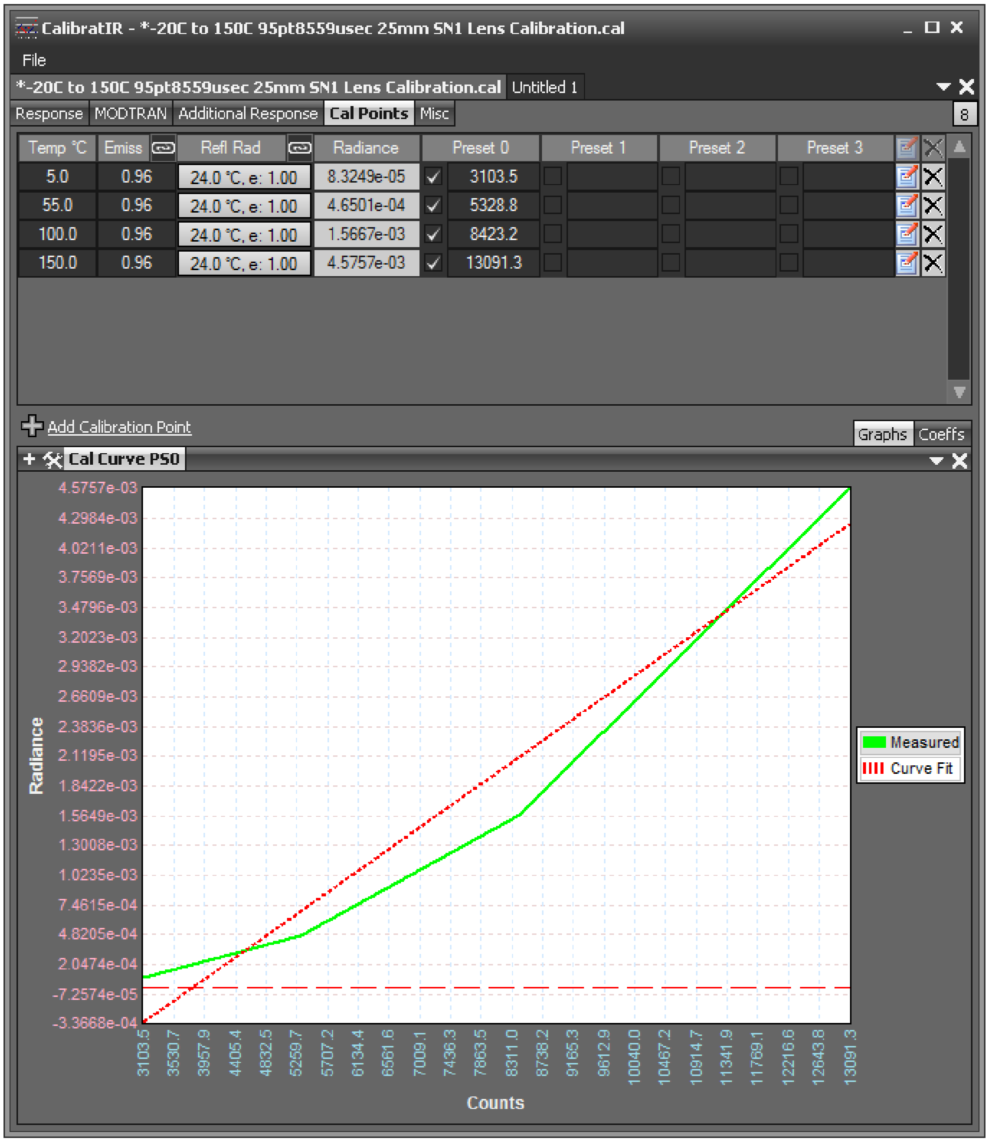

**Figure 27. Calibration graph in CalibratIR software for an SLS camera with band limits incorrectly set to 3 and 5µm. The measured data are not even close to linear. This indicates a problem with the response file.**

## Calculating Radiance

In FLIR's ResearchIR software, the radiance is expressed in power units: watts/square cm/steradian, although it is written as W/(sr-cm$^2$). For the purposes of simplicity, I will refer to the power units of radiance as "wicks". This is an abbreviation of Watts/sq. cm/steradian. This is my own notation which I use in my radiometry classes – it is not at all an official name for a unit! Figure 28 shows a Planck function curve for a blackbody target at 300 °C . The pink-shaded area is the total in-band exitance for the target in the 3-5µm band. This calculation assumes a square band. This wavelength-integrated exitance is divided by pi to get the in-band radiance for a blackbody. For targets with an emissivity less than 1, the radiance is given by the product of the emissivity and the blackbody radiance. For the curve shown below, the in-band exitance is 0.1299 watts/sq. cm – this is the area shaded in pink. The radiance is the exitance divided by pi steradians, or 0.04135 watts/sq. cm/sr. One can shorten that to 41.35 milliwicks.

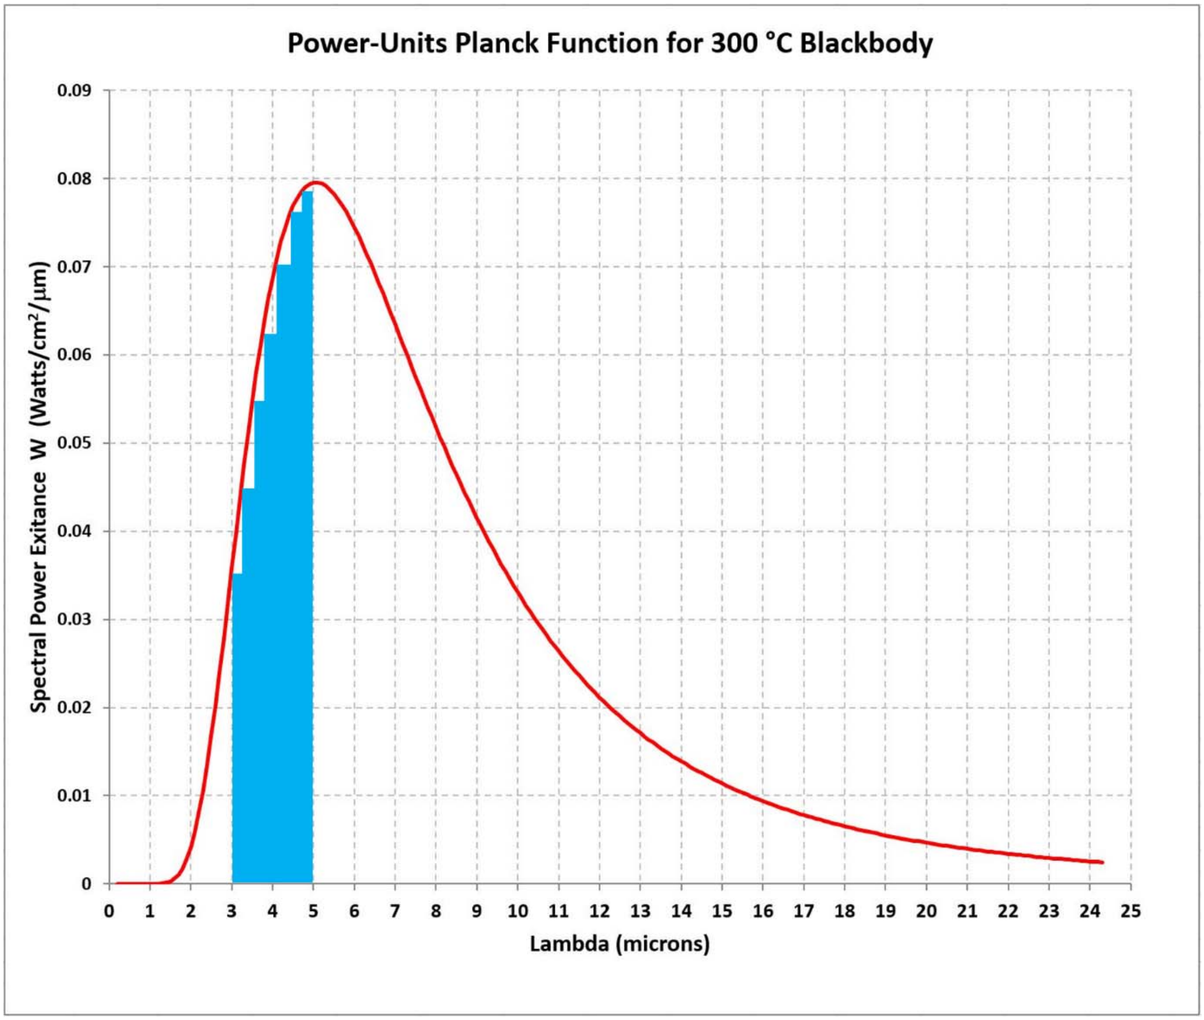


**Figure 28. Planck function for a 300 °C blackbody emitter. The pink area is the 3-5µm in-band exitance.**

Consider a laboratory cavity blackbody used for calibration. The cavity blackbody has a factory-specified emissivity of 0.99 and is set to a temperature of 300 °C . The radiance in the 3-5µm band is 99% of the 300 °C in-band blackbody radiance, or 40.9 milliwicks. The calibration software determines this radiance. The calibration procedure involves setting up a series of sources in a sequence of radiance values (different emitter temperatures) and then presenting them to the camera. In Figure 29 below, a FLIR camera is being presented to a series of cavity blackbodies at different temperatures to develop a set of calibration points.


**Figure 29. A set of cavity blackbodies being used for calibration of a FLIR camera**

The camera responses in digital counts are measured for each temperature point, the in-band radiance is calculated, and an empirical curve fit is made to relate radiance and counts. This curve is found to be highly linear over a substantial range of the camera's dynamic range. Figure 30 shows an example of a calibration curve for a different type of IR camera called a See-Spot with a bandpass of 4.5 to 4.8µm. The measured data points lie very close to a straight line, with an $R^2$ value of 0.9999. This graph can be used as-is to make simple radiance measurements. If a target gives 7000 digital counts, the in-band radiance can be read off the graph as 0.15 milliwicks.

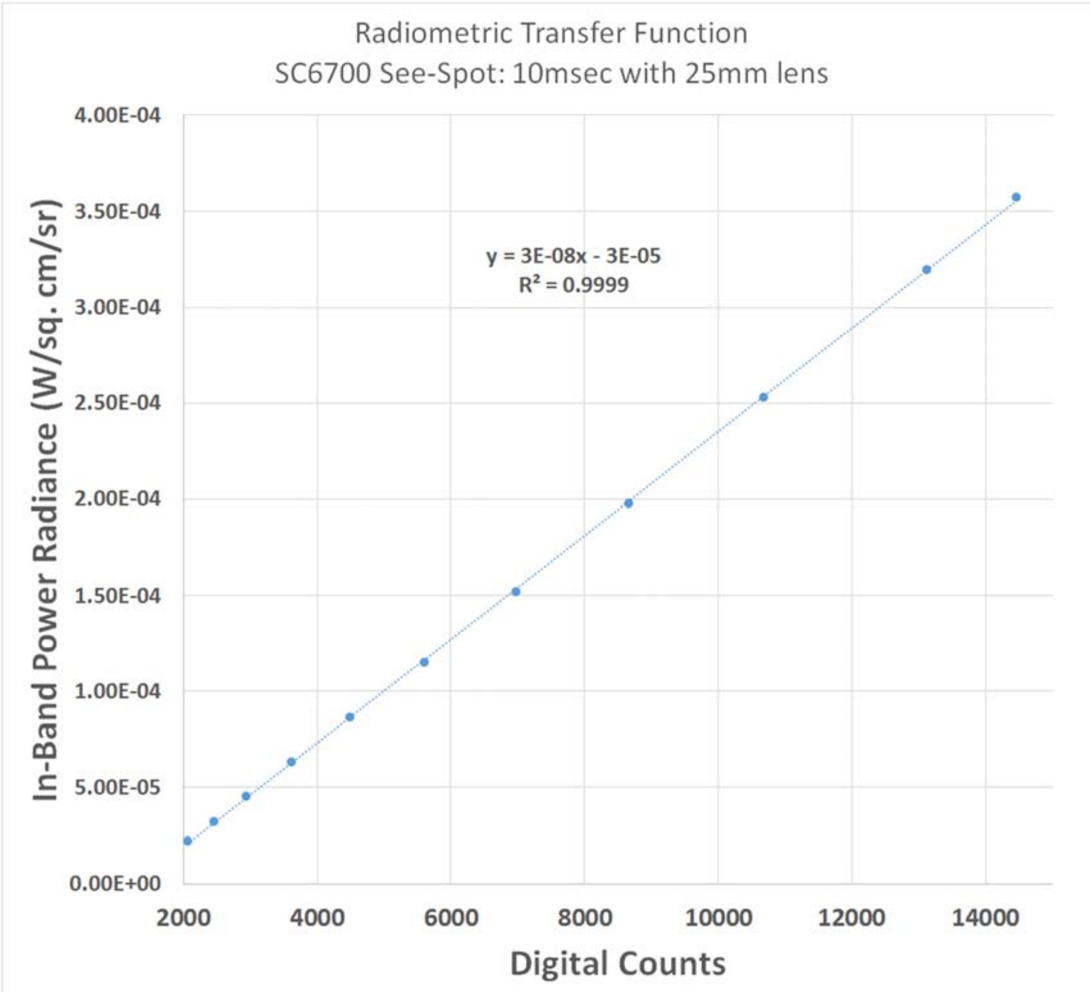


**Figure 30. FLIR SC6700 See Spot camera radiometric transfer function for a 10ms integration time and a 25mm lens.**

The slope and intercept of the curve are determined during the calibration process, then these coefficients are used by ResearchIR to convert counts to radiance, enabling one to make measurements of unknown radiance sources. The camera is pointed at a target of interest, like the exhaust manifold on a propeller plane engine, and the digital counts measured for pixels subtended by the manifold are converted into radiance units using the coefficients.

From there, the software converts the radiance into an apparent temperature. The degree to which this apparent temperature can be corrected to read out the actual surface temperature depends on how well the user characterizes the setup. For instance, the emissivity of the target and the reflected temperature are used to correct the radiance, and hence the apparent temperature. Other effects that modify the radiance the camera sees include atmospheric transmission and the air path temperature. When atmospheric transmission is not 100%, radiance from the target is attenuated. At the same time, the atmosphere can contribute radiance to the camera measurement. A hot humid atmosphere can itself emit a significant amount of path radiance which will skew target temperature measurements if it is not compensated for correctly. The goal of temperature measurement is almost always to measure the physical temperature of a target. The only way to do this correctly is to determine the radiance emitted by the target surface, then derive a temperature from that measurement using the in-band emissivity of that surface. The camera measured radiance is composed of the target emitted radiance as well as unwanted sources of radiance like radiance reflected off the target surface and radiance emitted by the air path onto the camera sensor.

## Calculating Temperature

Once the radiance has been "corrected back to the target", and adjusted for the target emissivity and reflected radiance, the conversion of the equivalent blackbody radiance to an apparent temperature can be done with a lookup table, where a series of temperature values and the radiance they give are arranged in a table so that the software can "look up" the temperature for a given radiance value. The integral of the Planck function for a discrete bandpass is not solvable in closed form. Only in the end cases can it be solved in closed form. One of these is the case of a very narrow wavelength range. The other case is where one integrates over all wavelengths, the integral gives the Stefan-Boltzmann Law where the total exitance is equal to the Stefan-Boltzmann constant times the absolute temperature to the fourth power. Figure 31 below shows a lookup table in 5 °C step increments for a 3-5µm blackbody emitter bandpass. If the camera reads a radiance of 0.555 milliwicks after correcting for other radiance sources and the target emissivity, the target temperature is 60 °C . The lookup tables can be generated with 1 °C increments and intermediate temperature values can be interpolated. Another approach is to fit a higher-order (order 6, for example) polynomial to the radiance versus temperature data. This can backfire near the endpoints, since higher-order polynomials are inherently oscillatory.

| Blackbody Emitter Temperature (°C) | 3-5µm In-Band Power Radiance (Watts/cm2/sr) |
|---|---|
| 0 | 6.44E-05 |
| 5 | 7.97E-05 |
| 10 | 9.79E-05 |
| 15 | 1.19E-04 |
| 20 | 1.45E-04 |
| 25 | 1.74E-04 |
| 30 | 2.09E-04 |
| 35 | 2.49E-04 |
| 40 | 2.95E-04 |
| 45 | 3.48E-04 |
| 50 | 4.08E-04 |
| 55 | 4.77E-04 |
| 60 | 5.55E-04 |
| 65 | 6.43E-04 |
| 70 | 7.42E-04 |
| 75 | 8.53E-04 |
| 80 | 9.77E-04 |
| 85 | 1.11E-03 |
| 90 | 1.27E-03 |
| 95 | 1.44E-03 |
| 100 | 1.62E-03 |

**Figure 31. Lookup table relating emitter temperature and radiance for a 3-5µm bandpass**

These lookup tables are not linear at all, which is why it is not standard practice to convert directly from counts to temperature. Since the relationship is non-linear, there is no obvious way to correct for linear effects like those introduced by atmospheric transmission, where the radiance is modified by multiplying it by a transmission value less than or equal to 1.cm[2]

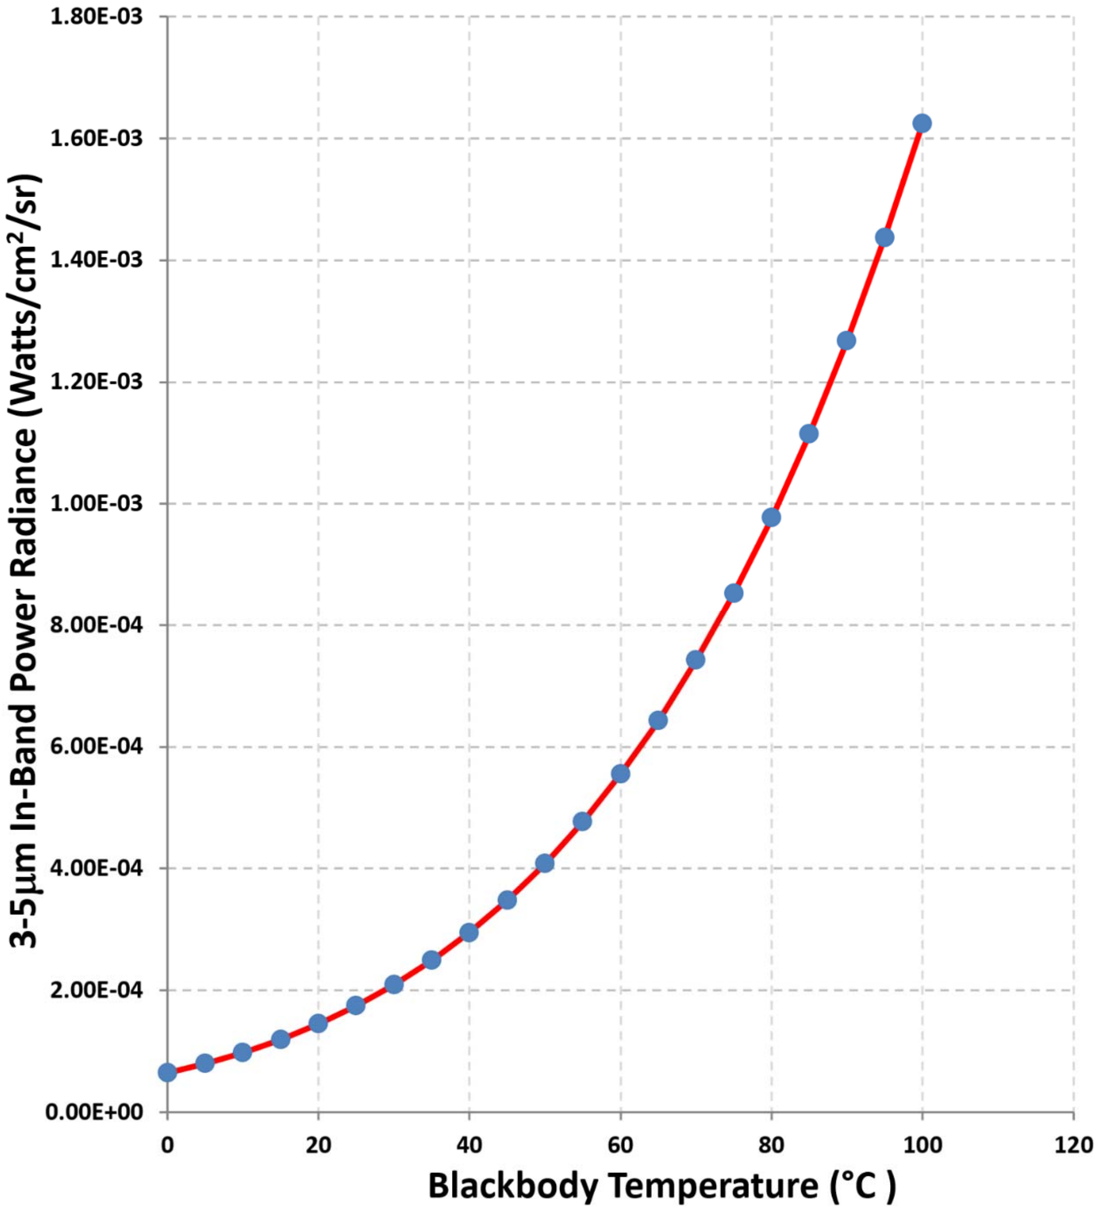


**Figure 32. A plot of the lookup table in Figure 31: the 3-5µm in-band radiance versus emitter temperature from 0 to 100 °C in 5 °C steps.**

## Stability

Radiometric IR cameras are designed to have excellent stability over time. If one points the camera at a stable blackbody source, the digital counts in a region of interest drawn on the source will stay stable to within a few digital counts over long periods of time, assuming that the self-radiation from the camera and its optics has also stabilized.

## Repeatability

A repeatable measurement means that an infrared camera that is calibrated can be turned off, then turned on a few days later and pointed at a source of known radiance and the measurement will be the same to within statistical noise. Getting good absolute repeatability will depend on various factors. For long integration times where cameras are more sensitive, one will only achieve repeatability when the internal self-radiation and optics radiation is the same as it was during calibration. This is why it is important to turn on cameras well in advance of any calibrations or subsequent measurements to let everything stabilize.

An example: Consider an InSb camera set up for 3-5μm imaging with an f/2.5 cold shield. The integration time is 1.6ms. This is a sensitive camera scenario where saturation of the sensor will occur at about 55 °C scene temperature. The camera with the lens installed is turned on for several hours and stabilizes in temperature before it is radiometrically calibrated at 1.6ms integration time. A 40 °C target with emissivity of 0.99 is found to give 8307 counts.

Now suppose the camera is turned off and a few days go by. I want to make a measurement with the camera, so I turn it on, but I am impatient and start measuring that same 40 °C target only 15 minutes after turning on the camera. The camera and optics are at a lower temperature than they were during calibration. The 40 °C target now only gives a count value 8007 counts. The 300-count drop in counts compared to what is expected is due to the optics being ~5 °C colder than during calibration. If one does not account for this effect, then measurements will be in error. In this case, the target will appear to be colder than it really is. We will revisit this later with some actual data from a camera. FLIR has implemented a proprietary temperature drift correction to mitigate this effect in the latest families of science cameras.

## Short Integration Time and $T_{drift}$

If I now turn the integration time down to 100μs, the camera will behave very differently in that it will no longer be sensitive to optics or camera ambient temperature to nearly the same degree at it was at 1.6ms integration time. If I calibrate the camera at 100μs integration time, then I may be able to make accurate radiometric measurements much sooner after the camera is turned on from a cold start, and still get excellent repeatability. This is because the imaging sensor will be much less responsive to changes in the temperature of the optics and the camera interior. As soon as the camera produces an image, it can be used to make measurements, and measurements made an hour later after the camera interior heats up from 22 °C to 37 °C (for example) should be very much the same for a stable temperature target. The temperature drift correction is no longer as important to the outcome. This has implications for radiometric calibrations done by camera users which do not have the factory temperature drift correction.

## Radiant Intensity Measurement

Radiant intensity, or intensity, as it is called in ResearchIR, is a measurement of a target that integrates the radiance of the target over its emitting area to give the radiant intensity in watts/steradian in a particular direction. The direction part is very important to consider, because some IR sources will have radiant intensities that are extremely dependent on direction. A highly collimated laser beam will have enormous intensity on the beam axis, but almost negligible intensity off axis. A collimated laser beam directed into an IR camera will focus to the size of the blur spot on the sensor surface and can easily damage it.

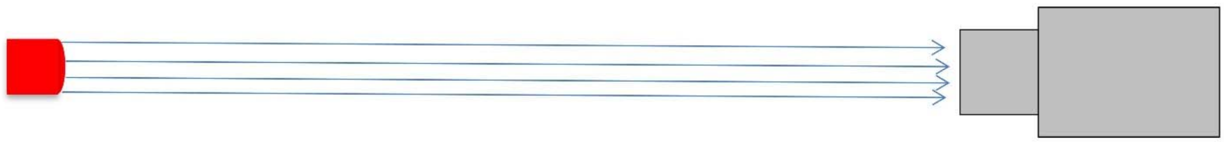

**Fig. 33a. Collimated laser source sending light rays to an IR camera. The off-axis radiant intensity is negligible.**

The opposite of a collimated source is a spherically symmetric IR source like the sun. The sun will radiate its power into a full 4pi steradians of solid angle and the intensity will be independent of direction.

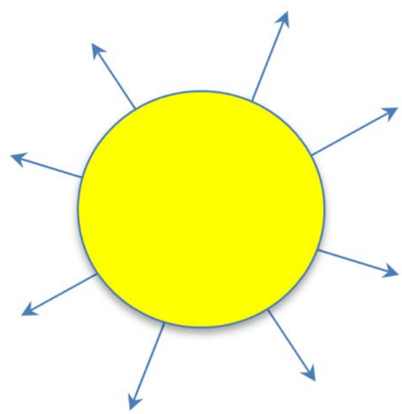

**Fig. 33b. The sun's disc is a very isotropic IR source: radiant intensity is the same from any view angle**

The radiant intensity of a target can be used to determine how detectable it will be to an infrared sensor on a missile, for example. Many missile seekers have a low spatial resolution IR sensor. They will fly toward a hot spot without ever resolving spatial details of the hot spot. The actual measurement of interest for this sort of target signature work is to measure the **radiant intensity contrast**. This is the radiant intensity of a target, minus the radiant intensity of an equal area of scene around the target.

Radiant intensity of a very small target is most often measured with respect to the camera aperture, and not "corrected back" to the target. By small target, I mean a target that is less than 10x10 pixels in size. When a target is that small in the image, it will appear to have less radiance than it actually does, due to optical effects which are discussed in a later chapter. If the target is more than 1000m away from the camera, as is often the case with radiant intensity measurements of aircraft, there is also a great deal of uncertainty about the effect of the air path transmission on the measurement, because air path transmission is a parameter that strongly affect radiance measurement and is not easy to measure for >1000m ranges. This air path topic will also be addressed in a later chapter. Given these optical effects of target apparent size and air path transmission losses, the best one can do usually is measure the radiant intensity at the camera aperture, and then leave it at that. Thus, small, distant target measurements are of apparent radiant intensity, not source radiant intensity.

The following steps outline how to measure radiant intensity:

- First have both radiance and spatial calibrations active in ResearchIR software. The spatial calibration will be used to determine the physical size of the emitting area on the target.
- Use Radiance as the displayed units in the top right corner of the ResearchIR window
- Draw an ROI around a target, and then turn on segmentation to suppress background pixels. Adjust the segmentation lower limit setting to turn all the pixels blue around the target, so that they won't be counted in the radiance measurement statistics.
- Measure distance to the target from the camera, then enter the distance value into the Stats entry in ResearchIR to get the correct single pixel area, and target area. Here is an example where the distance to the target in the ROI is 10.4m. The checkbox must be checked in order for the operator to enter a numeric value, as shown in Figure 33c:

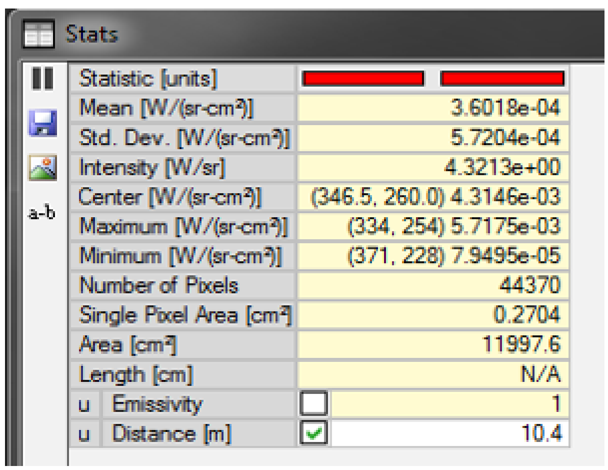

Stats

| Statistic [units] | |
|---|---|
| Mean [W/(sr-cm$^2$)] | 3.6018e-04 |
| Std. Dev. [W/(sr-cm$^2$)] | 5.7204e-04 |
| Intensity [W/sr] | 4.3213e+00 |
| Center [W/(sr-cm$^2$)] | (346.5, 260.0) 4.3146e-03 |
| Maximum [W/(sr-cm$^2$)] | (334, 254) 5.7175e-03 |
| Minimum [W/(sr-cm$^2$)] | (371, 228) 7.9495e-05 |
| Number of Pixels | 44370 |
| Single Pixel Area [cm$^2$] | 0.2704 |
| Area [cm$^2$] | 11997.6 |
| Length [cm] | N/A |
| u Emissivity | 1 |
| u Distance [m] | 10.4 |

**Fig. 33c. Entering distance into the Stats window**

- ResearchIR returns a ***radiant intensity*** of ROI in watts/steradian – it is the sum of the radiant intensities of each pixel in the ROI. In this example, the Intensity shown in Figure 33c is 4.3213 W/sr. The value is given to 4 decimal places, but this is **not** an indicator of the absolute accuracy of the measurement, only the resolution.

The equation used to determine the radiant intensity is as follows:

**Radiant Intensity = SUM over ROI(pixel area at that distance * Radiance of pixel)**

**Radiant Intensity = Mean Radiance of ROI * Area of ROI**

In the example above, the mean radiance is 3.60e-4 W/sq. cm/sr and the area is 11,998 sq. cm. The product of those two values is 4.321 W/sr.

The calculation of radiant intensity from radiance can be deduced by dimensional analysis. The units of radiance are watts/sq. cm/steradian. If a pixel has a particular radiance value, and the area of that pixel at the target range is multiplied by the radiance value, then the units of the product are Watts/steradian: the area units in the numerator and denominator cancel out.

Here is an example of a measurement of radiant intensity with and without segmentation:


**Figure 34a. Radiant intensity measurement of ROI without a segmentation filter applied. The RI is 8.30 W/sr.**

The radiant intensity is 8.30 Watts/steradian. All the pixels in the rectangular ROI are being counted. If the segmentation parameter is adjusted to mask out all the pixels below a radiance threshold, then the radiant intensity drops down. This segmented measurement is more meaningful, since it gives the radiant intensity of the hot portions of the scene, not the ambient portions.

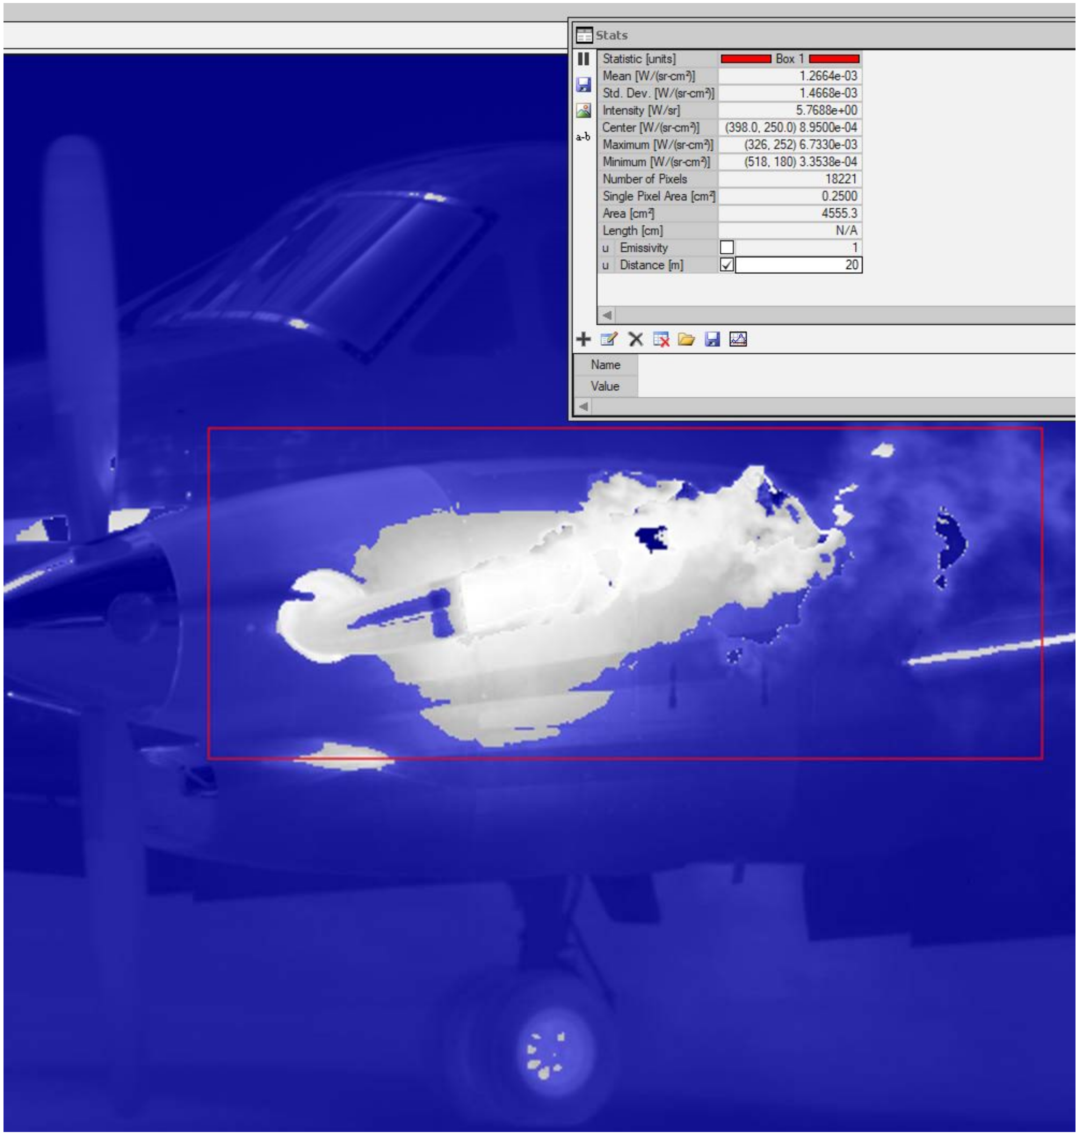


**Figure 34b. Radiant intensity measurement with segmentation filter applied. The Intensity drops down to 5.77 W/sr.**

Figure 35a-b shows RI measurements of the plane before and after engine start. Taking the difference between the two radiant intensity values gives the radiant intensity "signature" of the engine start process, which is a metric that could be used by surveillance aircraft to determine what is happening at an airfield. The difference is about 15 W/sr in this example.

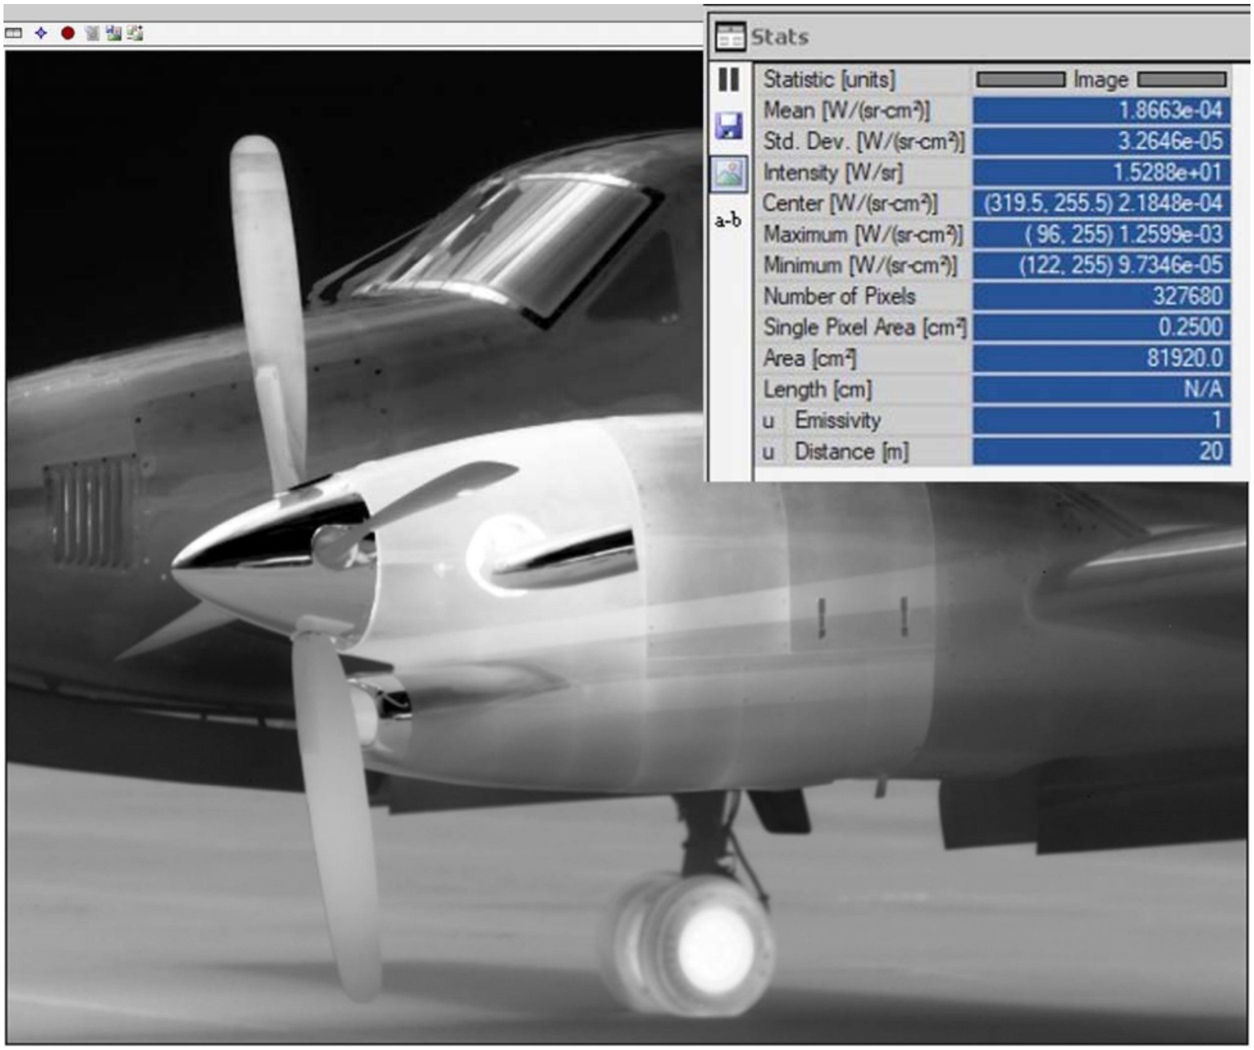


**Figure 35a. Beechcraft King Air before engine start. The RI of the whole image is 15.29 W/sr.**

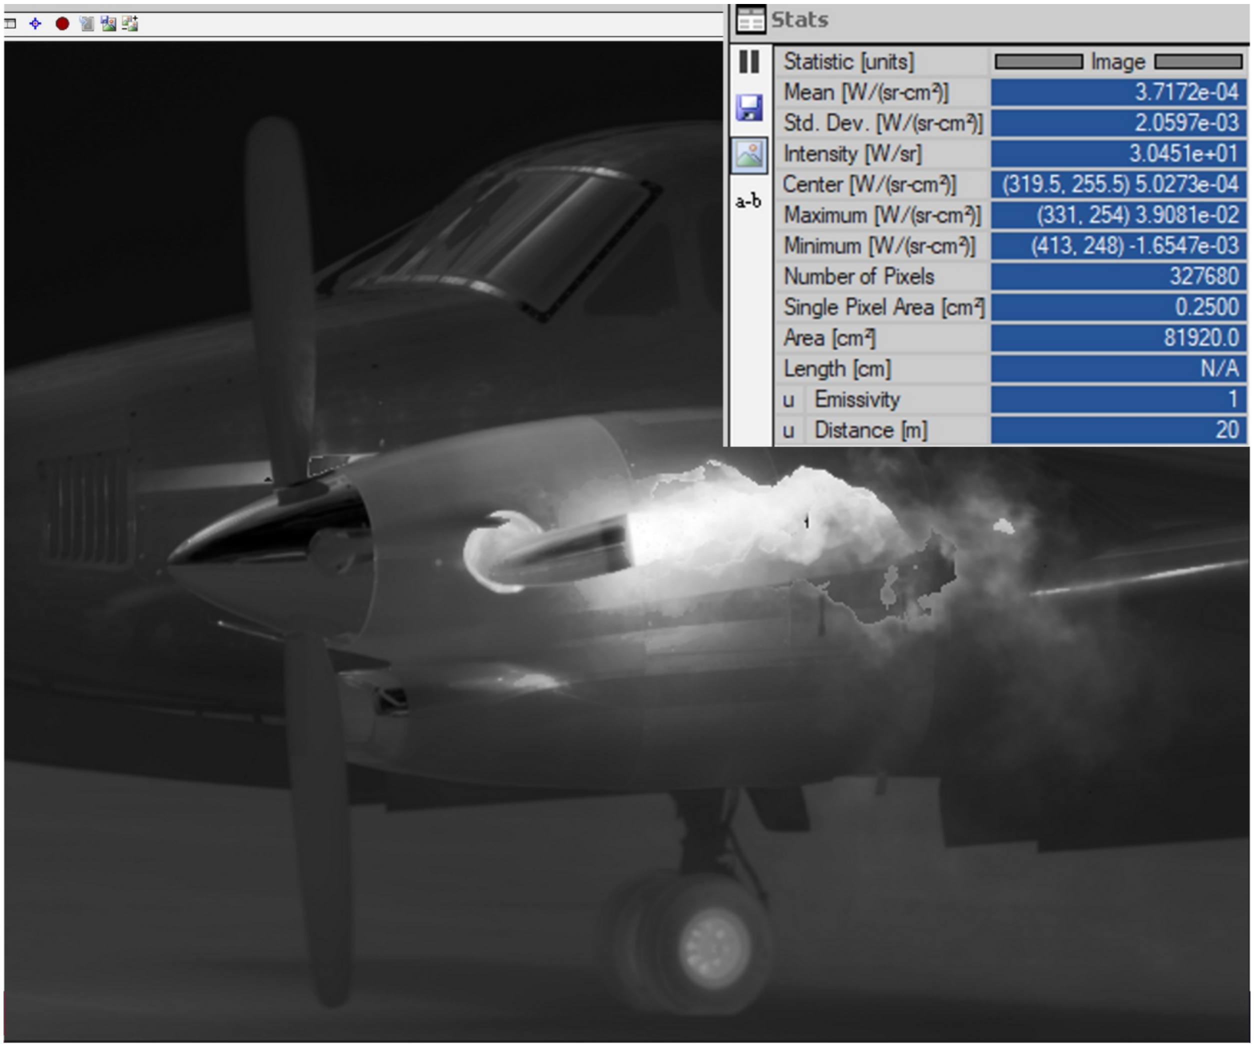


**Figure 35b. Beechcraft King Air a few seconds after engine start. The RI of the whole image is 30.45 W/sr.**

## Radiant Intensity of Point Targets

Radiant intensity is a useful way to characterize a target that is very small relative to the camera system's IFOV. A point target is a target that is much smaller in angle than a single IFOV angle in terms of its physical size and its distance from the camera. An example of this would be a target that is 1m square at a range of 6000m from a camera with an IFOV of 1 milliradian. The angle subtended by the target is 1/6 of an IFOV, making it look essentially like a point source to the camera system. When a target like this is viewed by a camera with a 1 milliradian IFOV, the resulting pattern on the sensor is known as a blur spot, and its size is determined by diffraction and aberrations in the lens, as well as the finite size of the detector elements in the array. This will be discussed in more detail in a later chapter. A blur spot produced by imaging a point source will typically be 3-4 pixels in diameter with a midwave IR camera system. Because the target is so small in the image, it makes little sense to talk about the radiance of any given pixel in the image of the target. A radiant intensity measurement combines the radiance of the pixels in the blur spot into a single measurement. The radiant intensity measurement also will be found to be less sensitive to the position of the blur spot relative to the grid of pixels. As a blur spot moves across the grid of pixels on a sensor, the count value of any single pixel can fluctuate as the energy moves onto and then off the pixel.

**Note:** It is absolutely critical to enter the correct distance to the target in order to get the correct radiant intensity measurement! ResearchIR does not "know" the distance to the target unless the operator enters the value. For a given blur spot measurement, the radiant intensity returned by the software will increase with the entered distance. This makes sense, because a target has to have a higher and higher radiant intensity to produce the same digital count values at a larger distance. The radiant intensity calculation treats the camera pixels in an ROI like they are all summed together, like a single big detector. The radiant intensity measurement is a measurement of the total power in Watts within an ROI striking the clear aperture of the lens, divided by the solid angle of the clear aperture of the lens as seen from the target.

## Superframing

**Superframing** is a term that refers to high dynamic range imaging with infrared cameras. High dynamic range is very often required with infrared radiometry, particularly in the midwave IR waveband, because the in-band radiance of a target is a strong function of temperature, particularly from ambient terrestrial temperatures (~0-30 °C ) up to typical aircraft target temperatures of >300 °C .

Superframing starts by putting the IR camera into a mode where the integration time is changed from frame to frame, cycling through a set of preset integration times. Depending on the data format selected (SFMOV or ATS files), the images from the presets are either captured into independent video files (SFMOV), or a single file (ATS). These individual data files can be processed into a single movie composed of **superframes**, which are high dynamic range images. Superframing is very useful for applications involving imaging both hot and cold objects and scenes simultaneously. But the temperature dynamic range of a midwave IR camera is limited by the well capacity of the unit cell circuitry in a typical ROIC and the typical noise performance of the sensor. This limitation imposes a radiance dynamic range of about a factor of 18 at a given integration time, meaning the camera can measure a radiance as low as X and as high as 18X. The image will be noisy at radiance values below X, and radiance values above 18X will saturate the unit cell circuitry. A factor of 18 range in measurable radiance imposes limits on the temperature range for a given integration time. The standard Niceville InSb science camera calibration ranges from -20 °C to 350 °C , and that temperature span can be covered by 4 overlapping sub-ranges:

1. *-20 °C to 55 °C*
2. *35 °C to 150 °C*
3. *80 °C to 200°C*
4. *150 °C to 350 °C*

Superframing in Niceville science cameras can be done with 2, 3 or 4 presets. The price of using more presets is speed – the superframe frame rate is much slower than what is possible with a single preset. For example, the X6901sc camera can run at 1004 frames per second for integration times shorter than ~0.9ms. If the camera is set up to run in superframe mode with two presets, the frame rate will be less than half that value, simply due to the overhead of acquiring two sub-frames per superframe. We recommend using the minimum number of presets required for the application.

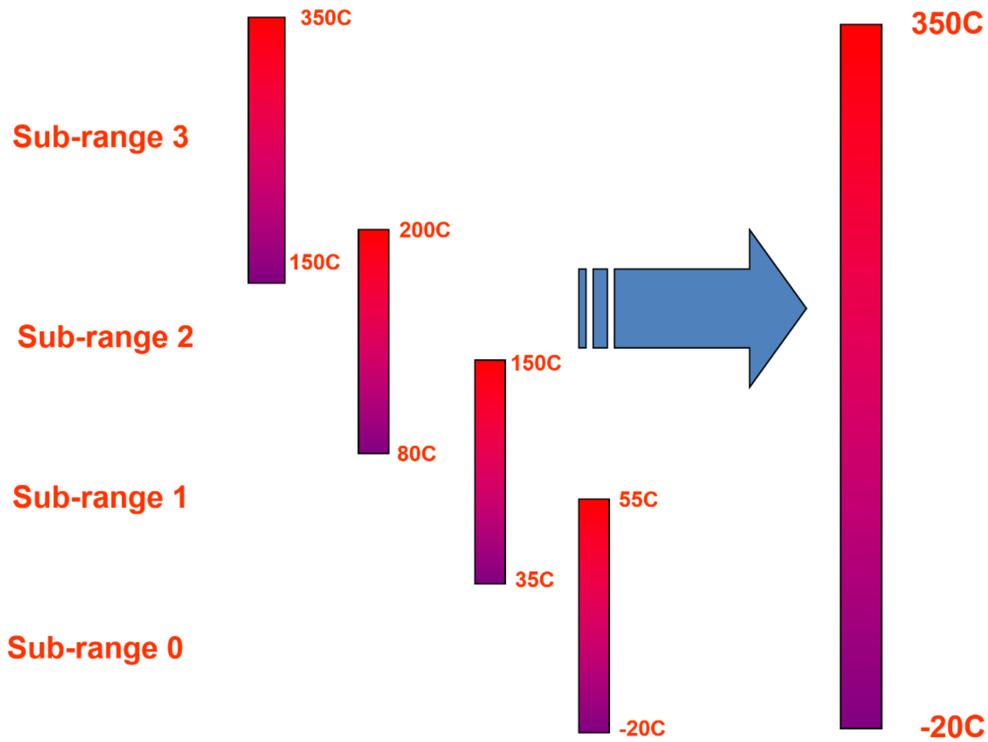


**Figure 36. Diagram of superframing showing sub-ranges for standard 3-5µm camera calibration**

Here is an example of superframing in action: four presets of a Beechcraft King Air running up its engines, captured by a 640x512 InSb camera at the Santa Barbara Airport.

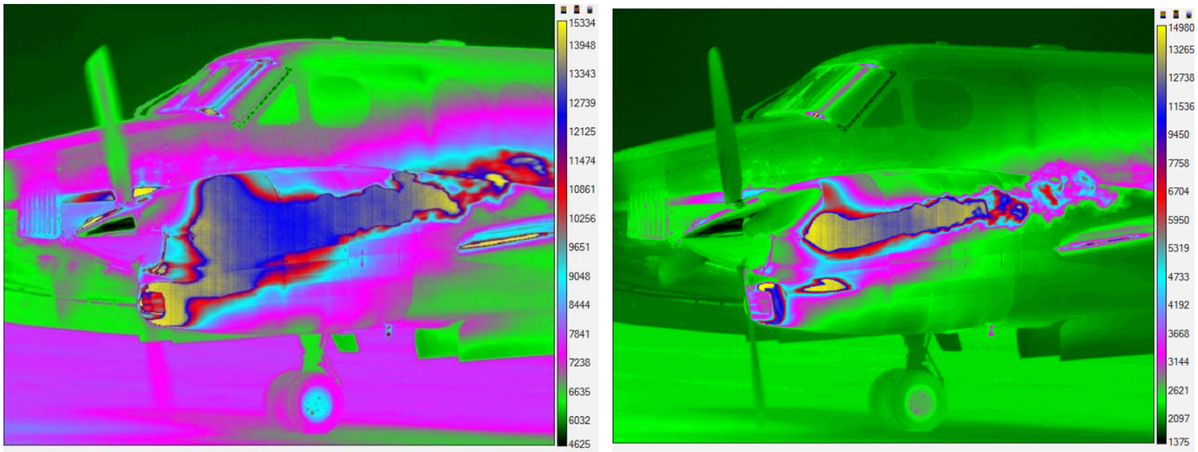

**Figure 37a-b. 2ms and 0.5ms integration time images of Beechcraft. The fixed-pattern noise on the exhaust manifold and plume is due to saturation of the sensor.**

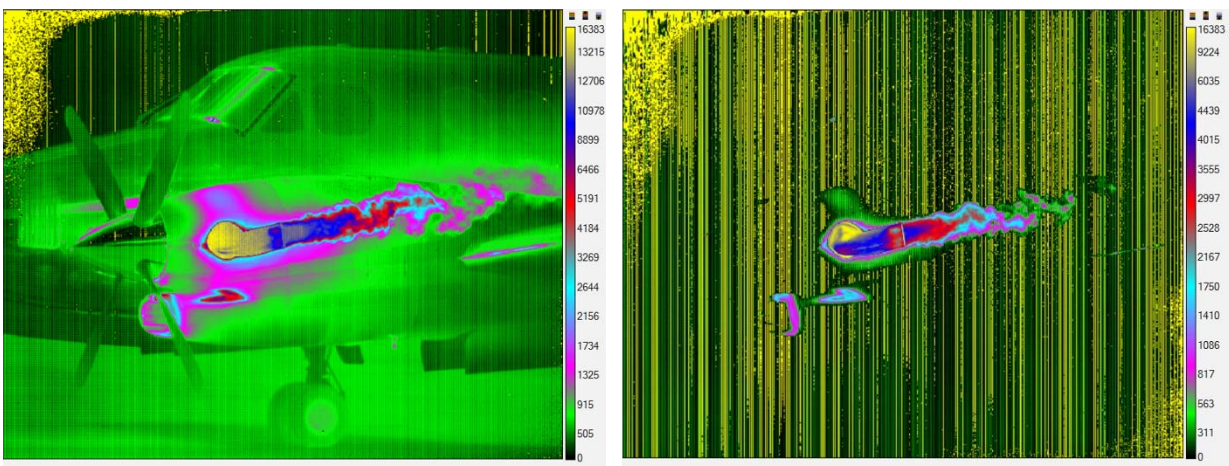

**Figure 37c-d. 0.125ms and 0.03ms integration time images of Beechcraft. The vertical lines and artifacts in the upper left corner of these images are caused by nonlinearity in the ROIC circuitry. These non-linearity artifacts are all but absent from FLIR science cameras produced in the last 15 years with more advanced ROIC designs.**

You can see that most of the pixels in the first two presets look ok – they aren't saturated except on the hot exhaust manifold and plume. Those pixels will be used in the superframe to construct the cooler parts of the image. The saturated pixels won't be used – they are useless for radiometry, other than to give you a lower limit on the radiance of those pixels. In the second two presets, most of the pixels have a very low signal level and those parts of the image are also useless for radiometry, but we already have those parts of the image taken care of, so we don't need them. What we do need are the pixels on the exhaust manifold and plume.

In order for superframing to work, the presets have to have associated radiometric calibrations. Those calibrations can be user-generated, or factory calibrations. Of course, each preset also needs its associated NUC. The superposition of pixels from the sub-frames starts with the conversion of the pixel values into radiance units. The appropriate radiance-valued pixels are then tiled into a superframe. Then the next set of subframes is processed in the same manner. Figure 38a shows the four radiometric transfer functions for these presets plotted on linear axes. The second plot in Figure 38b has a logarithmic radiance axis, which shows the overlap between the ranges more clearly.

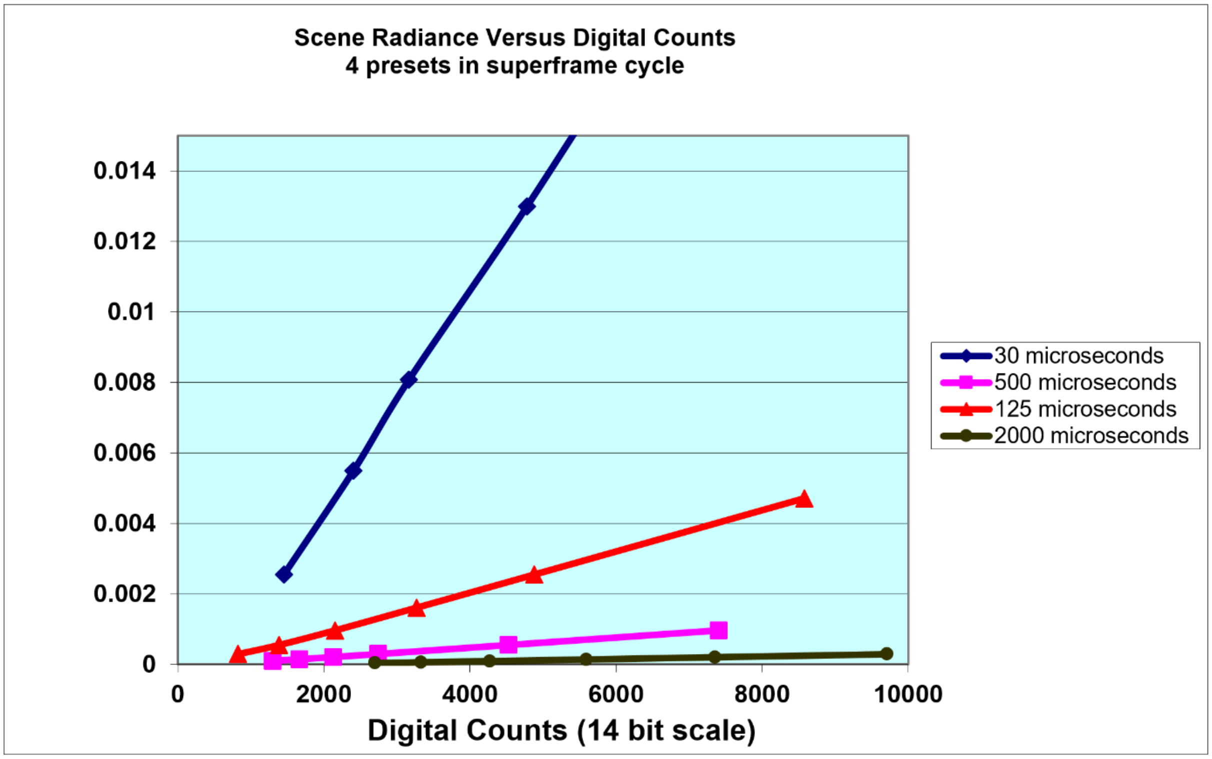


**Figure 38a. Radiometric transfer functions for superframing**

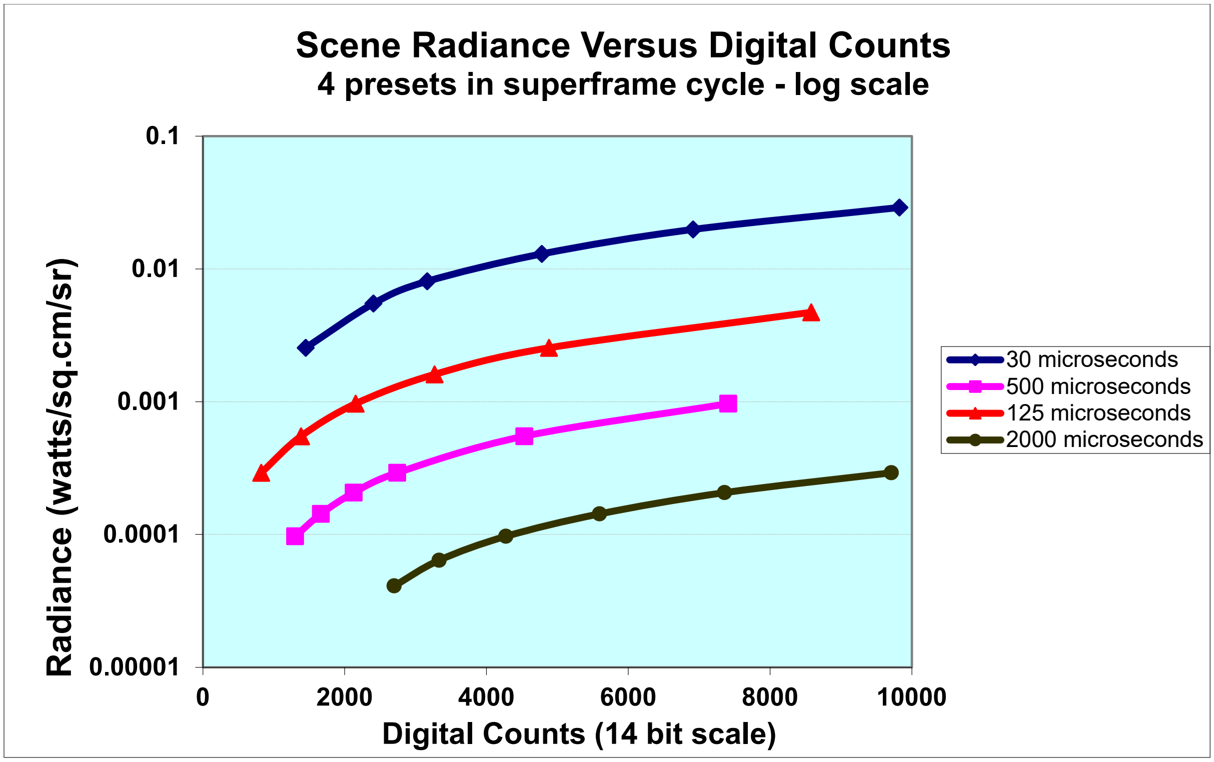


**Figure 38b. Semilog plots of radiometric transfer functions for superframing**

ResearchIR software can process superframing image data into superframes that are expressed in in either radiance units or in temperature units.

Here is a superframe of the King Air in radiance units. The image contains measurements that span a factor of ~100 in radiance, from the cold sky to the hot exhaust pipe. There are some unavoidable artifacts around the exhaust plume that are due to changes in the scene that occur between presets.

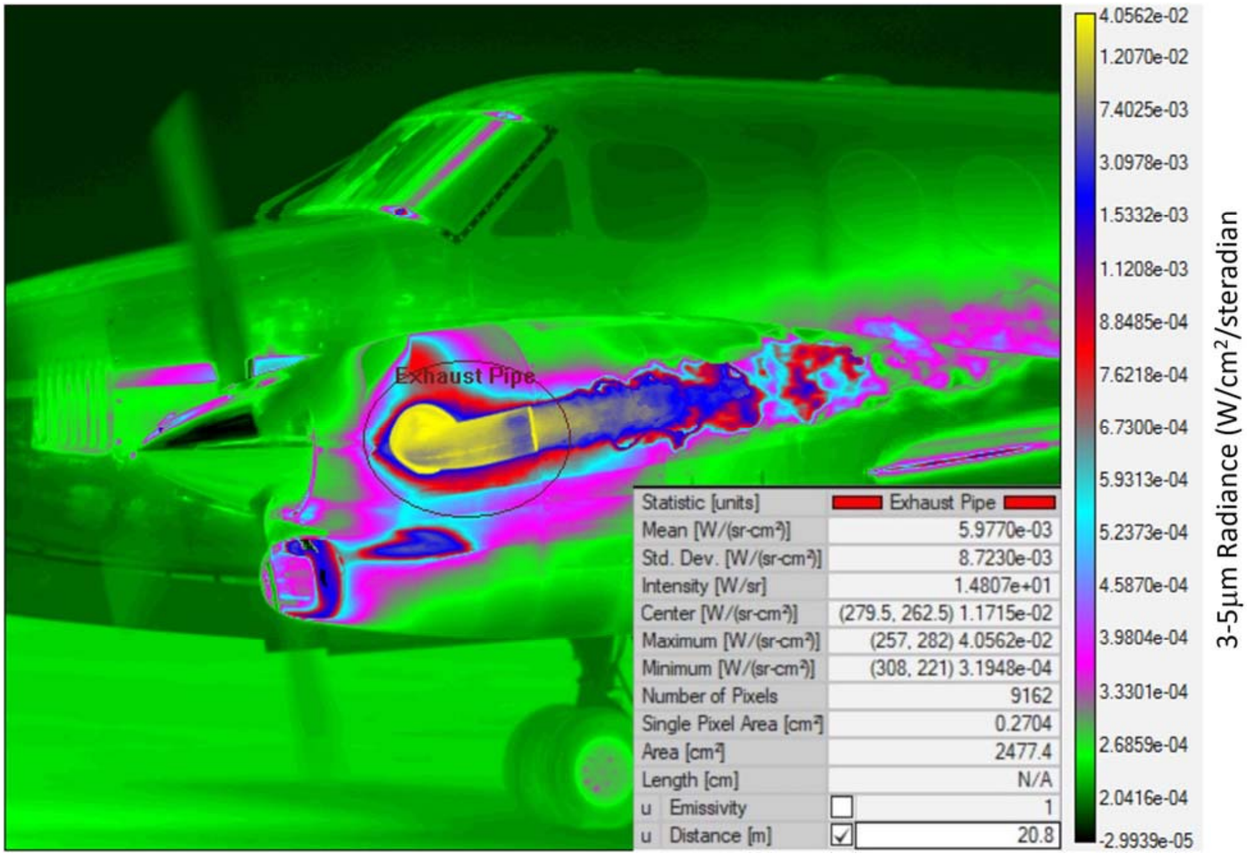


**Figure 39a. Radiance superframe of Beechcraft King Air and ROI radiance statistics**

Figure 39b is a temperature superframe of the same input data. The statistics window shows that the apparent mean temperature of the exhaust pipe ROI is 120 °C . The pixel map in Figure 39c shows what presets were used to generate the pixels in the superframe.

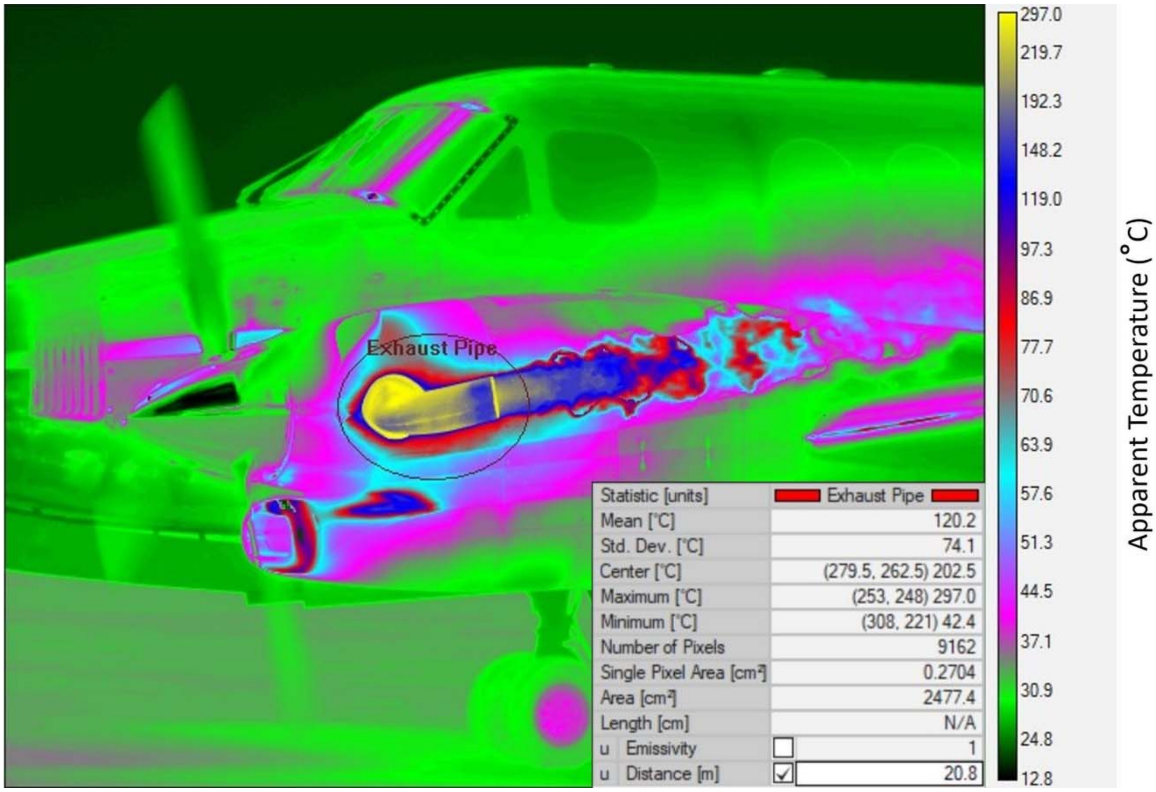


**Figure 39b. Temperature superframe with measurements of exhaust pipe ROI**

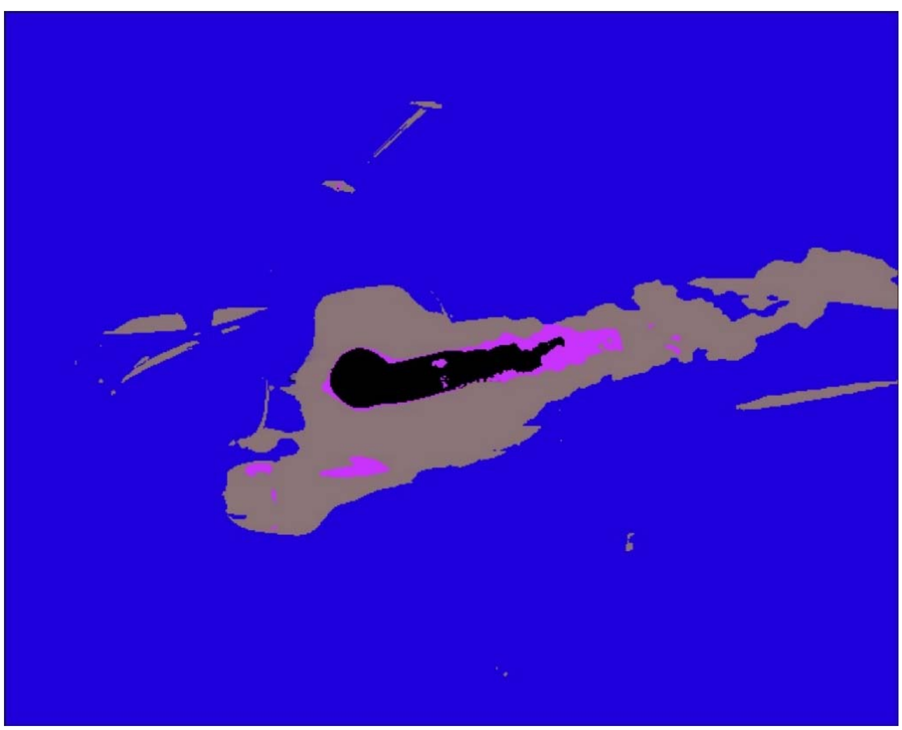

**Figure 39c. Pixel map showing presets used to create superframe. Blue: PS0: 2ms , Gray: PS1 0.5ms , Purple: PS2 0.125ms , Black: PS3 0.03ms**

## Summary

Radiometric calibration of an infrared camera requires that the digital counts for a pixel are linearly proportional to the irradiance that impinges on that pixel.  This is turn ensures that the camera's response to scene radiance is linear.  A radiance calibration requires the use of accurately calibrated blackbodies with known in-band irradiance.  Cavity blackbodies are the best choice because their emissivity tends to be very flat with wavelength and close to unity. Radiometric calibrations are done with a lens installed on the camera.  For best results, the camera and lens should be mated together, and the camera should be turned on and allowed to reach thermal equilibrium before the calibration begins.  The FLIR Niceville cameras incorporate a feature called $T_{Drift}$ that corrects for radiance offsets that arise due to differences in the camera interior/optics temperatures from their values during calibration.  $T_{Drift}$ really helps in situations where a calibrated camera is operated in very cold or very hot conditions at long integration times.  The dynamic range of a camera calibration depends on the waveband of the camera, and for wide ranges of temperature in a scene, it may be necessary to do superframing, which is a cycling of integration times through a set of preset values to build up a high dynamic range data set that can be reduced to a single frame.

# Chapter 2 : How to Make Non-Uniformity Corrections


Introduction ........ 2—67

Why do you need a NUC? ........ 2—67

Facts about NUCs ........ 2—68

What is in the NUC? ........ 2—68

- Gain Table ........ 2—69
- Offset Table ........ 2—72

Result of Applying NUC ........ 2—73

Sensor Linearity and NUC Performance ........ 2—75

- Digital Gain and Offset ........ 2—78

Evaluating a NUC ........ 2—79

Making NUCs with FLIR ResearchIR ........ 2—88

Creating a New NUC in ResearchIR ........ 2—90

- Things to consider before you start making NUCs ........ 2—91
- The NUC Process in ResearchIR ........ 2—96

Saving the Camera State File and the NUC ........ 2—115

NUC Stability ........ 2—118

Summary of NUC Longevity ........ 2—122

PC Side Correction NUCs ........ 2—124

Creating a PC Side Correction ........ 2—126

- PC-Side Correction Pack: ........ 2—132
- Offset Update ........ 2—136

Changing the Focus Setting ........ 2—142

Why NUCs are Made with Multiple Image Frames ........ 2—145

Commercial Blackbody Sources ........ 2—148

- Facts about Blackbodies ........ 2—148
- High-Temperature Area Blackbodies ........ 2—151
- Cavity Blackbodies ........ 2—153

Factory NUC Corrections and Radiometric Calibrations ........ 2—154

Troubleshooting a NUC .......... 2—156
What happens when the NUC Sources are not in the Camera System's Linear Range? .......... 2—159
NUCing over a Smaller Range .......... 2—163
Managing NUCs .......... 2—166
Load NUC Options, Digital Gain and Offset .......... 2—171

## Introduction

The infrared camera term NUC is an acronym for non-uniformity correction. It is commonly pronounced “nuke”, not “nuck”, at least in the US. We say that an image has been “nuked”, not “nucked”. Infrared camera sensors have a high degree of intrinsic non-uniformity for a variety of reasons. This non-uniformity needs to be corrected, not only to make the image look good to the eye, but to be able to extract quantitative data from the image, particularly radiometric data. Figures 1a-b show two images that start with the same raw digital data. The image on the left is uncorrected, and the right image is corrected with a NUC made with laboratory blackbody sources.

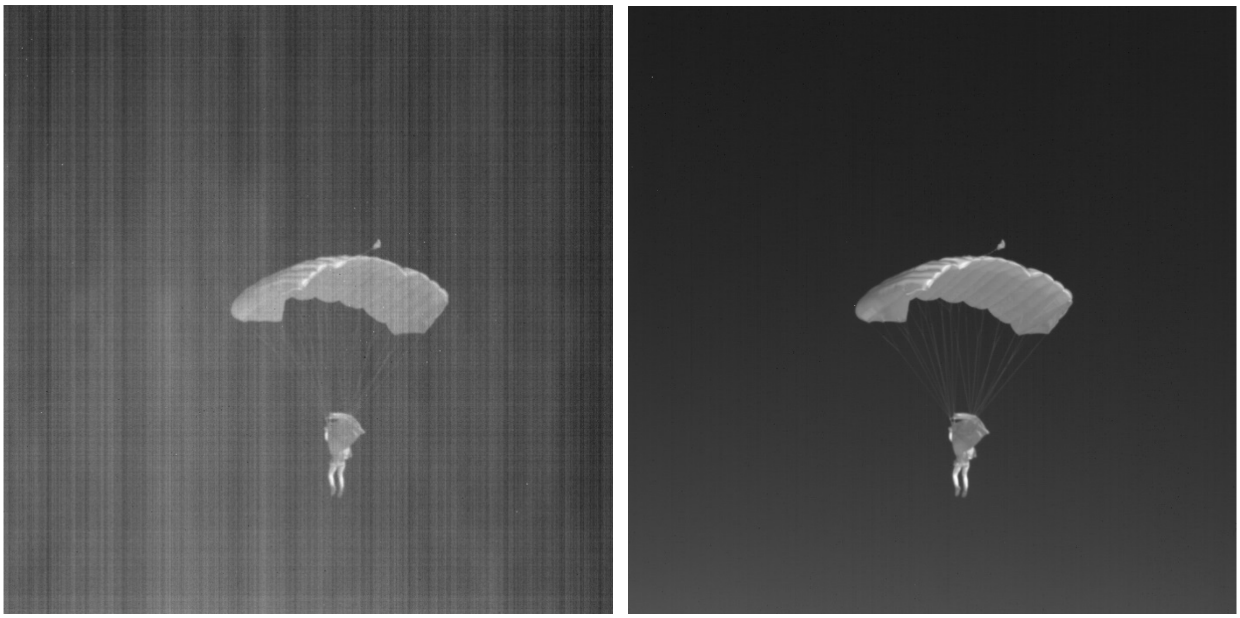

**Figure 1a-b. Uncorrected image (a) and the same image with NUC correction (b)**

The NUC process can be done inside the camera, or in the host PC that is collecting image data from the camera. The former method is far more common, and for the vast majority of IR cameras users, the NUC is something they don’t have to deal with, since it is embedded into the camera. The only aspect of a NUC that may be apparent to camera users is the flat field correction or offset update, as it is often called.

For the user of FLIR science cameras, particularly the ones made in the US, it is possible to create NUCs “from scratch”. These NUCs can exist in the camera, or in the host PC. The image in Figure 1b was NUCed using NUC files that were created in the laboratory after the data had been collected.

## Why do you need a NUC?

- The semiconductor materials that are used for IR detectors are inherently non-uniform in their responsivities and dark current values over the distance scales of focal plane arrays, especially in comparison to silicon detectors. The readout integrated circuits I the focal plane arrays also have non uniformities in their column amplifier offsets.

- You need a uniform camera image to make the images look good to the eye and remove spatial noise, which is very distracting and can hide important details.
- If you are trying to detect very small items of interest in the image, the uniformity of the image must be quite good, or you will not be successful.
- The NUC is also very important for radiometry. The more uniform the image, the less likely you get spatial dependencies on radiance measurements. You ***do not*** want your target to change its apparent radiance as it moves around in the field of view. Always do a good NUC before radiometric calibration!

### Facts about NUCs

- The NUC removes both low spatial frequency "shading" and high frequency fixed pattern spatial noise.
- The NUC does NOT remove temporal noise. Only time averaging does that.

### What is in the NUC?

The NUC is a set of files that are used to apply a linear correction to each pixel in the imaging sensor for a particular integration time. The NUC is a gain and an offset for each pixel. In order for the NUC to work correctly, it has to be applied to pixels operating within the LINEAR range of the sensor. Recent-production FLIR science cameras made in Niceville, Florida are set up so that the digital image output is linear over its entire range, so this is less of an issue, but it is very important to consider for some older FLIR science cameras, particularly those made before ~2008. In those older cameras, the user would observe nonlinearity in the output at the top and the bottom of the digital counts output range, especially when imaging ambient scenes with shorter integration times.

The NUC is most commonly composed of three data files: the gain table, the offset table and the bad pixel table. The first two tables correct for the non-uniformities which are due primarily to optical shading caused by the lens on the camera and the variations in column amplifier offset in the ROIC.[6] The bad pixel table is a table that identifies pixels that are defective in some way. There are various classes of bad pixels, but the most severe bad ones are either short circuited or open circuit, and these pixels will appear as white pixels (digitizer is railed) or black pixels (zero voltage) respectively. There are also pixels that flash on and off randomly; they are called "twinklers" in ResearchIR. The bad pixels are replaced by the nearest good neighbor pixel.

---

[6] There are other secondary effects as well, like the polishing marks on the InSb material, and internal reflections in the camera lens.

### Gain Table

When a lens is put onto a camera with a focal plane array sensor, there is a variation in irradiance across the sensor's area that is caused by geometry. This geometric non-uniformity is circularly symmetric, and it is due to the restriction of incident light rays to a cone called the ray bundle. The restriction is caused by the circular aperture of the optical system. The aperture has a projected area which varies as $\cos(\theta)$, where theta is the angle of incidence of light relative to the optical axis, as shown in Figure 2. For rays coming down the optical axis, the irradiance on the aperture is maximum. At higher angles, the factor $\cos(\theta)$ decreases and the irradiance decreases accordingly. At the detector, the projected area of the detector also decreases with $\cos(\theta)$. Finally, the distance from the center of the aperture to the detector increases as $\cos(\theta)$, so the irradiance (which is proportional to the distance squared) drops off as $\cos^2(\theta)$. The final effect is that the image irradiance falls off as $\cos^4(\theta)$.

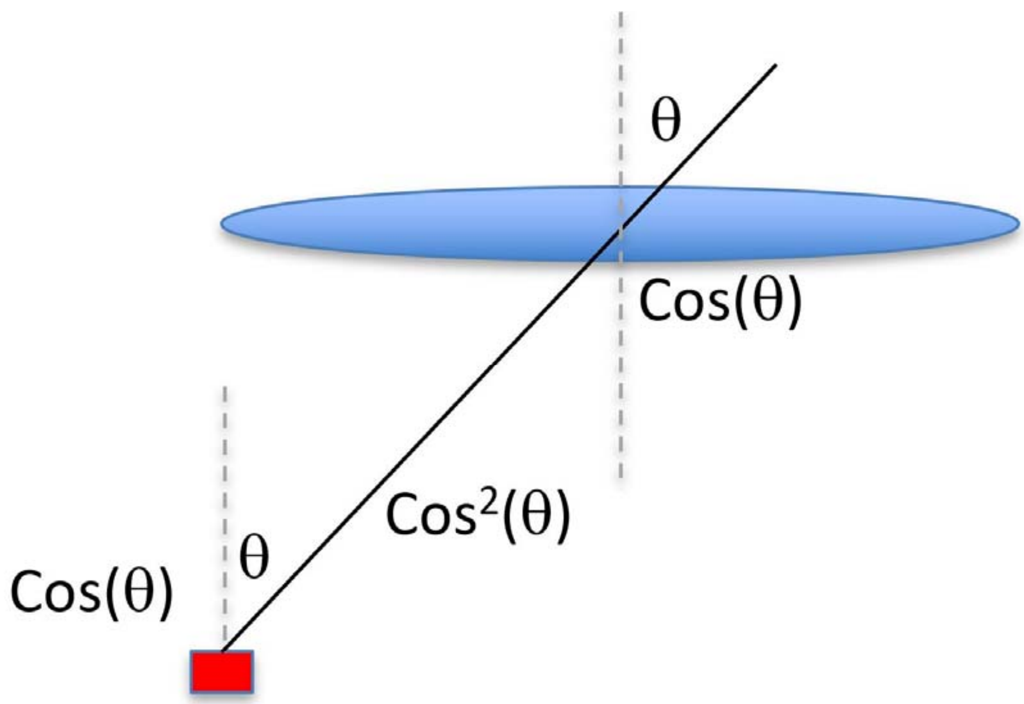


**Figure 2. Diagram of the coldstop (blue) on a cooled infrared camera and a single detector (red), showing the contributions to the $\cos^4(\theta)$ dependency in irradiance on the sensor. The system is not shown to scale and the angle theta is greatly exaggerated.**

This drop-off has the effect of making the corners darker than the center of the image. For a 640x512 sensor with 15µm pixels and a coldstop distance of 20mm, the maximum angle is 17 degrees to the array corner. The plot of $\cos^4(\theta)$ from 0 to 17 degrees looks like Figure 3, and the fall-off of response from the center out to the corners is 19%; i.e. the corners get 81% of the irradiance of the center of the array when the field of view of the camera is filled with a uniform source.

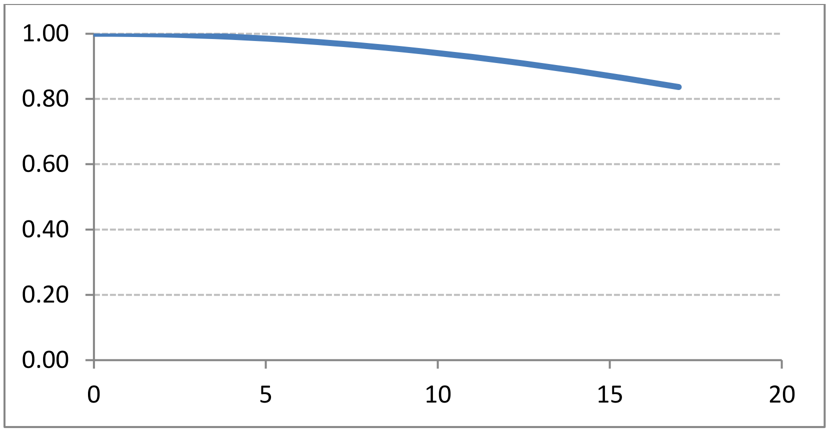


**Figure 3. Plot of $Cos^4(\theta)$ for a 640x512 pixels sensor with 15µm pixel pitch and a coldstop located 20mm away from the sensor**

The gain table values are computed to correct the gain of each pixel to the mean gain value for the array. The gain terms are numbers close to 1 in size, typically. In the case of the 640x512 pixel array discussed above, the corner pixels will have gain values that are greater than 1, which has the effect of brightening the corners. The center pixels will have gain values less than 1, to turn down the brightness of the central pixels. The mean of the gain table is 1; this value will ensure that the image's mean count value is not changed by gain correction. The mean of the offset image is zero, again by definition. The important underlying concept is that applying the gain and the offset should **not** change the mean value of an image of a uniform source over a linear range of source radiance. To that end, every pixel in the sensor is corrected to the mean value of the pixels in the image. If the NUC did anything other than that, it would add (or subtract) counts to the mean and it would "break" the radiometric calibration of the camera.

Figure 4a shows an image of the gain table for a typical NUC table, with Figure 4b showing a line profile drawn across the diagonal. The corner gain values are ~1.2 and the center values are ~0.9. There are a few bad pixels along this line profile which have zero gain. These pixels will be replaced with nearest neighbors so that they won't show up in the final image, where they would be distracting to the viewer and might confuse viewers by mimicking a small object. More importantly, bad pixels can skew statistics, particularly standard deviation values for regions of interest.

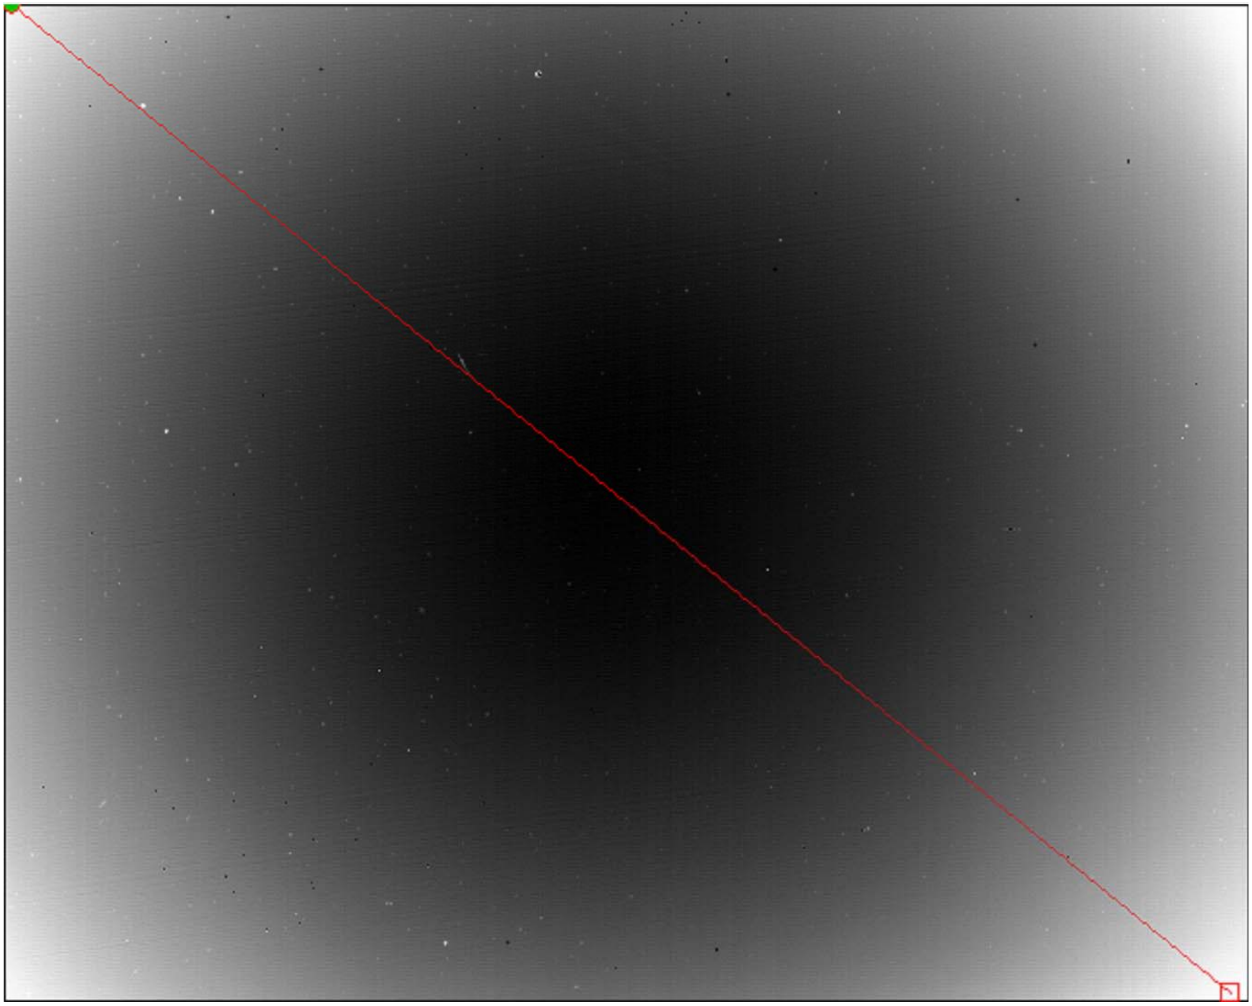

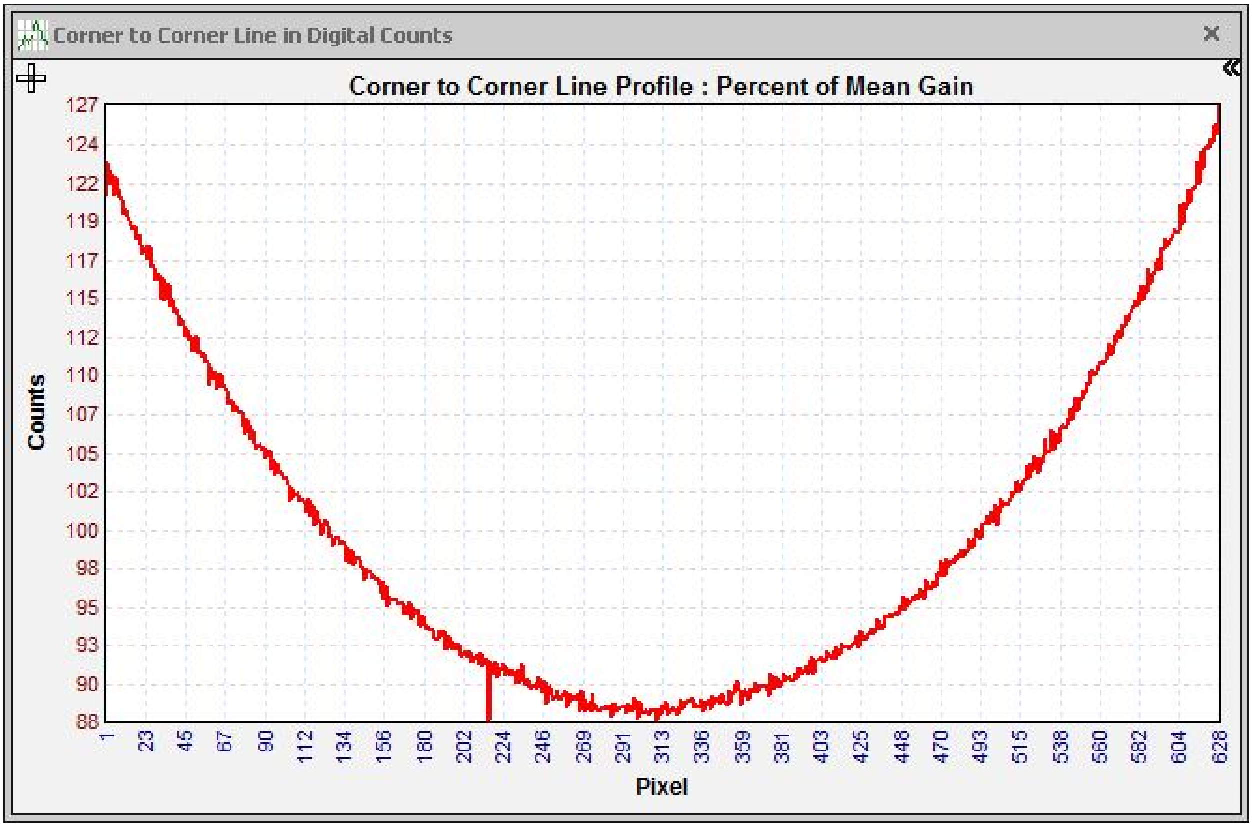


**Figure 1a-b. Gain table from a NUC (a), and a diagonal corner to corner line profile in percent of mean (b). The mean of this table is 1.**

## Offset Table

The gain table is used to apply a gain correction to the raw, uncorrected image data. The gain table values corresponding to each pixel in the image are multiplied by the digital pixel count values across the whole image. Then the digital counts offset for each pixel is added. The equation is as follows:

$$\mathbf{N(i,j) = G(i,j) * U(i,j) + O(i,j)}$$

> Where (i,j) is the coordinates of each pixel in the image, N is the NUCed image, G is the gain table, U is the uncorrected image and O is the offset table.

Just as with the gain table, the Offset table is also an image in its own right, as shown in Figure 2a-b. The offset values in the table are added to each pixel after the gain has been applied. For InSb arrays, the offset table values are typically around ~500 counts in magnitude and are a mix of positive and negative values. The mean of the offset table is 0. The offset corrects the gain-corrected image to its mean level, preserving radiometry by not adding or subtracting net counts to the gain-corrected image mean.

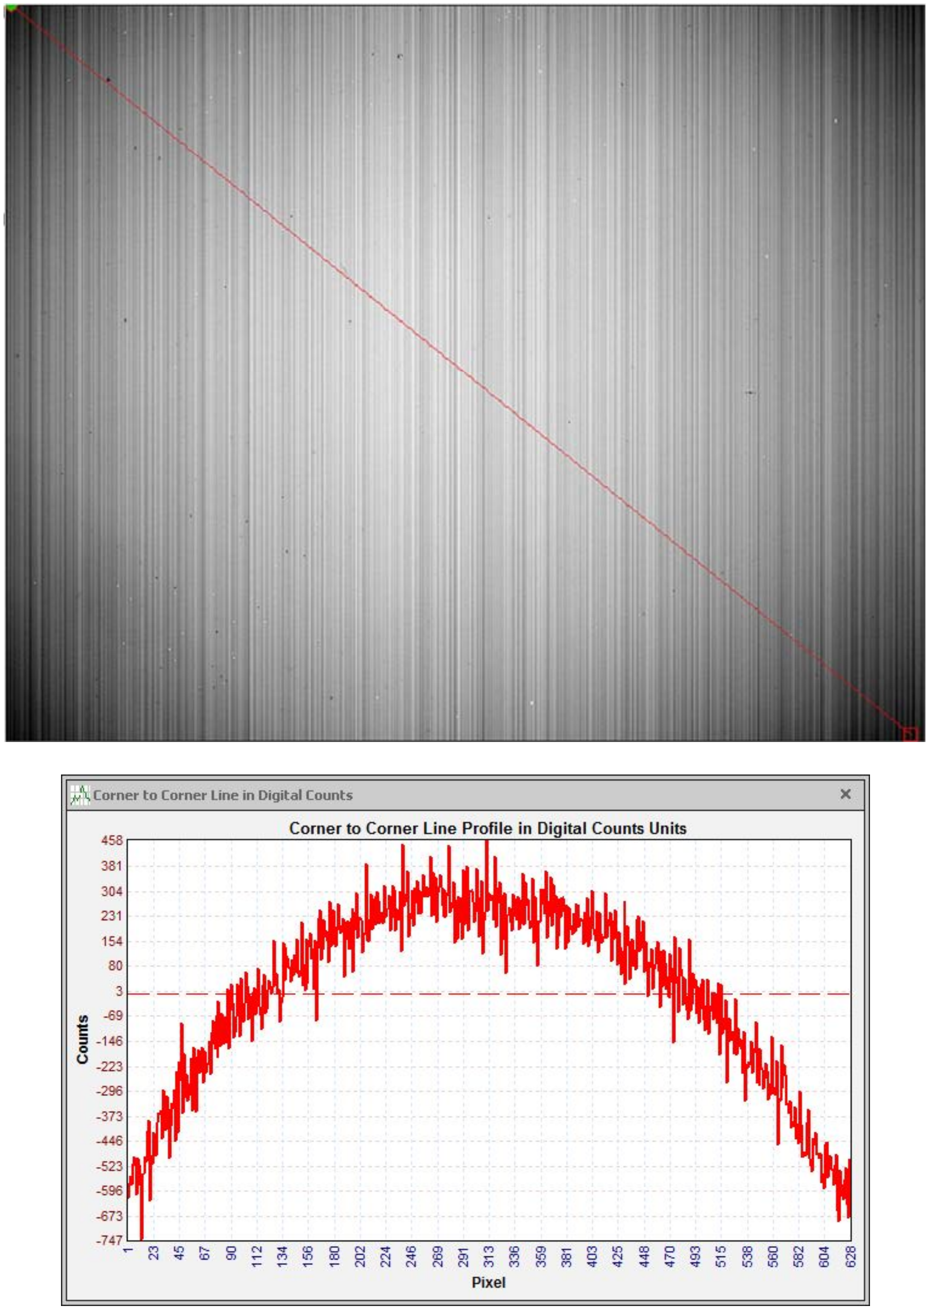


**Figure 2a-b. Offset table for a camera NUC (a) and a diagonal corner-to-corner line profile in digital counts (b). The mean of this table is zero counts.**

## Result of Applying NUC

The preceding examples of gain and offset tables were generated from a dataset that I took with a 3-5μm 640x512 pixel InSb camera. Figure 6 is an image taken by this camera of a 25 °C laboratory blackbody source with no NUC applied, followed by Figure 4, which is the same image data with the gain and offset tables applied. The image mean stays at the same value, 8808.8 counts for both the uncorrected and corrected images, but the standard deviation decreases dramatically from 477 counts to 4.0 counts.

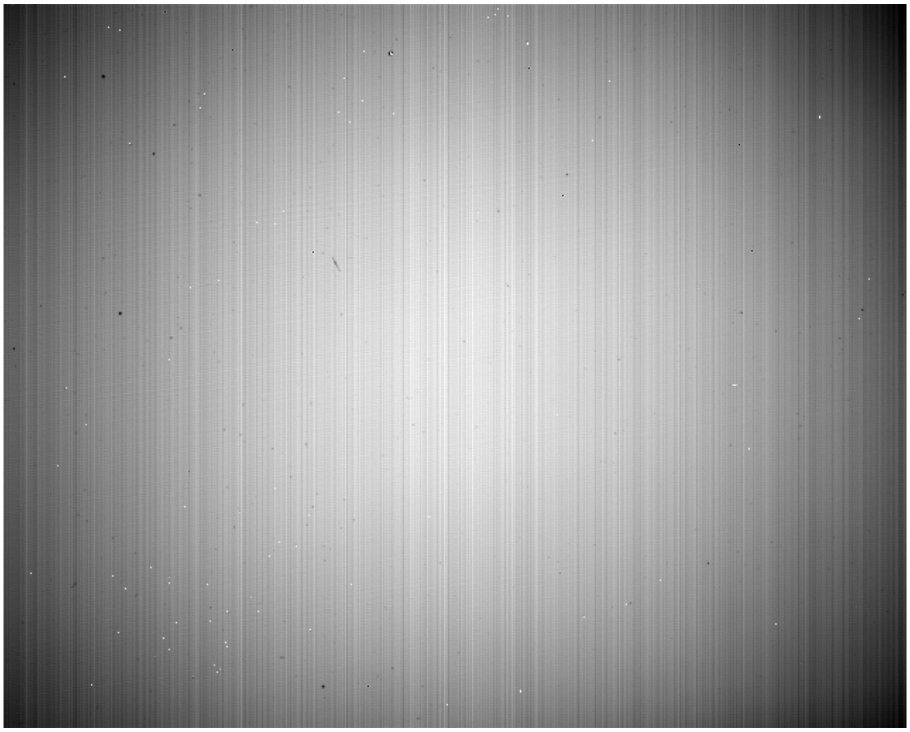

**Figure 3. Uncorrected midwave IR image of a 25 °C laboratory blackbody. Image mean = 8808.8 counts, Image standard deviation = 477 counts.**

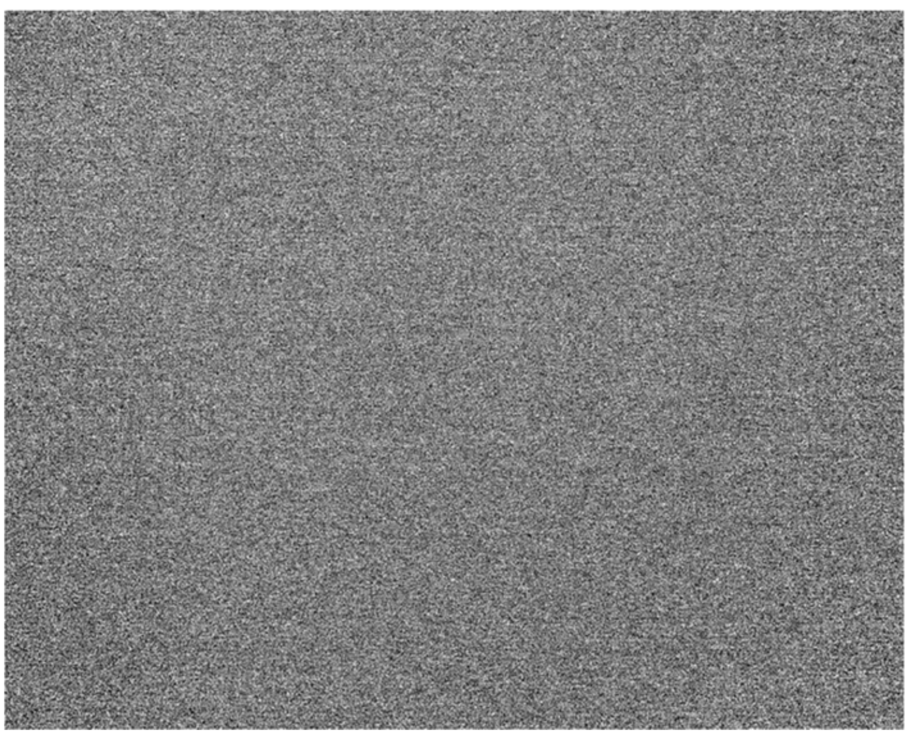

**Figure 4. NUC-corrected midwave IR image of a 25 °C laboratory blackbody. Image mean = 8808.8 counts, Image standard deviation = 4.0 counts.**

A NUC will change the values of individual pixels in the sensor array to make them all have the same (or nearly the same) count values when the camera is pointed at a uniform source. All the pixels will track together as the source radiance changes, because both pixel gains and offsets are corrected.

## Sensor Linearity and NUC Performance

**In order for the NUC to work properly, the imaging sensor must be exposed to scene radiance values within the sensor's linearity limits AND the NUC gain and offset tables need to be generated with two sources that are both within the linear range of the sensor.**

These linearity limits are determined by the sensor design. On newer FLIR science cameras made in Niceville, Florida, the digital output data is linear over the entire accessible range of the output. When the NUC is turned off, the digital image data output varies from about 500 to 1000 counts all the way up to 16,383 counts, which is the maximum value for a 14-bit digitizer[7]. On older FLIR science cameras, the limits are usually 5% to 80% of the well capacity of the integration capacitors in the ROIC unit cell circuits, or a digital count range of about ~1,000 to ~13,000 counts. It is generally fine to use an ambient-temperature NUC source if the integration time of the camera is long enough that the ROIC integration capacitors will have some "fill" to them when the camera is exposed to this source. But at short integration times, an ambient temperature NUC source may not get you much response above very low well fill. The NUC may then not work well for higher scene temperatures. In both cases, the "hot" NUC source must swing the digital counts above the "cold" source by at least a thousand counts. The following radiometric measurements demonstrate this.

Figure 5 is an example of the radiometric transfer function for a FLIR SC6000 science camera operating at a short integration time of 47 microseconds. Note the non-linearity at the top and bottom of the curve. The labels on each point are the temperatures of the blackbody sources used to obtain these points. This is an older ROIC design (the ISC0309), and the non-linearity is worse at the bottom of the well compared to newer ROIC designs used in FLIR-US science cameras in current production.

[7] 16,383 is 2 to the 14$^{th}$ power minus 1.

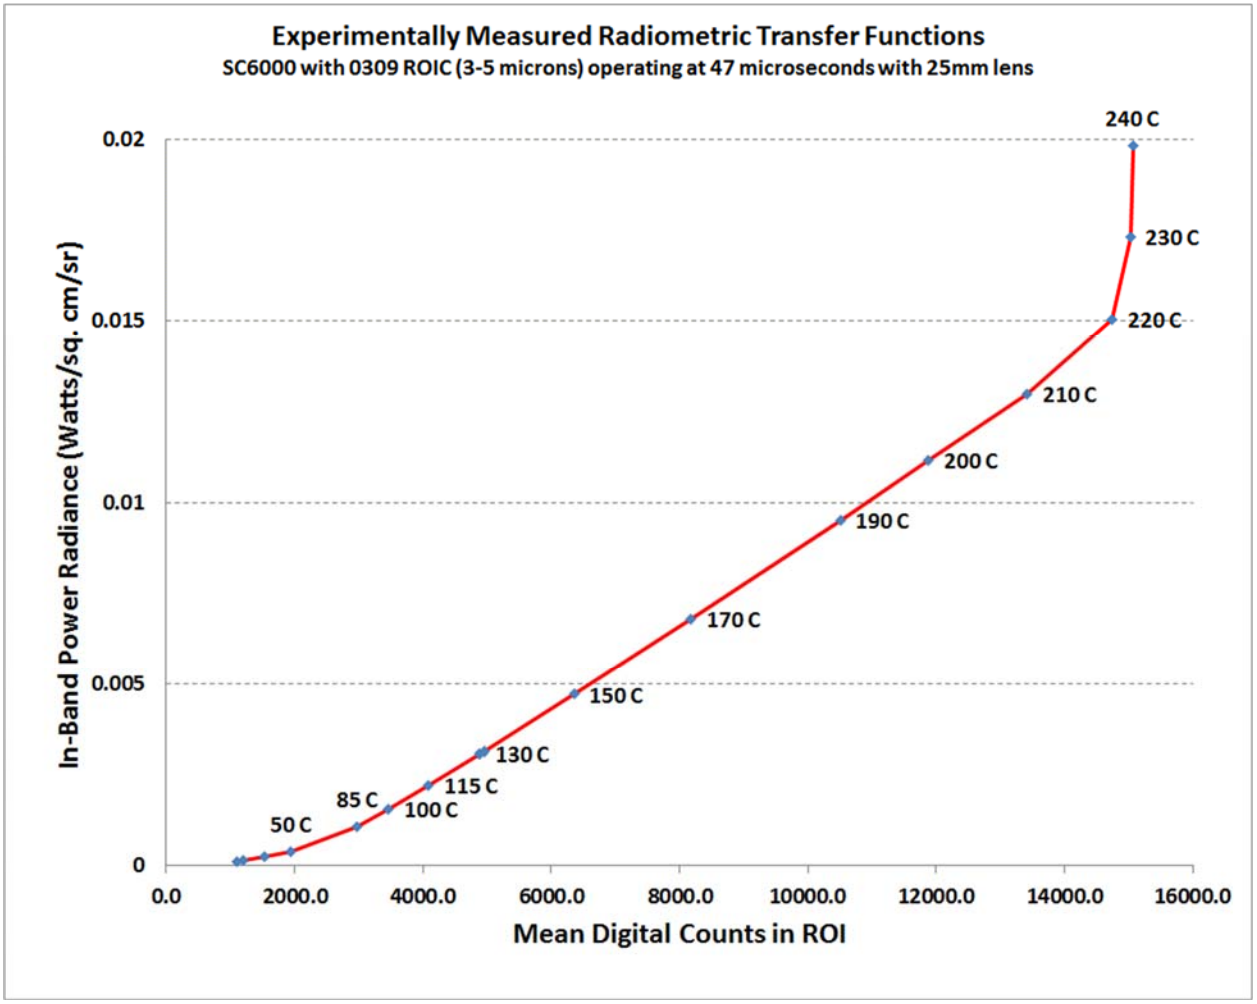


**Figure 5. Radiometric transfer function for SC6000 operating at 47 microseconds**

In order for a standard two-point NUC to work properly on this SC6000 camera at this short integration time, the graph tells us that this camera will need to be operated between ~3400 counts and ~13,000 counts, which is the linear range for the ISC0309 ROIC, as shown in Figure 6. This corresponds to about 20% to 80% of the well capacity. The equation in the box (with the $R^2$ value of 0.9998) is the radiometric transfer function, or RTF. By only using the RTF points between 100 °C and 210 °C, the goodness of fit is very close to 1.

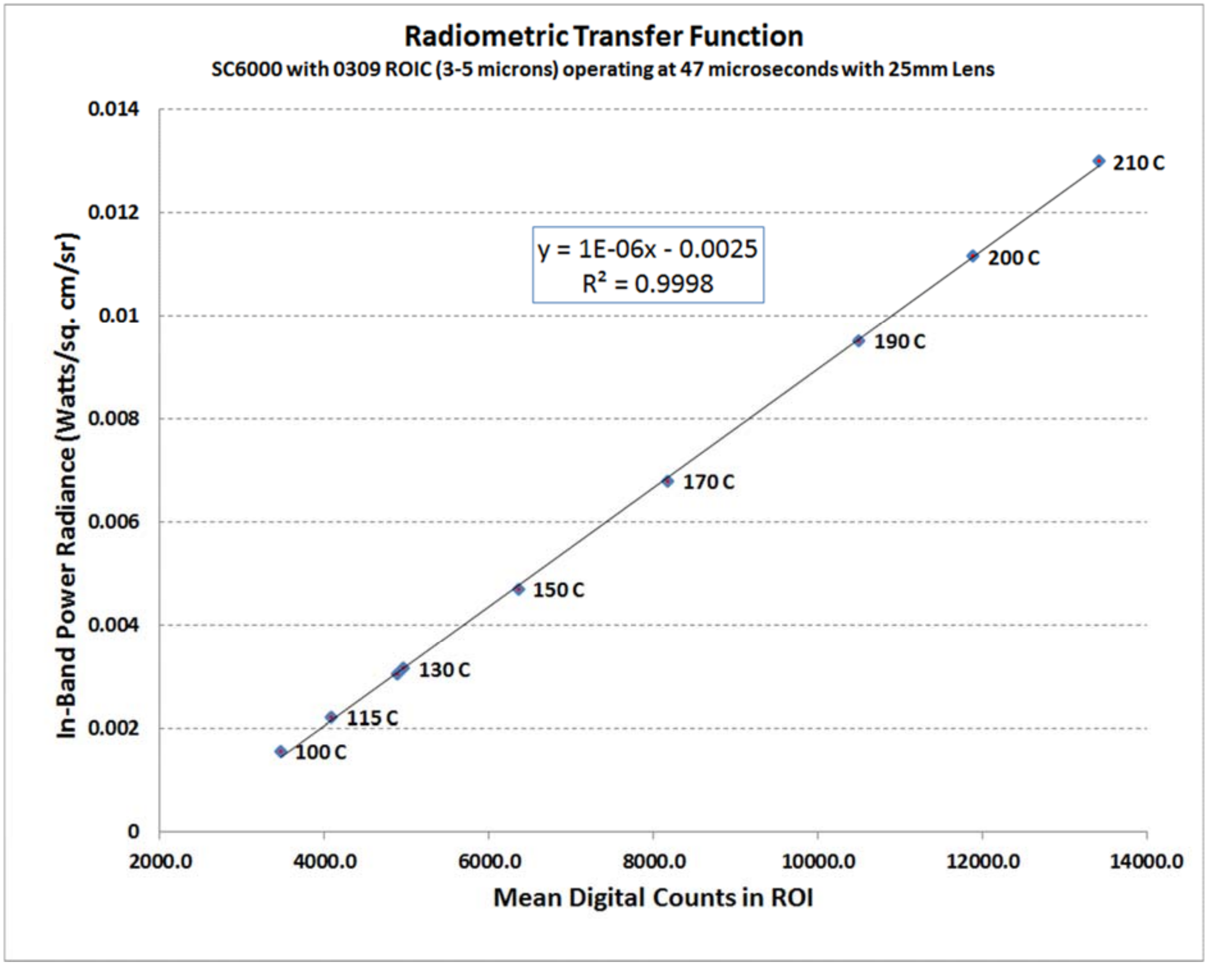


**Figure 6. Linear region of Radiometric Transfer Function for SC6000 operating at 47 microseconds**

The two NUC sources need to produce digital counts within this (20-80% well fill) linear range, or else the NUC that is generated will not work well! The linear region for this camera system configuration starts for a blackbody emitter at 100 °C, which is the first point at ~3500 counts in Figure 6. This value of 3500 counts is about 20% of the dynamic range of the camera. The NUC will work well up to a target temperature of 210 °C or 13000 counts, which is about 80% of the dynamic range. The 210 °C and 100 °C scene temperatures represent about a factor of 8 in power radiance. The two NUC sources should be 100 °C and 210 °C to get the maximum signal swing. In a pinch, a 100 °C source and a 130 °C source would probably work, although there might be more non-uniformity than in the 100 °C/210 °C example.

**If you use an ambient temperature blackbody as the "cold" source for a two-point NUC under these short integration time conditions, the NUC will not perform well.**

It is well worth determining the linear range of your camera system by measuring the radiometric transfer function. You can measure it using laboratory blackbodies and the CalibratIR utility included with ResearchIR.

Knowing the linear range will enable you to craft your own NUCs that meet the linearity criteria. The linear range of the ROIC is determined by the exposure, and the range holds for different integration times, so you generally only have to measure it once. The exception to this is extremely short integration times (<10 microseconds), where the linear portion of the digital counts range will become limited in range. Experimental data on a FLIR SC8303 camera operating at 0.5 microseconds showed linear behavior up to a count value of ~4600 counts, which was reached for a 600 °C scene temperature. Higher scene temperatures increased the count value, but a plot of scene radiance versus digital counts became highly non-linear.

## Digital Gain and Offset

An important factor in this linearity measurement is that the digital counts that one gets out of the camera for a given ROIC well fill will depend on the settings for digital gain and offset, which are global camera settings that apply to all the presets. For most applications, these two parameters should be set to 1 and 0 respectively. In the case of the graph above, they were set that way in the SC6000 camera. These parameters are accessed in the camera control GUI and are saved in the camera State file. In the example screenshot in Figure 7, the digital gains are all set very close to 1 and the offsets are small. These particular values are responding to the $T_{drift}$ correction algorithm in the camera, since this particular camera has a factory calibration in it.

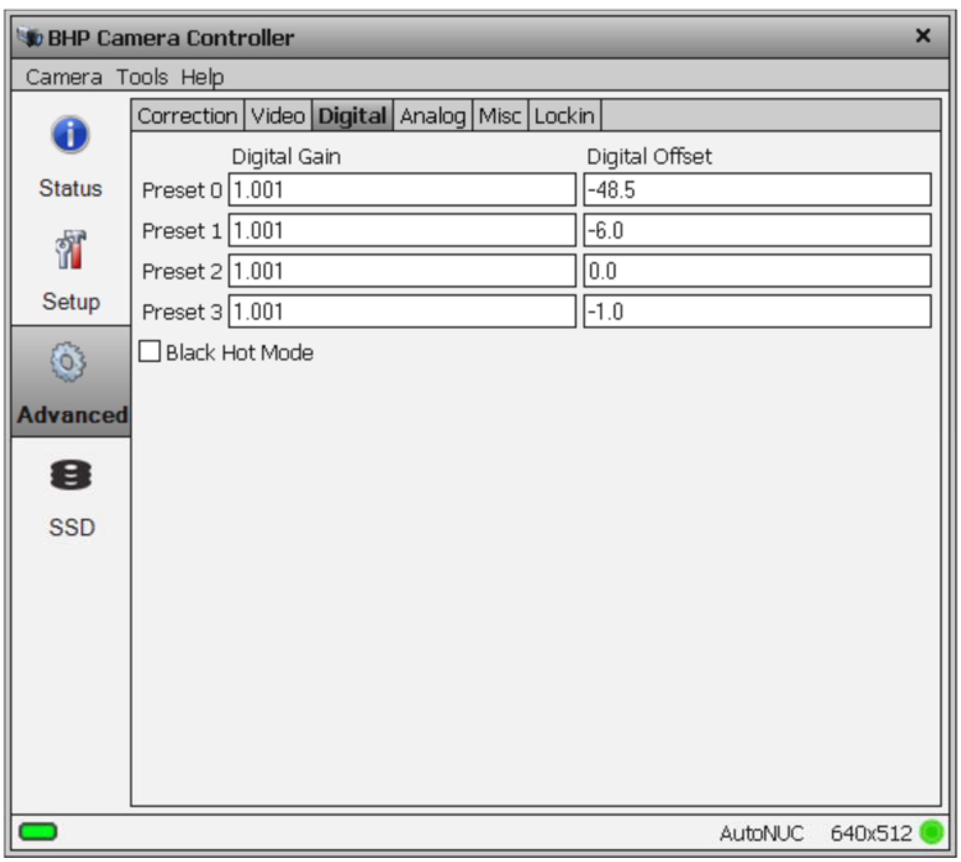


**Figure 7. Digital gain and offset table for presets in an SC6000 camera with a Factory calibration**

A good conservative rule of thumb is that you should be within 20% and 80% of the well capacity of the camera to get good linear performance, which is vital to the success of the NUC's correction, and for accurate radiometry. This 20-80% range equates to mean digital count values of about 3,000 to 13,000 counts for the FLIR-US science cameras, both new and older units. The newer camera that are using newer ROIC designs can be used from ~500 counts to 16,383 counts with good linear performance. This count range corresponds to a factor of ~20 in power radiance. Again, these are rules of thumb only, and the best course of action is to use blackbodies to measure the radiometric transfer function for your camera under the conditions you want to use it (correct integration time, lens, warm filters, etc.), then stay within the linear range of this RTF.

The end of this chapter has some examples of what happens when you generate a NUC using sources that are not in the linear range of the camera system.

## Evaluating a NUC

The most commonly used figure of merit for a NUC is the ***corrected uniformity*** of the camera when viewing a uniform radiance ("flat") scene. The corrected uniformity is the standard deviation of an image in counts divided by the mean of the image in counts.

The standard way to calculate the corrected uniformity is as follows:

1. Take a sequence of 64 frames after the NUC is performed, while the camera is pointing at a uniform blackbody.
2. Then the frames are averaged to reduce temporal noise. This can be done in MATLAB. ResearchIR also has a filter function which has a frame averaging option. Up to 16 frames can be averaged. I prefer more frames than 16, but it works.
3. Then compute the standard deviation of the time-averaged image and the mean count value of that time-averaged image.
4. Then divide the standard deviation by the mean to get the corrected uniformity.

Figure 8 is an image of the clear sky with no NUC applied. There is a line profile drawn across the diagonal of the image.

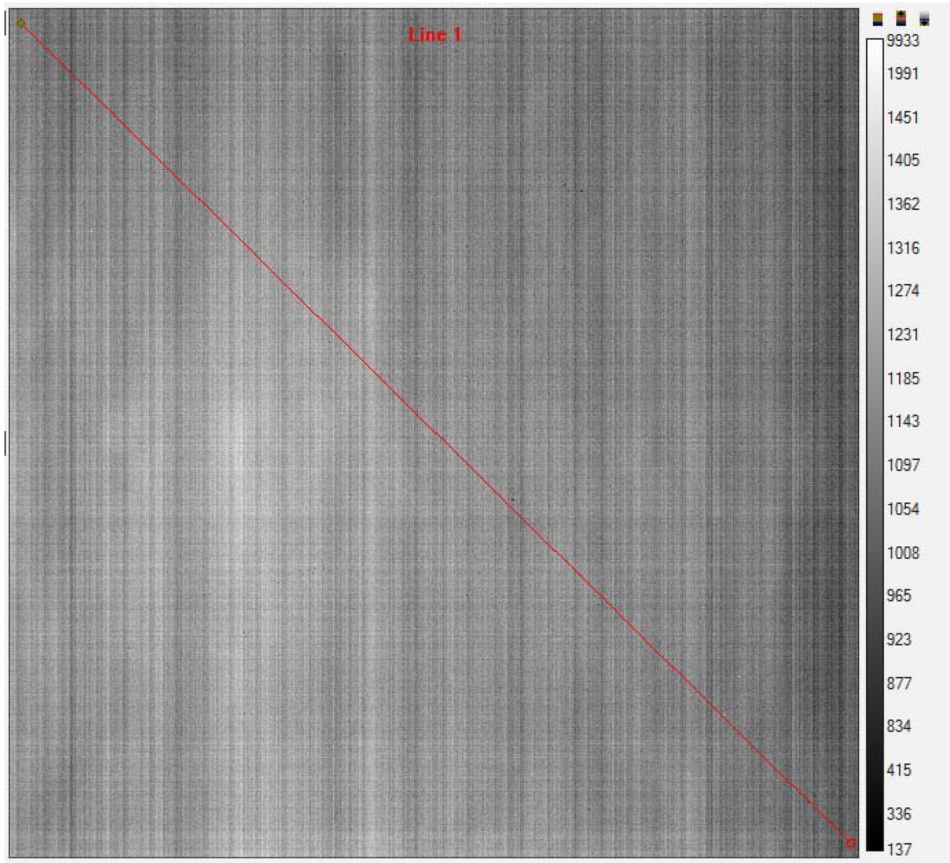


**Figure 8. Uncorrected midwave IR image of the empty sky**

The line profile is a good way to evaluate the NUC. You can see there is about 400 counts of nonuniformity from one corner to the other in Figure 9:

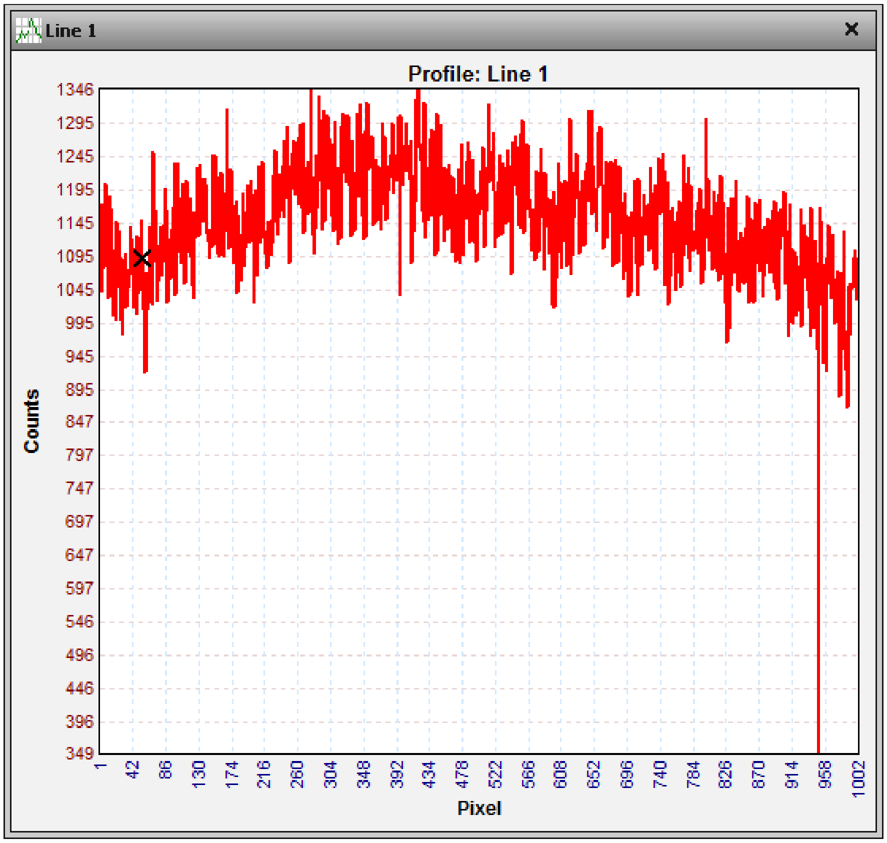


**Figure 9. Corner to corner line profile of uncorrected sky image**

The corrected uniformity is 0.078. This is calculated from the image statistics shown in Figure 10. The standard deviation value is 90.2 counts and the mean value is 1152 counts, and 90.2/1152 = 0.078.

Stats

| Statistic [units] | Image |
|---|---|
| Mean [counts] | 1151.8 |
| Sum [counts] | 1207775323.0 |
| Std. Dev. [counts] | 90.2 |
| Center [counts] | (511.5, 511.5) 1159.0 |
| Maximum [counts] | (460, 564) 9933.0 |
| Minimum [counts] | (991, 556) 137.0 |
| Number of Pixels | 1048576 |

**Figure 10. Image statistics for uncorrected sky image**

Figure 11 is an image of a parachutist made with the same camera close to the same time as the clear sky image, also with no NUC applied. The image quality is poor because of all the column-to-column spatial noise.

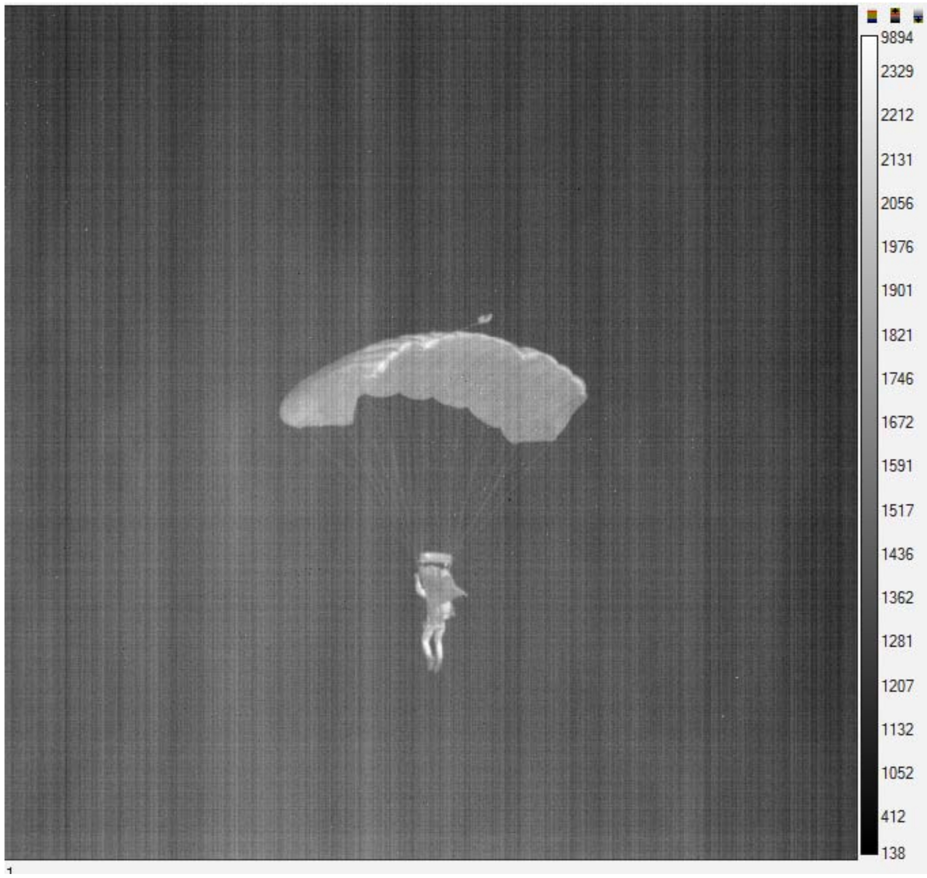


**Figure 11. Parachutist imaged with SC8000 camera - no NUC is applied.**

Now I will apply a “hillbilly” NUC to the empty sky image back in Figure 8 – the result is in Figure 12. The NUC was done in the field with an ambient temperature calibration plate for the first source, and the palm of my hand for the second source. My palm is not nearly as uniform as a laboratory blackbody and the low frequency spatial noise is distracting.

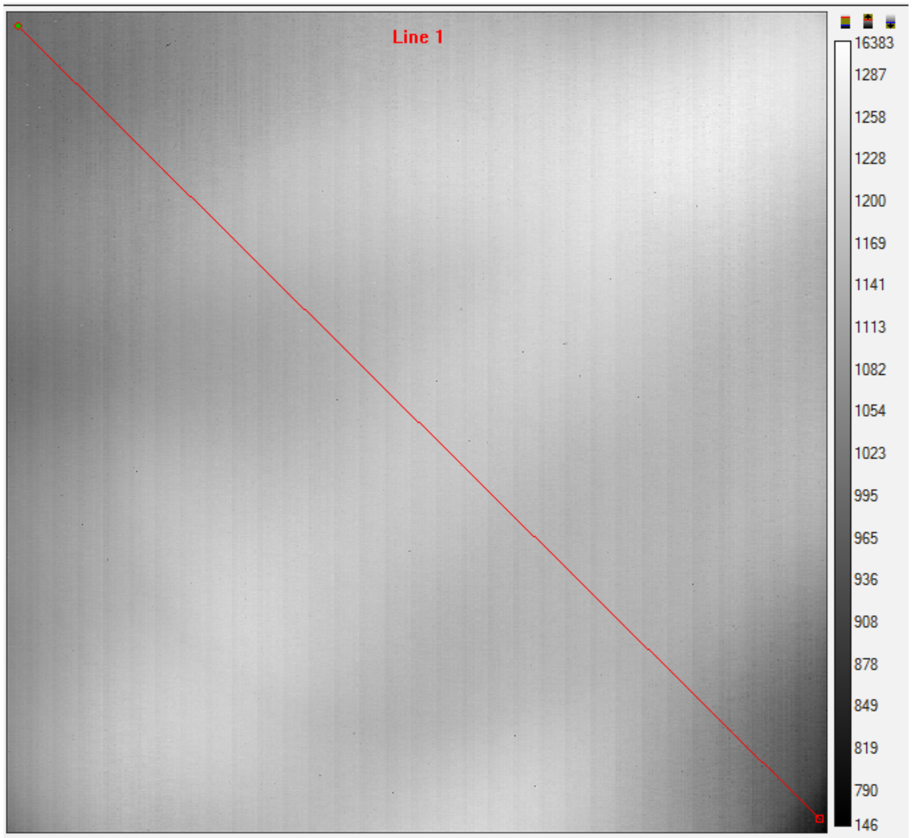


**Figure 12. Sky image with NUC made from the palm of my hand as the hot source and an ambient temperature calibration plate for the cold source**

Still, the image is better than with no NUC. We note that the high spatial frequency noise has been reduced significantly, but there is still low frequency spatial noise (which is called **shading**). You can also see a faint circular pattern in the image – this is probably a way-out-of-focus reflection of the cold stop off the skin on my hand. The non-uniformity is from my palm not being uniform in its surface temperature and having some reflectivity. The line profile in Figure 13 shows that the high spatial frequency noise (column-to-column variation) is now much less pronounced.

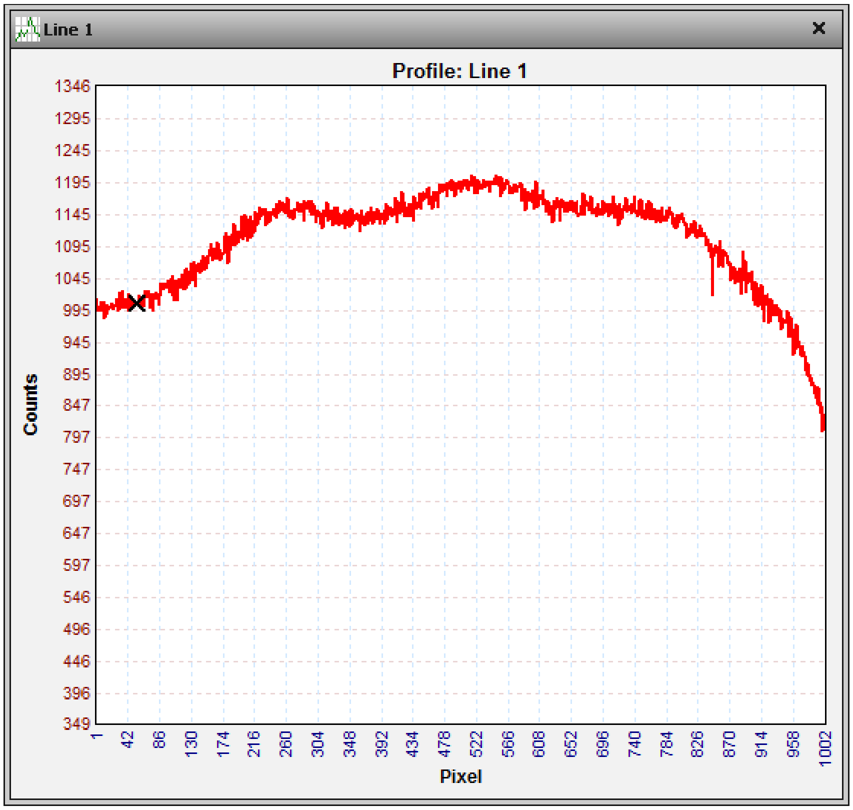


**Figure 13. Corner to corner line profile on sky image with "hand-made" NUC**

The corrected uniformity is now 0.067, or 77.4/1152. The statistics are shown in Figure 17.

Stats

| Statistic [units] | Image |
|---|---|
| Mean [counts] | 1151.9 |
| Sum [counts] | 1207856095.0 |
| Std. Dev. [counts] | 77.4 |
| Center [counts] | (511.5, 511.5) 1187.0 |
| Maximum [counts] | ( 60, 11) 16383.0 |
| Minimum [counts] | (162, 573) 146.0 |
| Number of Pixels | 1048576 |

**Figure 14. Image statistics for image corrected with "hand-made" NUC**

The parachutist image in Figure 18 is more uniform looking with this handmade NUC applied to it, but it could be better still.

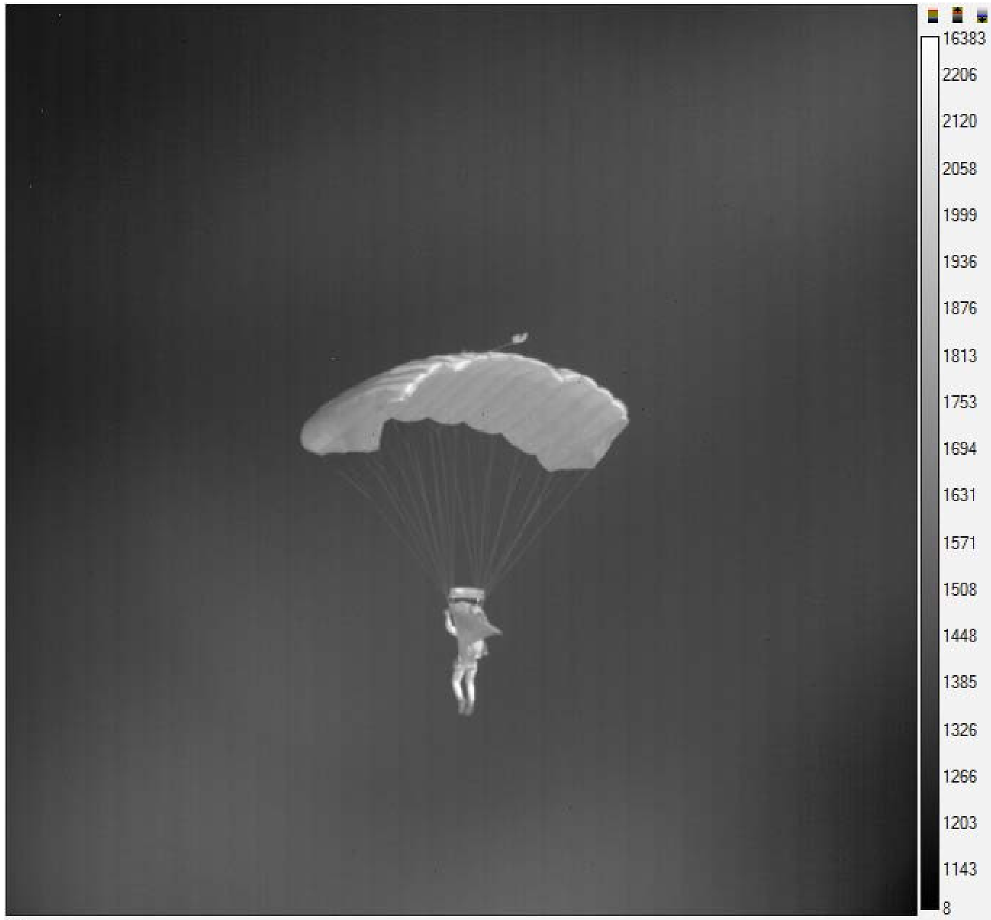


**Figure 18. Parachutist image with field NUC applied. The NUC was done with an ambient temperature calibration plate and the palm of the author's hand.**

Now I will apply a laboratory NUC to the empty sky image, as shown in Figure 19a. The NUC was performed using a high-quality laboratory blackbody source on the camera a week after the image data was acquired.

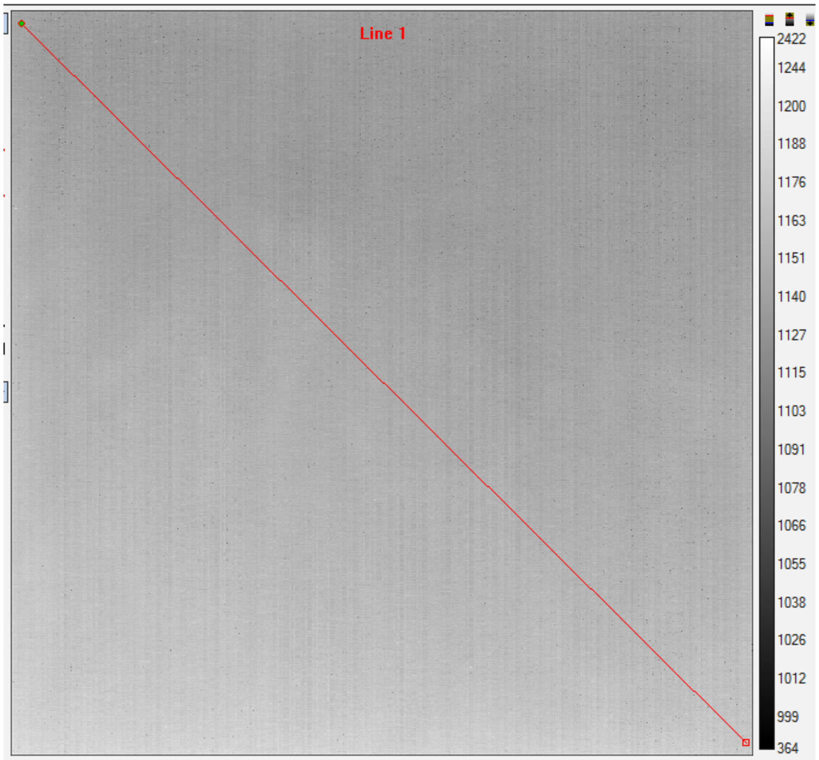


**Figure 19a.  Sky image with laboratory NUC**

Note the empty sky image is slightly darker at the top compared to the bottom, as shown in the slight slant of the line profile in Figure 19b:

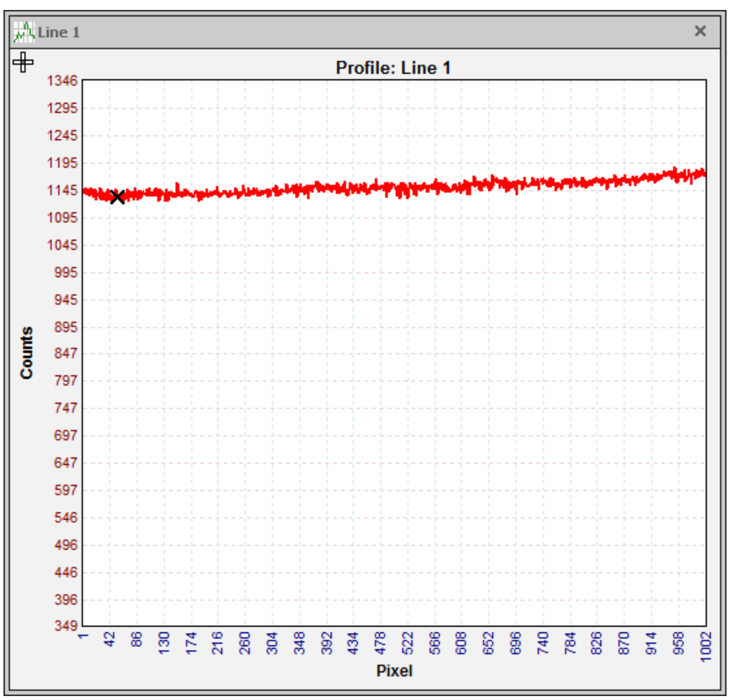


**Figure 19b.  Corner to corner line profile on sky image with laboratory NUC**

The Corrected uniformity is 0.014, or 15.6/1151, as calculated from the image statistics in Figure 15. The residual gradient from top to bottom is real – the clear sky below zenith does not have uniform IR radiance in elevation. The closer you get to the horizon, the longer the air path, and the higher the radiance.

**Stats**

| Statistic [units] | Image |
|---|---|
| Mean [counts] | 1150.6 |
| Sum [counts] | 1206504338.0 |
| Std. Dev. [counts] | 15.6 |
| Center [counts] | (511.5, 511.5) 1149.0 |
| Maximum [counts] | ( 40, 119) 2422.0 |
| Minimum [counts] | (220, 879) 364.0 |
| Number of Pixels | 1048576 |

**Figure 15. Statistics for laboratory NUCed image of sky**

The parachutist image looks good in Figure 21. It is hard to quantify the performance of a NUC with something complex in the scene, but the human viewer knows that the NUC is good. All NUCS should be evaluated using a very uniform area blackbody to get a real quantitative evaluation.

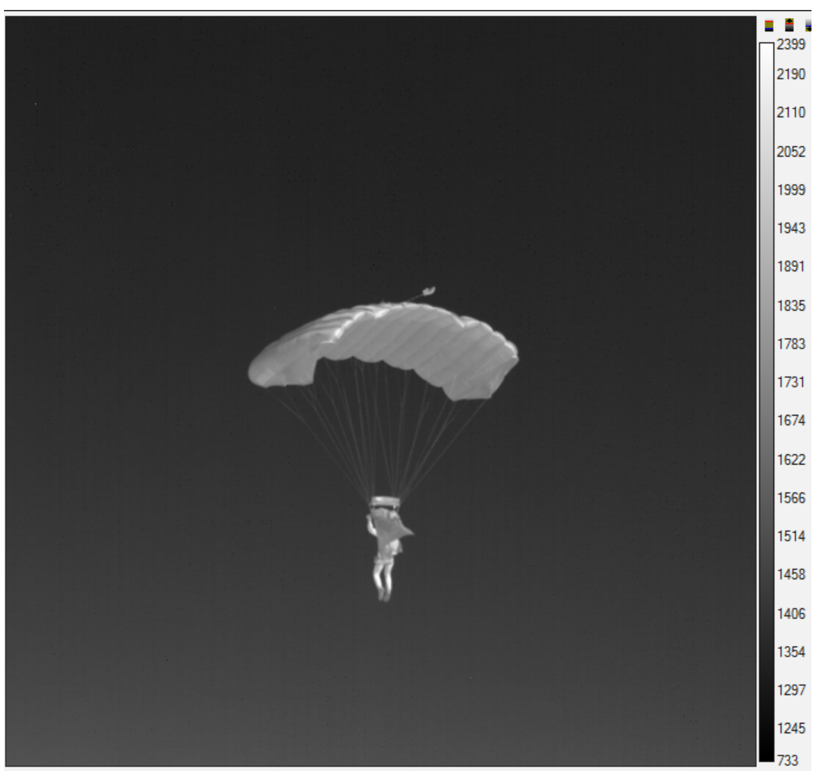


**Figure 21. Parachutist image with laboratory NUC applied.**

## Making NUCs with FLIR ResearchIR

FLIR ResearchIR is set up so that users can make their own non-uniformity corrections using uniform sources of infrared radiation. These NUCs can live inside the camera, or in the host PC.

FLIR science cameras from the Niceville facility in Florida all come with a set of standard NUCs inside. When the camera is connected to ResearchIR, and the camera controller is opened, you can see what the status of the camera is with the Status window, as shown in Figure 22. This shows the integration time of the camera, among other things. The integration time is intimately linked to the NUC. The NUCs are created with a particular integration time setting, and in order for them to work properly, the camera must be set to the right integration time when the NUC is loaded. This is automatically done for in-camera NUCs, but is not automatic for PC Side Corrections, which are NUCs made by the user that live on the host PC, not in the camera. This camera is operating at an integration time of 40ms.

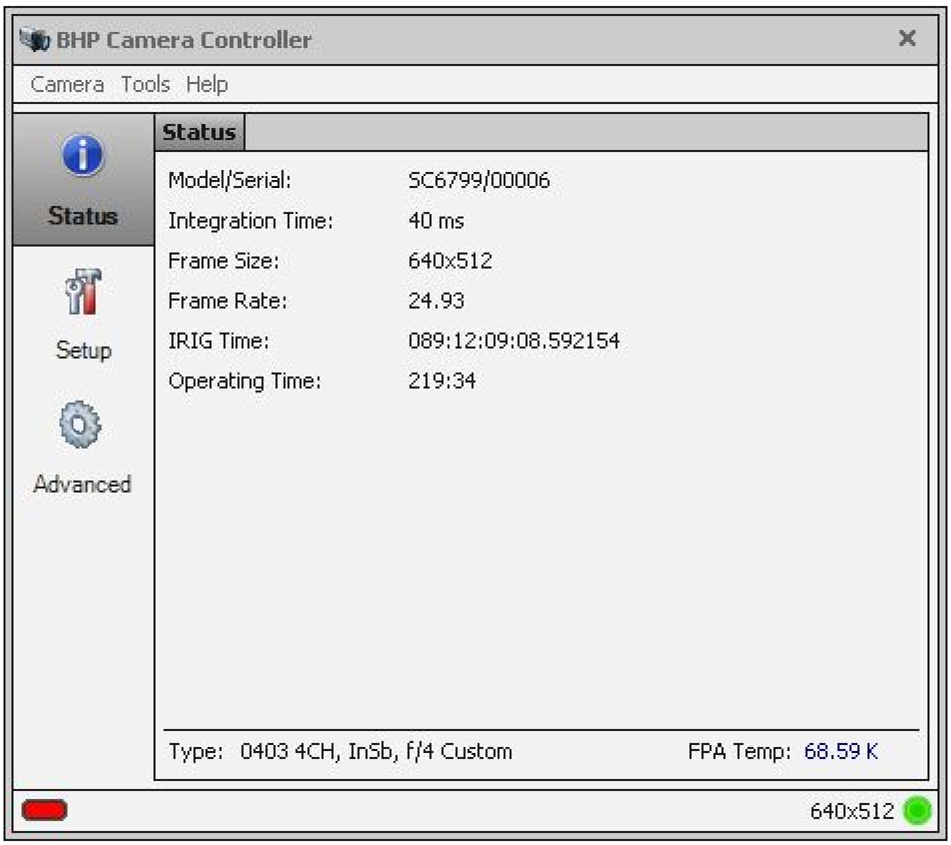


**Figure 22. Status window for SC6700 See-Spot at power up**

The Setup window (Figure 23) confirms that the camera is in Preset 0 and that the integration time is 40ms. Preset 0 is the first on the list, and the radio button next to it is indicating that that is the active preset with a blue dot. There are four available presets in these cameras.

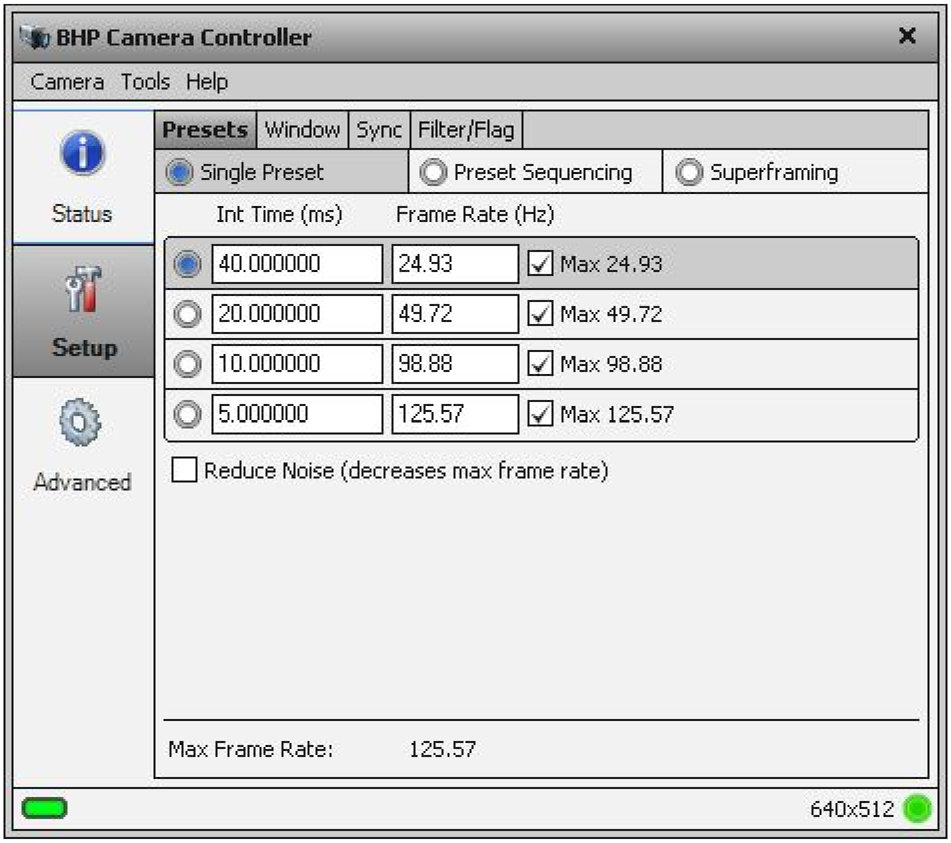


**Figure 23. Setup window with Presets tab showing.**

Figure 24 is a screenshot of the control GUI in ResearchIR showing the factory NUCs that came with this camera, a FLIR SC6700 See-Spot:

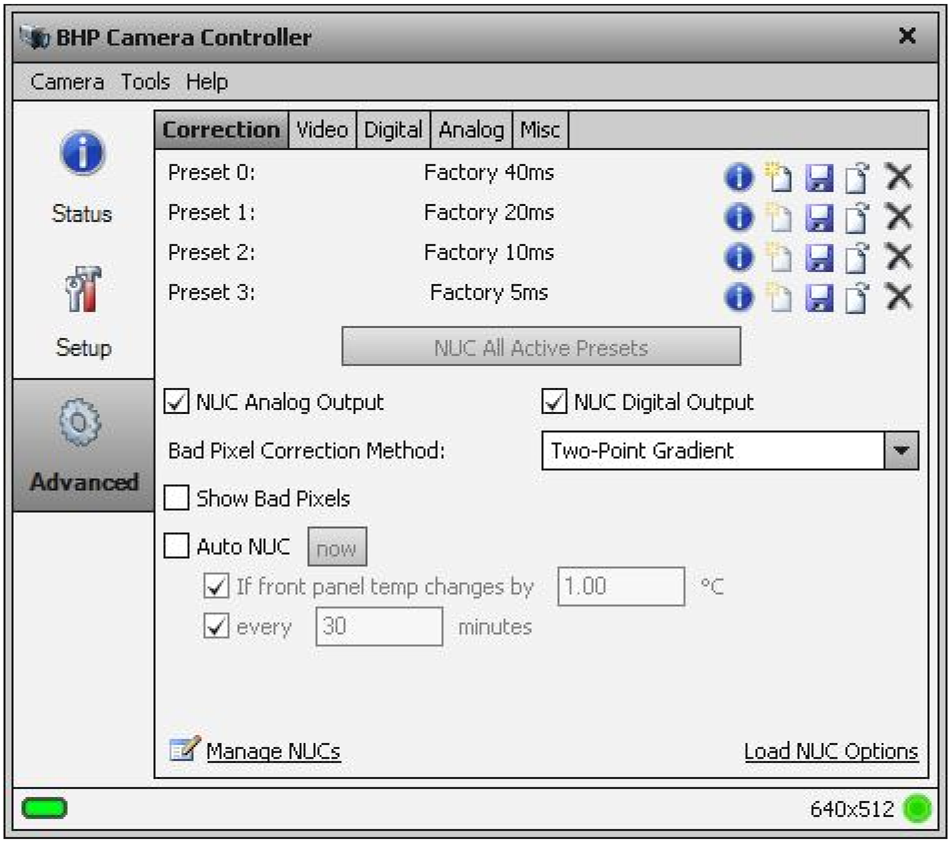


**Figure 24. Correction tab of camera controller in ResearchIR**

These NUCs have integration times that follow a typical pattern of variation in factory NUCs: Preset 0 is 40ms, then Preset 1 is 20ms, and so on, decreasing by a factor of 2.

Factory calibrated science cameras come with both radiometric calibrations and factory NUC tables inside the camera. The factory NUCs are done by skilled technicians using a wide variety of high-end laboratory area blackbody sources. These sources are calibrated on a regular schedule so that the NUCs are done with repeatable radiance levels. Replicating a set of NUC sources like this is expensive, and many customers will prefer to use factory NUCs. For those customers that have the blackbodies, and that want to do their own NUCs through their own choices of lens (and optional filters held behind the lens or in a filter wheel in the camera), the following is a step-by-step process for creating a NUC.

## Creating a New NUC in ResearchIR

There are two ways to create NUCs using ResearchIR software. There are NUCs that are stored in the camera itself, and there are NUCs that are stored on the host PC running ResearchIR. There are advantages and disadvantages to each. Some comments:

- Most users will use the factory NUCs that come with the camera if they are just trying to get a good image, are not trying to do radiometry and did not purchase a set of factory calibrations.

- Most users that are trying to measure temperature will purchase factory radiometric calibrations that are paired up with factory NUC files. They should use the factory NUCs.

- For power users that are also creating their own radiometric calibrations and want the greatest degree of control over the data, I recommend considering the use of PC Side NUCs, as they are called.

- For users that are not doing radiometry, user-created or factory camera-side NUCs are very convenient since they stay with the camera. This is very helpful when the asset is shared with other users.

- Some older FLIR cameras may not have NUCs loaded in them anymore, or the NUCs may have gotten corrupted. The user can create new NUCs using an area blackbody that can be stored in the camera (if that is supported) or on the host PC.

## Things to consider before you start making NUCs

1. If you are going to change the camera configuration by creating new NUCs, you will want to save the camera State file so that when it is next powered up, you can put it back into the same state that you left it in. So, for instance, the new NUC you will create will be the one that is loaded into a preset, and that preset will be the active one. The state file is easy to save. You go to Camera/State/Save As, and you will get this window (Figure 25):

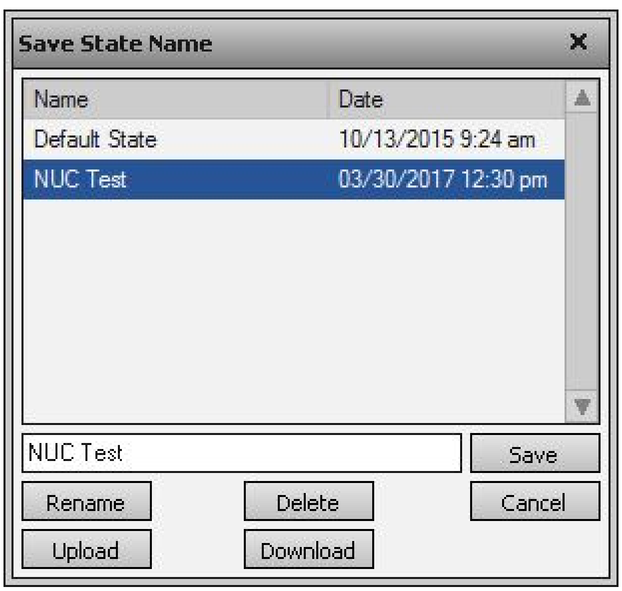


**Figure 25. Camera state Save As window**

I typed in the name NUC Test and hit Save. When you are finished with the NUC you are about to create, and the camera is in the configuration you want it in, then go to Camera in the camera controller and Save the state again. You can download the state file to the host PC for posterity, so that in case it gets overwritten later, you can have a clean copy to upload and restore the camera state. The state file has a *.rsc extension. It is a good idea to use the same name as the state file name in the camera controller. Here I have downloaded it to the host PC (Figure 26):

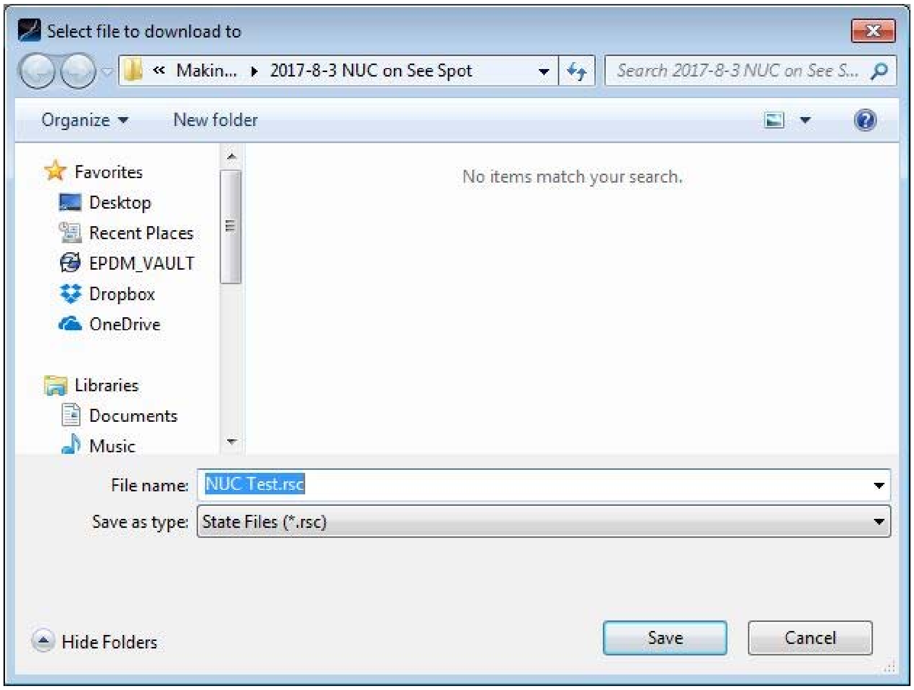


**Figure 26. Downloading the state file to the host PC.**

2. Set the camera clock to the correct time and date!! This is done in the Tools menu in the camera controller window. I set the camera time to the PC time after I make sure my PC time is correct first! If I have an IRIG time source connected to the camera, then I would set the PC time to the camera time.

3. Power up the camera with lens attached for at least an hour, ideally two hours before you NUC it. That stabilizes the camera and reduces thermal gradients across the lens body that can give you low spatial frequency shading in the image. The camera will stabilize the fastest if it is bolted to a heatsink, like a thick aluminum plate. Figure 27 is a time plot of two internal camera temperature sensor readings, and the image mean, all taken while the camera was viewing a stable 25 °C blackbody. The camera had been off all night, and then it was turned on and allowed to run. It took over two hours to fully stabilize. The camera was just sitting on a table without any external heatsink connected to the baseplate.

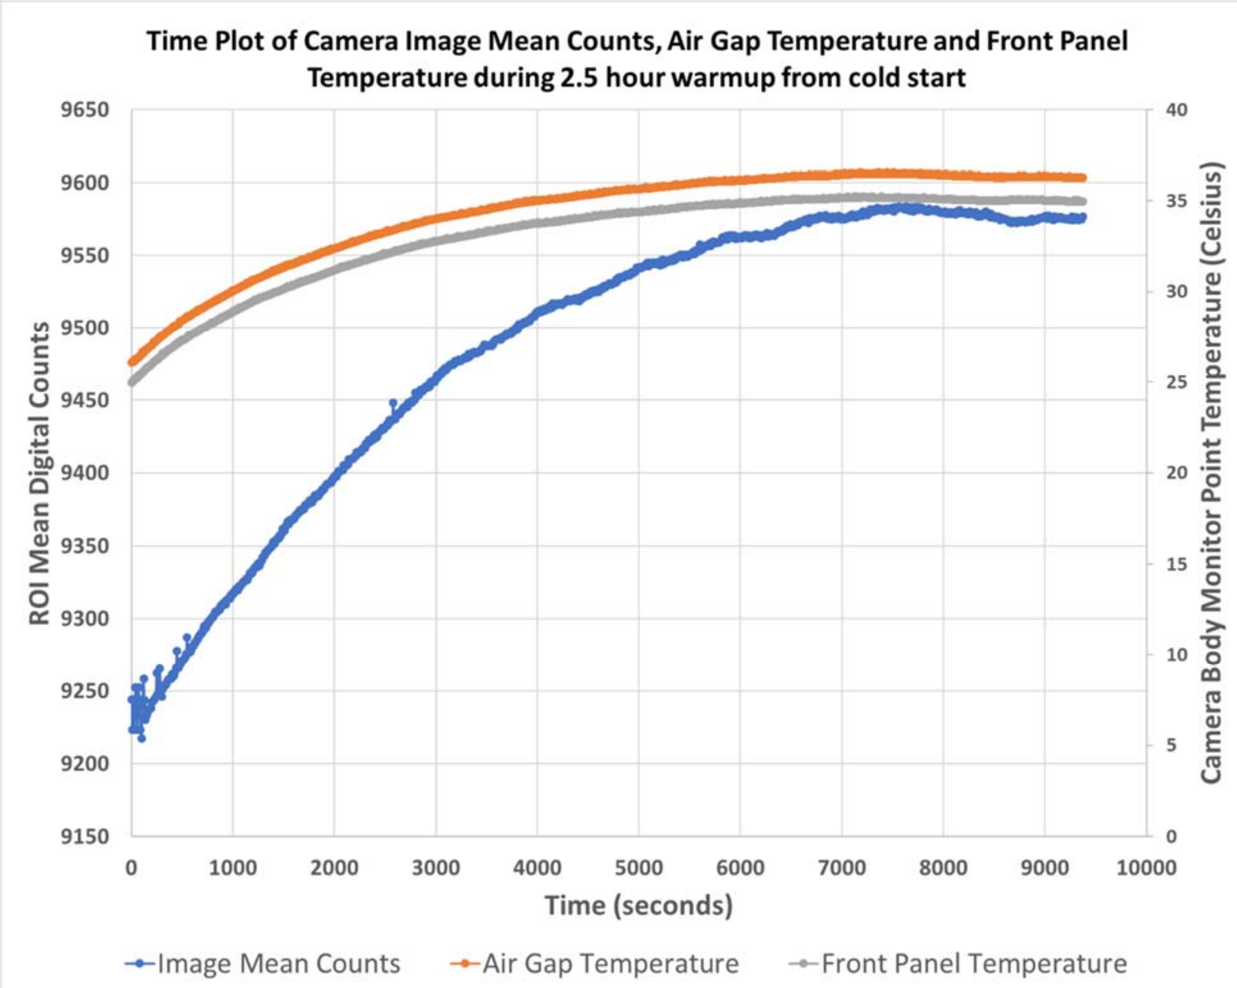


**Figure 27. It can take over two hours for a 10-pound FLIR science camera to stabilize thermally from a cold start if the camera is thermally floating. For maximum precision, let the camera stabilize before performing NUCs or radiometric calibration. Bolt the camera to a big heat sink and it will stabilize faster than this! This camera was just sitting on the table, so it heated up quite a lot.**

4. You want the camera to be as close to the operational state that it will be used in when you NUC it. If things have not fully stabilized when you NUC it, then it will never look as good as it could have. This is particularly true for long integration time, high sensitivity operating states.

5. Look at something as bright as the hottest thing you want to measure. Maybe a cavity blackbody if you have one that will go to the required temperature.

6. Determine the integration time you want to use, as shown in Figure 28, and configure the camera with optics (and any optional warm filters) you want to use.

7. Select the frame rate you want to use and do the NUC at that frame rate. The NUC can be slightly affected by frame rate!

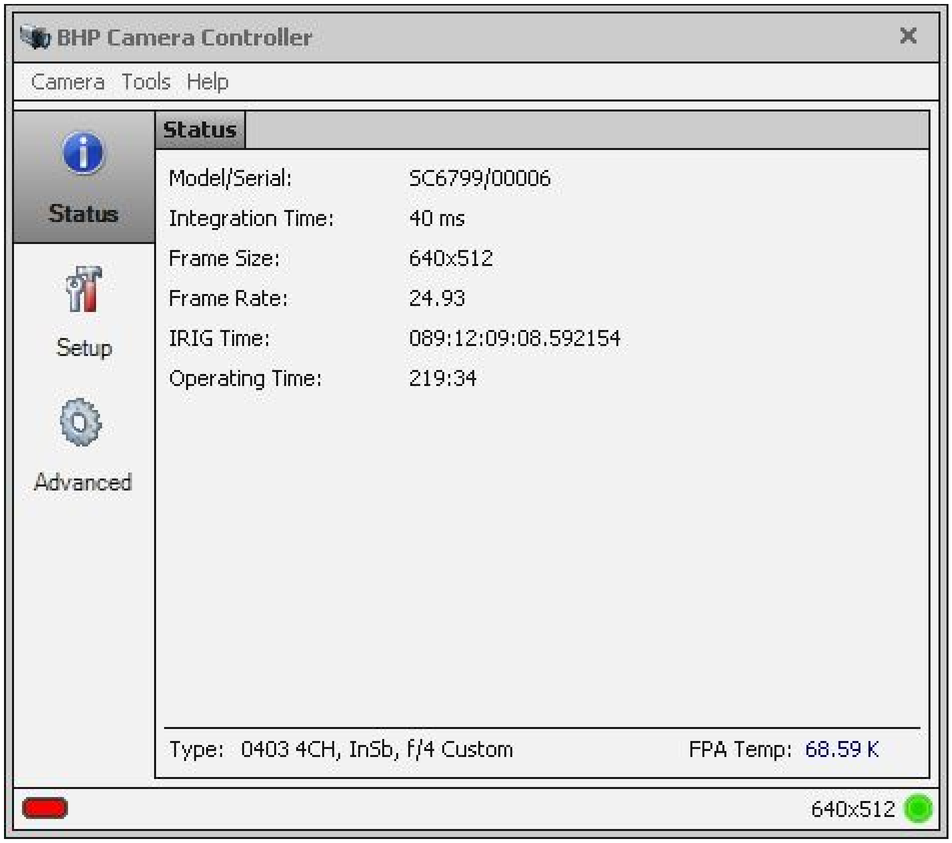


**Figure 28. The camera is set to 40ms for the integration time. The frame rate is 24.93 frame per second which is the maximum frame rate for that integration time.**

8. Two-point NUCs are always what you want to do. The one-point NUC option should not be used. Offset update is only done after you have a two-point NUC and it is temporary. If you power-cycle the camera, the offset update is deleted off the camera. The two NUC points should be done with sources that give digital counts at about 20% and 80% of the camera's dynamic range. This is in counts units, not temperature. For a 14-bit camera, these values are around 3500 counts and 13000 counts. For short integration times, try to get the lower temperature source to give you at least 500 counts of signal above the ambient background. That might be an absolute level of around 1000-1500 counts. The higher temperature source should be giving you mean counts of around 13,000 counts. MAKE SURE there are no saturated pixels in the image!

9. You will need an area blackbody that can overfill the camera's FOV, as shown in Figure 29. Using two blackbodies is very convenient, compared to waiting for a single blackbody to stabilize at a second temperature, which can take a while if the area blackbody is a high-temperature unit. The high temp area blackbodies don't have active cooling, and I have blown fans on them to cool them down. If you only have one high temperature unit for NUCing, do the lower temperature part of the NUC first, then set the

controller for the higher temperature, and wait for the blackbody to settle at the higher temperature.


**Figure 29. NUC station setup with camera lens about 5 inches from the 6-inch square blackbody emitter. The les is focused to infinity, which is a standard approach, although for the ultimate in uniformity, it is good to pre-focus the lens for the desired working distance, then do the NUC with the lens set to that focus, maybe even using Kapton tape to tape the lens ring in a locked position so it won't get changed accidently.**

10. For maximum uniformity, it is a good idea to preset the focus to the focus value that it will be used at during the application. on some of the older lenses, we noticed that there can be radial gradients that can appear if the camera focus is changed from what it was during the NUC process. Newer lenses are less prone to this. If you do change focus and rings appear, you can always do an offset update with a source in front of the lens to clean up that radial fixed pattern noise.

11. Put the NUC source close to the lens and **do not focus** on the surface or you will "NUC in" tiny imperfections in the surface.

12. For ambient scene imaging the NUC plate that comes with the camera is a handy NUC source. It has a finish that is designed to have a high emissivity.

13. Large optics are a challenge to NUC unless you have a really big area blackbody. The biggest blackbody I have used for NUCing is 12 inches square. If your optics are optically well-matched to the camera, then it *may* work to NUC with a smaller lens, then replace the small lens with the larger lens, and then do an external offset update with the larger lens.

### The NUC Process in ResearchIR

The camera control GUI within ResearchIR has two modes, Basic and Advanced. You need to be in the Advanced mode to create NUCs using the NUC wizard. ResearchIR will stay in Advanced mode indefinitely after it is set once. The install default mode is basic, so the user will have to go to the camera control GUI and pull down to find the menu option, as shown below in Figure 30:

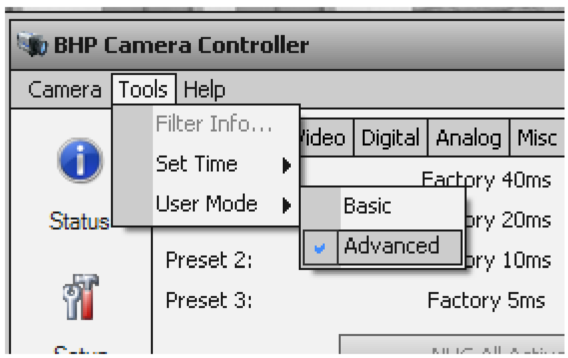


**Figure 30. Set the camera controller GUI to Advanced**

The camera is set to 40ms integration time. Make sure that you are in Single Preset mode and that the radio button next to 40ms integration time is selected, as shown in Figure 31:

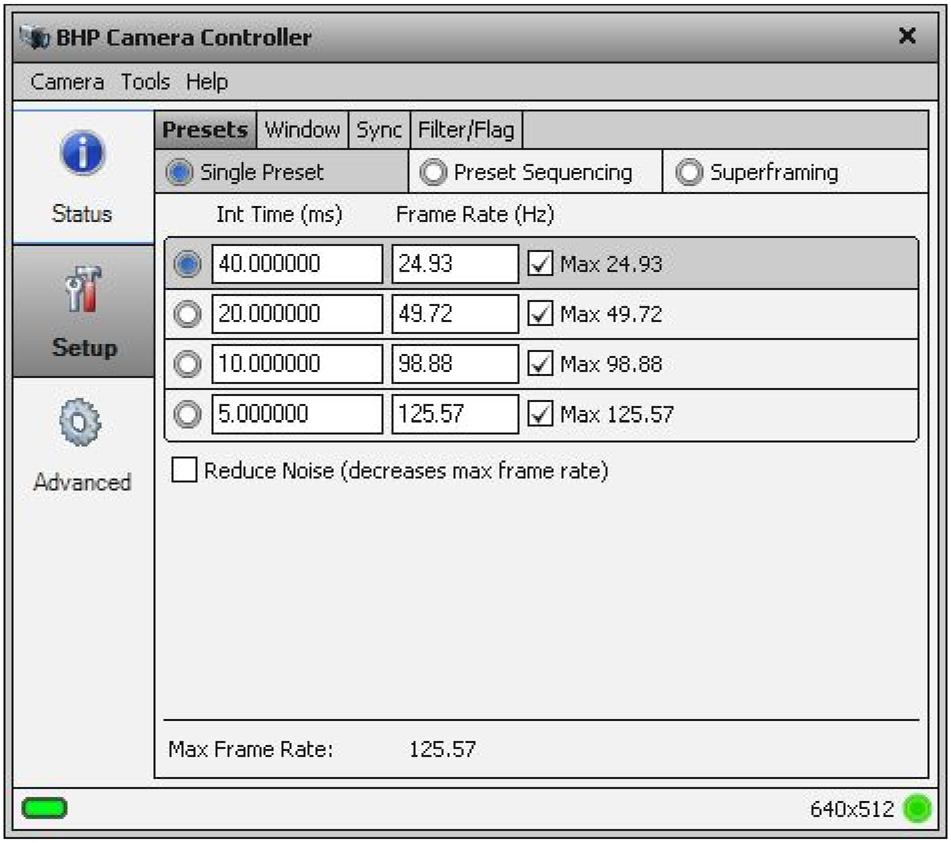


**Figure 31. The NUC will be done on Preset 0, which is the active preset. You can tell it is the active preset because the radio button next to the number 40 is blue.**

Turn off the effect of any NUC that might be in the camera. You want to be performing the NUC process with an uncorrected image to make sure you aren't clipping a lot of pixels on the 16383-count digitizer rail. You can see a live image during the NUC process. For this example, the camera I used was set up with a factory NUC at 40ms, and the NUC was being applied to the image, as shown in Figure 32:

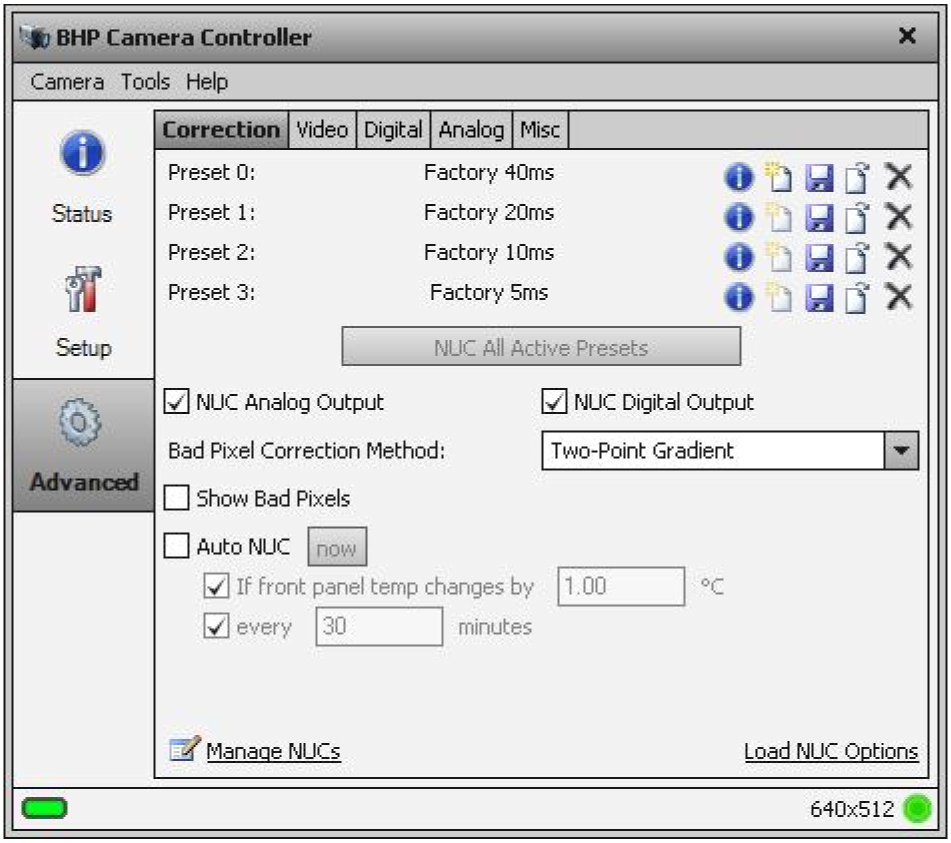


**Figure 32. There is a factory NUC for 40ms integration time loaded and the digital image data is being NUCed, because the NUC Digital Output checkbox is checked.**

I unchecked the NUC Digital Output checkbox so that the digital data stream is not corrected, as shown in Figure 33:

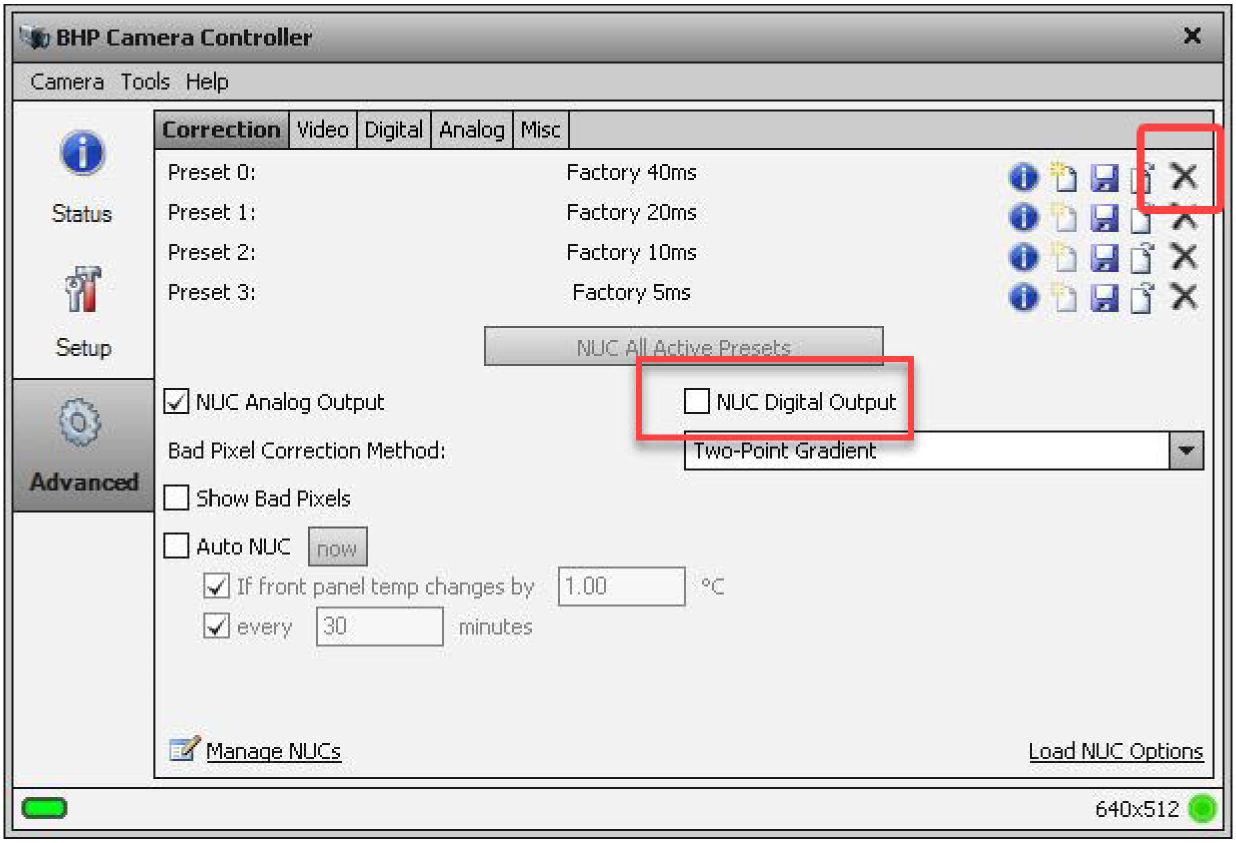


**Figure 33. The factory NUC is not being applied since the checkbox is now unchecked, but the factory NUC is still loaded in the camera. The "X" button will unload this NUC from the preset.**

Just to avoid any confusion, and since we are creating a new NUC anyways, it is a good idea to use the X button to unload the NUC that was loaded into Preset 0, as shown in Figure 34. Now there is no NUC name shown next to Preset 0:

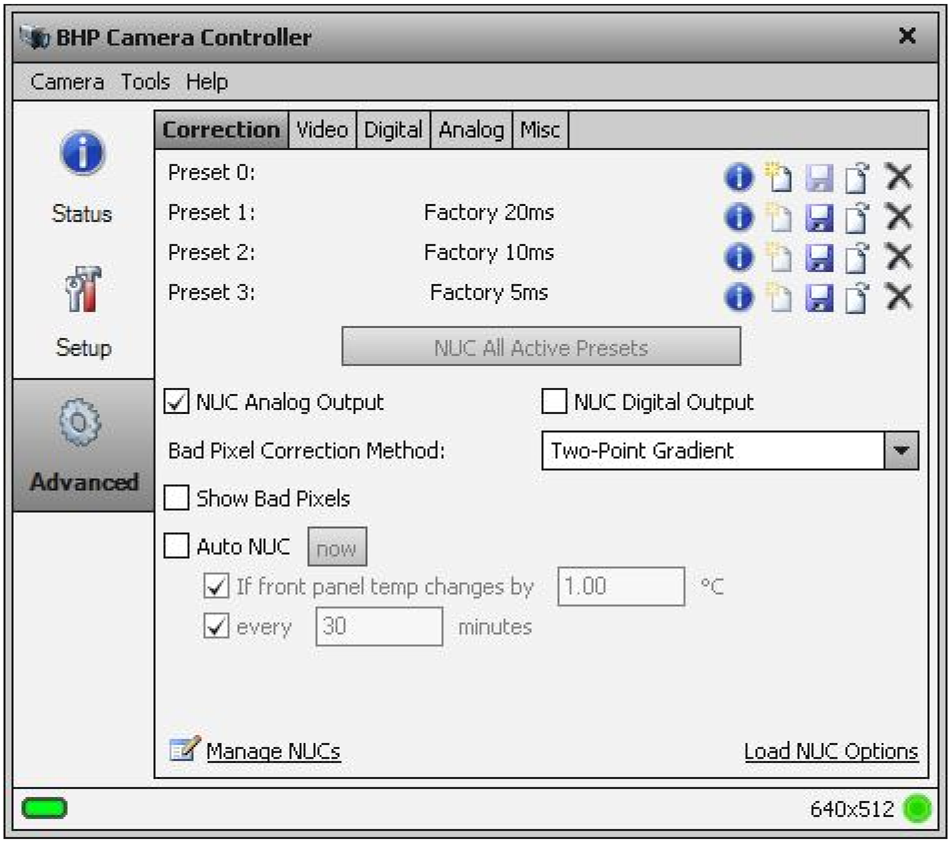


**Figure 34. The NUC in the camera has been unloaded, so the NUC name field in Preset 0 is blank. If we now do a camera-side NUC, the digital data won't be NUCed until we check the "NUC Digital Output" checkbox**

Press the little button that looks like a sparkling new sheet of paper to perform a camera-side NUC (Figure 35):

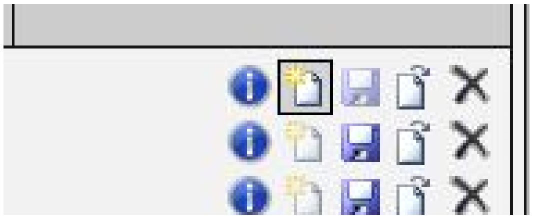

**Figure 35. New NUC button with focus on it.**

and you get this NUC creation screen (Figure 36):

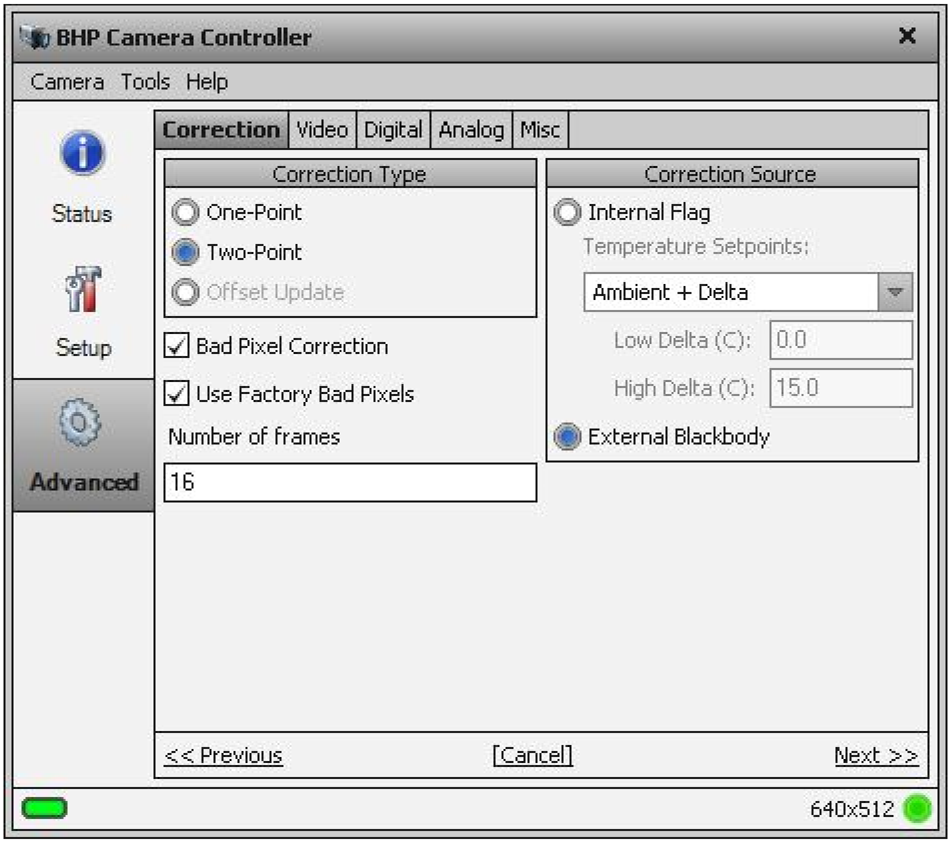


**Figure 36. NUC creation window**

We need to do a two-point NUC with the External Blackbody option. At this point, no NUC was loaded, which is why the Offset Update option is greyed out. That is not an option unless there is a two-point NUC loaded, or a one-point NUC.[8] The default values for Bad Pixel Correction and the use of Factory Bad Pixels are that the checkboxes are checked. The factory bad pixel map is embedded into the camera software. If it is used, then it is combined with any new bad pixels that are found in the NUC process. The default number of frames is 16, and this seems to work very well.

Hit the Next button and you will get the Bad Pixel Detection Parameters screen shown in Figure 37:

[8] We don't recommend ever doing a one-point NUC. It sets the gains equal to 1 and does an offset correction on a single source, and then it saves it to the camera. If the image mean count value changes much at all, the image looks like it has no NUC applied. It is a last-ditch option when it is not possible to create a two-point NUC because you only have one NUC source available.

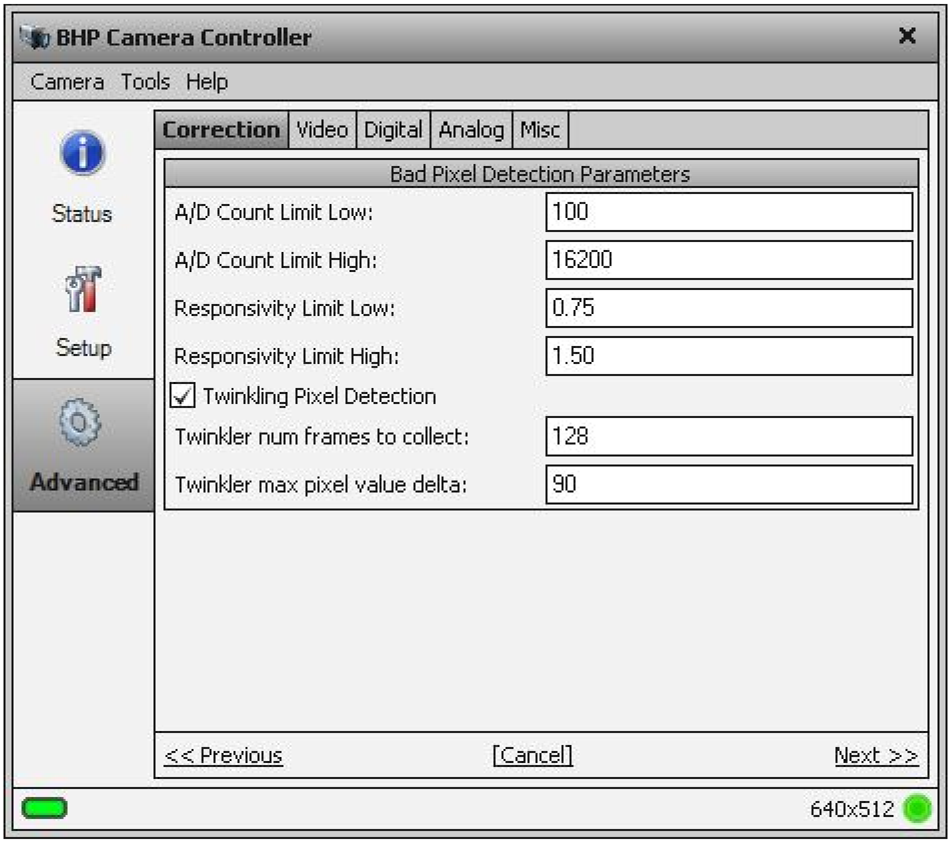


**Figure 37. Bad pixel detection parameters**

The above bad pixel detection parameters are the default values. They work well in most cases. Sometimes, a particular lens may have more roll-off on the corners than typical. In that case, the Responsivity Limit Low may need to be set lower than 0.75 or you will get a lot of bad pixels in the corners of the image. The "Twinkler max pixel value delta" is a parameter that defines what a twinkle is. It means that over a sequence of images (128 in this case, the default value), if a pixel deviates more than 90 counts from its mean value across the 128-frame sequence, then it is marked as bad. So bad pixels can be added to the bad pixel map because they are lower than 100 counts (probably an open circuit), or greater than 16200 counts (shorted out), or have too little or too much response relative to the mean response, or they are twinkling. After you have set these parameters, hit Next.

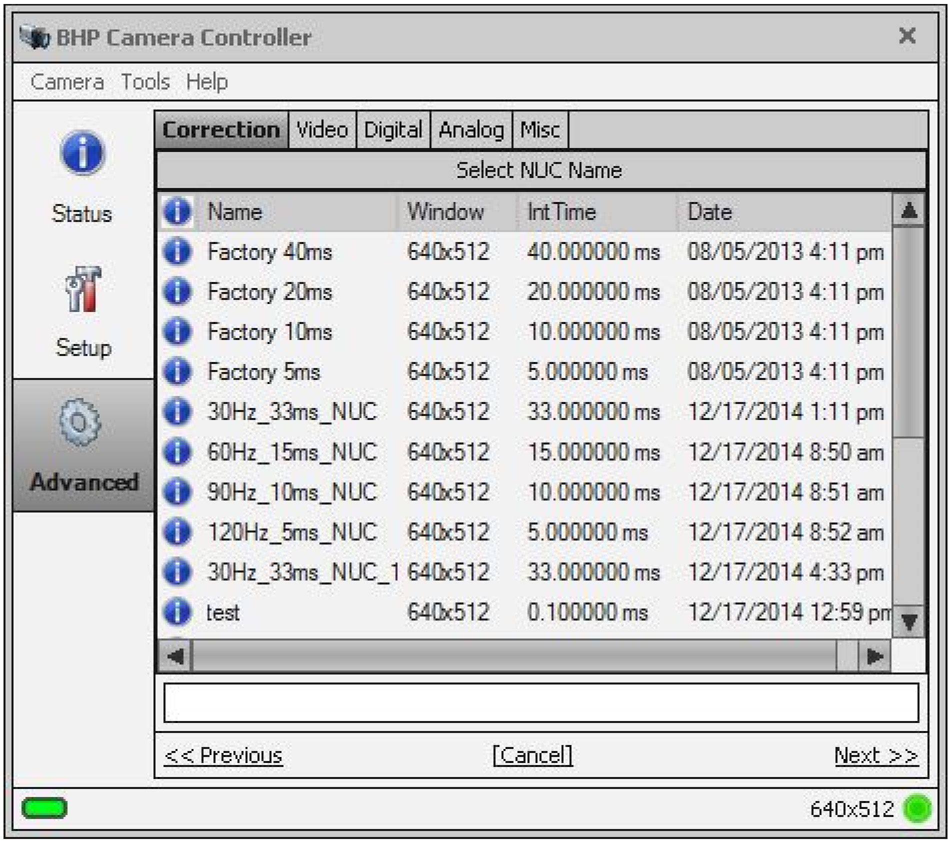


**Figure 38. NUC name input field**

Give the NUC a name in the next screen, shown in Figure 38. Note the Date column. This indicates the date that the NUC was created. MAKE SURE you set the camera time to the PC time, or IRIG time if you have that as an input, BEFORE you create a NUC. If you later look at the date and it is totally different from the day that you know you created it, then there will now be a point of confusion that should have been avoided by setting the clock correctly!

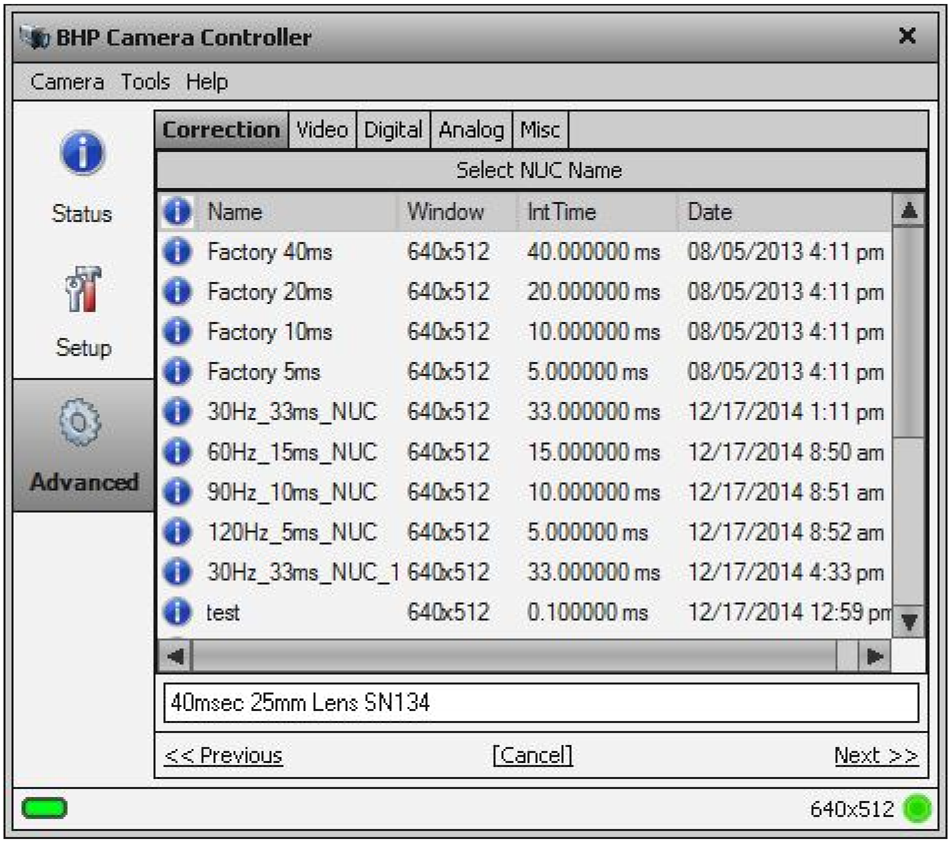


**Figure 39. NUC naming text field. The visible part of the text field is 63 characters long.**

I give the NUC a name that is descriptive and that captures what lens was used, down to the serial number in case I have an assortment of 25mm lenses, as shown in Figure 39. The name can be pretty long, since the field can hold 63 characters in the visible portion of the text field, but you can actually go longer. If the name is too long, it will crash the software and Windows won't like it because file path names can get too long for the operating system file directories. I tried all this out, so you don't have to!

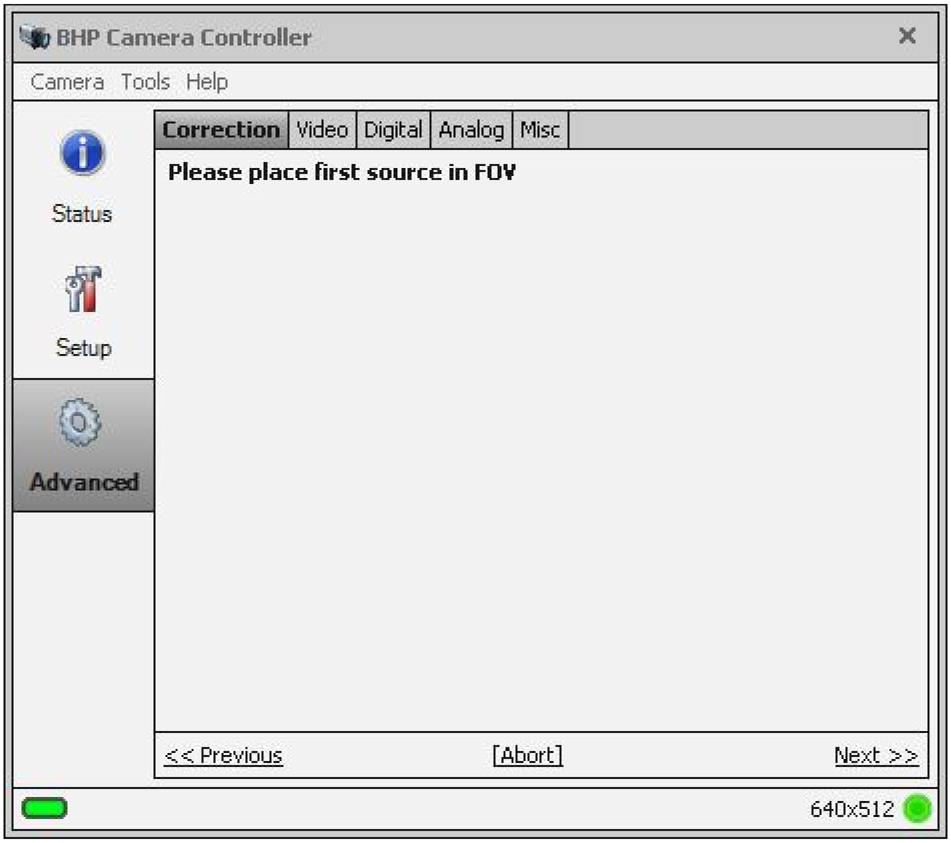


**Figure 40. NUC routine asking for the second source. The second source should produce a mean count value at least 1000 counts different from the first source. The second source does not have to be hotter than the first, it can be colder than the first. The software figures it all out.**

Now it will ask me to place the first source in the field of view, as shown in Figure 40. Make sure the source is very uniform and that the surface is out of focus. It is best practice to set the focus to infinity or whatever the focus position needs to be for the use case, then do the NUC, and that way the NUC is optimized for the use-case focus setting. If you end up changing the focus, you can do an offset update to get rid of any non-uniformity.

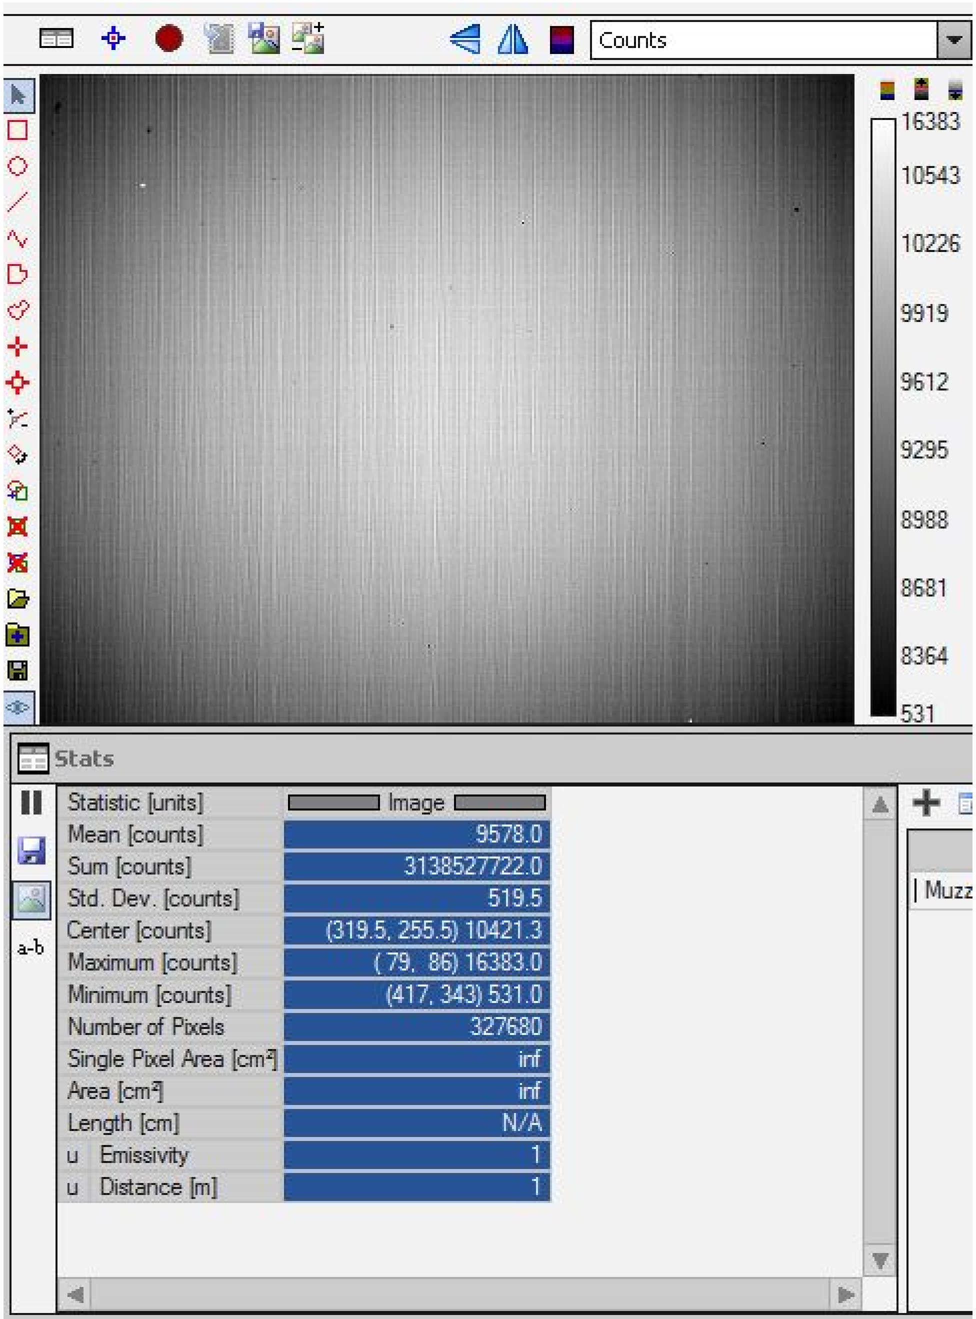


**Figure 41. Image of the 25C blackbody during the NUC process**

You can and should always be looking at the LIVE statistics for the images during the NUC process. Figure 41 is an image of the 25 °C blackbody, which is what I used for the first source. The mean counts are 9578 and the standard deviation is 519 counts. This is a high standard deviation compared to what it will be after it is NUCed. The standard deviation divided by the mean is a measure of the uniformity. We call it the uncorrected uniformity, since the image data is not NUCed. It is 0.054, or 5.4%. It probably should be called the *uncorrected* non-uniformity because the lower the number, the better the uniformity.  But that is not how the IR camera industry defined the term.

For the second source, I used a 35 °C setting on the blackbody and waited 5 minutes for it to fully stabilize.  I chose 35 °C because it gives me about 3000 counts of signal swing from the first source.  It is good practice to get at least 1000 counts of signal swing between the first and second NUC sources. I used a Temporal Plot of the image mean counts to monitor the stabilization at the new setpoint. I plotted the image mean with a measurement function, and watched the counts as the blackbody went up to 35 °C, overshot, came back, and settled, as shown in Figure 42.  The measurement function feature is described in the ResearchIR user guide.  This is a better way than trusting the blackbody controller to tell you when it thinks the blackbody has stabilized. The controller is using a temperature sensor buried in the metal block that the emitter is made of.  There is **always** a delay between the temperature it measures and the emitter surface temperature because of the time constant associated with thermal diffusion.

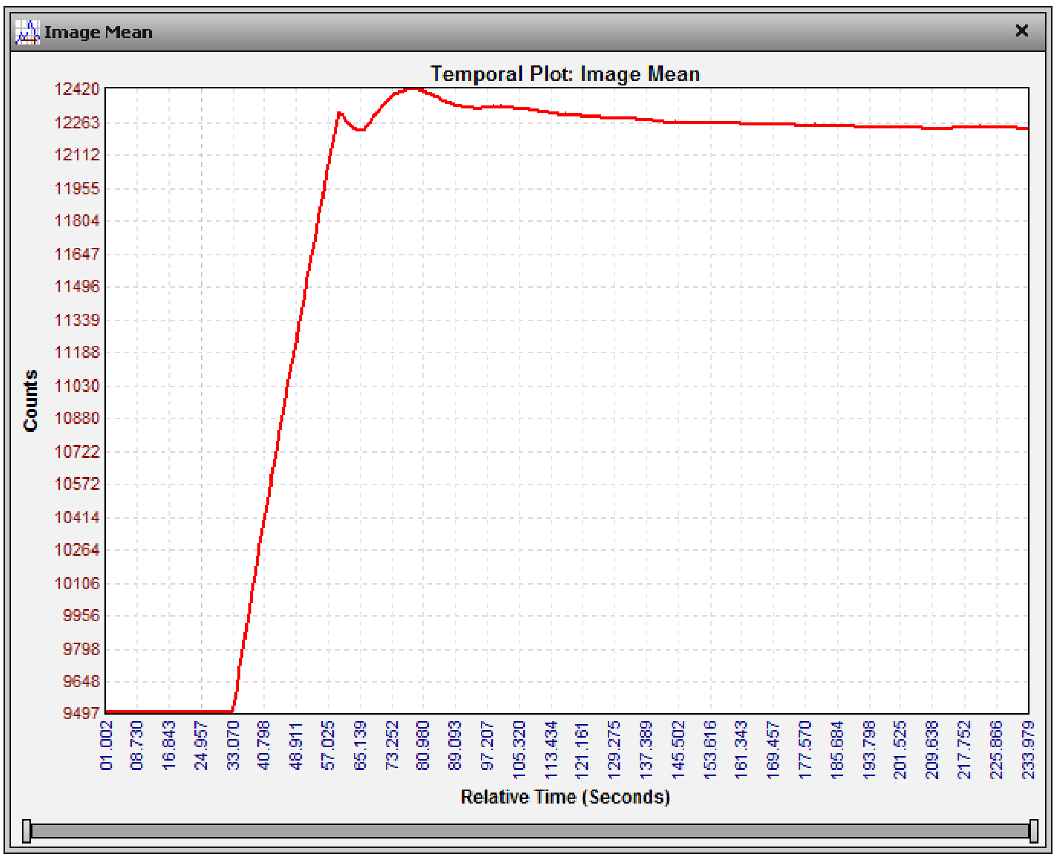


**Figure 42.  Blackbody ramping from 25 °C to 35 °C as seen by the SC6700 See-Spot camera.  To be on the safe side, it looks like I should wait about 3 minutes after changing the controller setpoint before acquiring the frames for the NUC second source. Then the blackbody will have stabilized in its radiance. It will also tend to be more uniform spatially.**

I also monitored the image statistics to make sure I got enough signal swing but did not saturate the image in the center – see Figure 43. You can use the Segmentation tool in ResearchIR and set the upper threshold to 16000, for example, and that will warn you if you start to saturate pixels in the center – they will turn red.

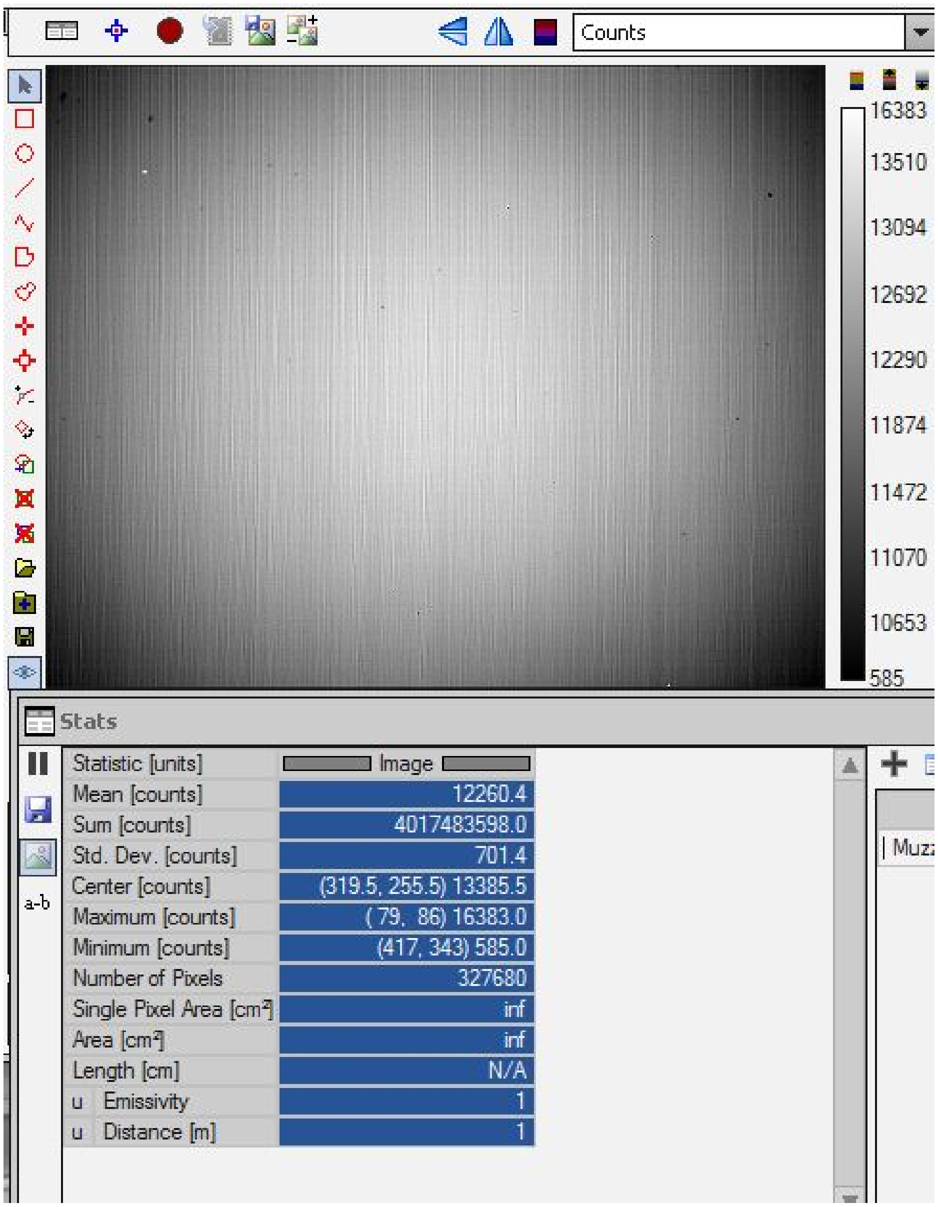


**Figure 43. Image of the 35 °C blackbody used as the second NUC source. The mean value is an acceptable IR radiation level – the sensor is within its linear range. There are a few saturated pixels in a cluster, but they are bad pixels and they will get replaced.**

When the counts weren't changing much, and the changes were random in direction, I then took the 35 °C NUC data by hitting Next, as shown in Figure 44:

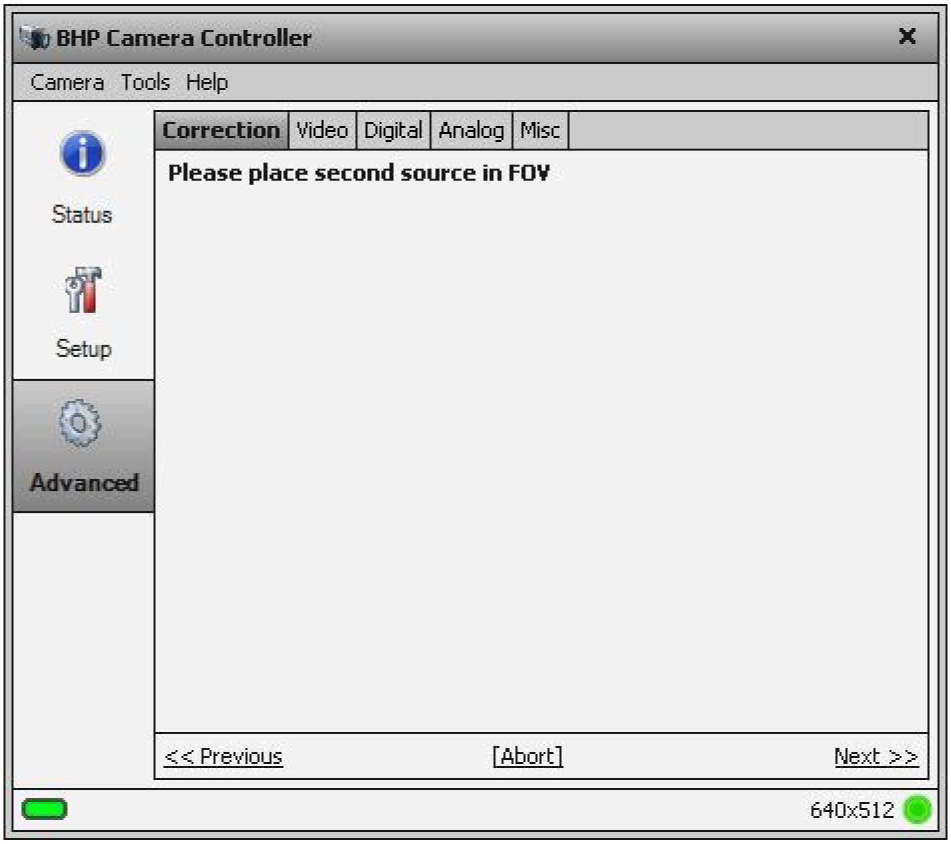


**Figure 44. Put the second source into the field of view and hit Next.**

The camera now collects the twinkling pixel data, which takes a little while, depending on the frame rate of the camera preset. See Figure 45:

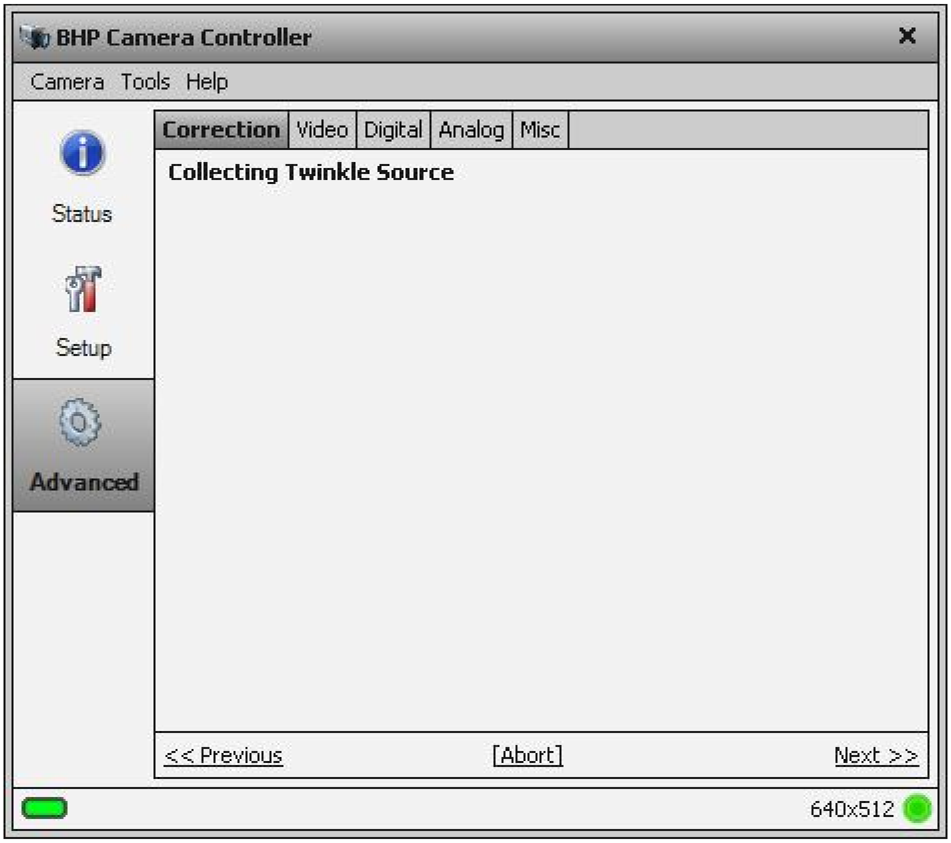


**Figure 45. The camera will take 128 frames and look for twinkling pixels. These are pixels that have values that depart from the mean over a 128-frame sequence by more than 90 counts.**

When the Twinkling pixel detection has stopped, the NUC wizard shows the statistics for the NUC (Figure 46). You can see how many bad pixels there are, and how many were marked as bad for the various criteria. I had chosen to use the factory bad pixels, so the factory bad pixels are included in the statistics. The NUC process only captured one additional new pixel into the map. The factory bad pixel map contains all the pixels that were "dead on arrival" – either open or shorted. This is an excellent array with only 63 bad pixels. At very long integration times, we sometimes see more bad pixels show themselves.

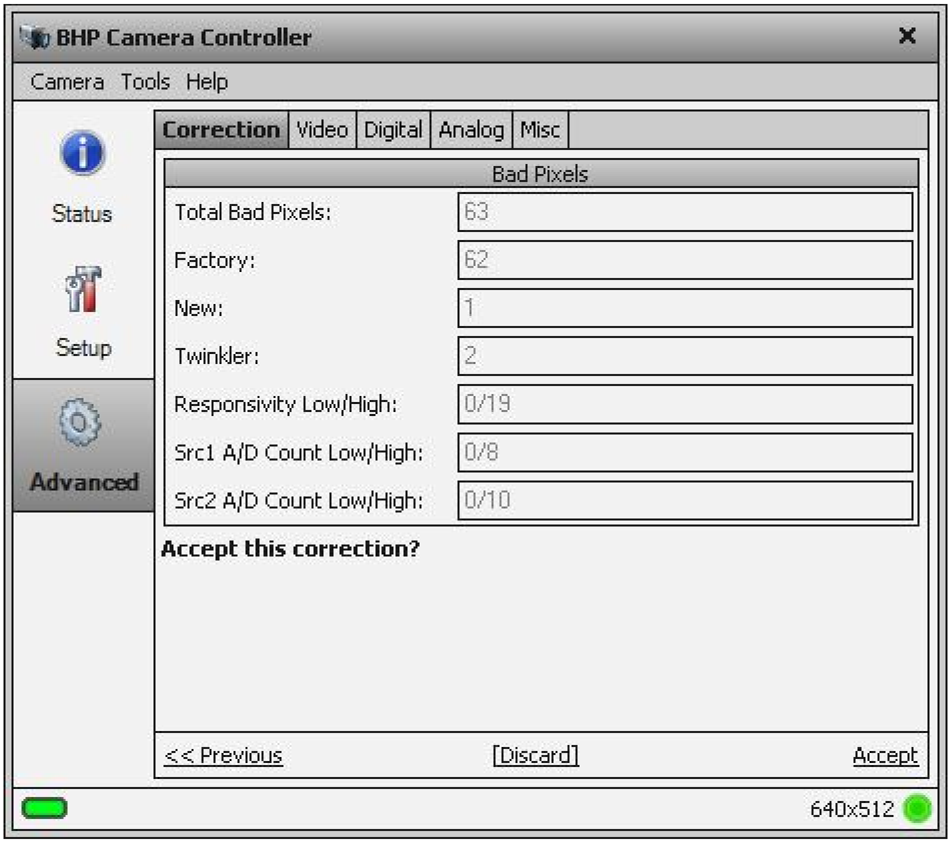


**Figure 46. If the NUC looks good, you can accept it. You should check the NUC Digital Data checkbox, so you can see what the second source looks like with the new NUC applied. This is a good NUC from the perspective of the number of bad pixels. There were already 62 factory bad pixels, and this only added one to the total.**

It is a good idea to check the NUC Digital Output checkbox and see the image to make sure the NUC looks good (Figure 47):

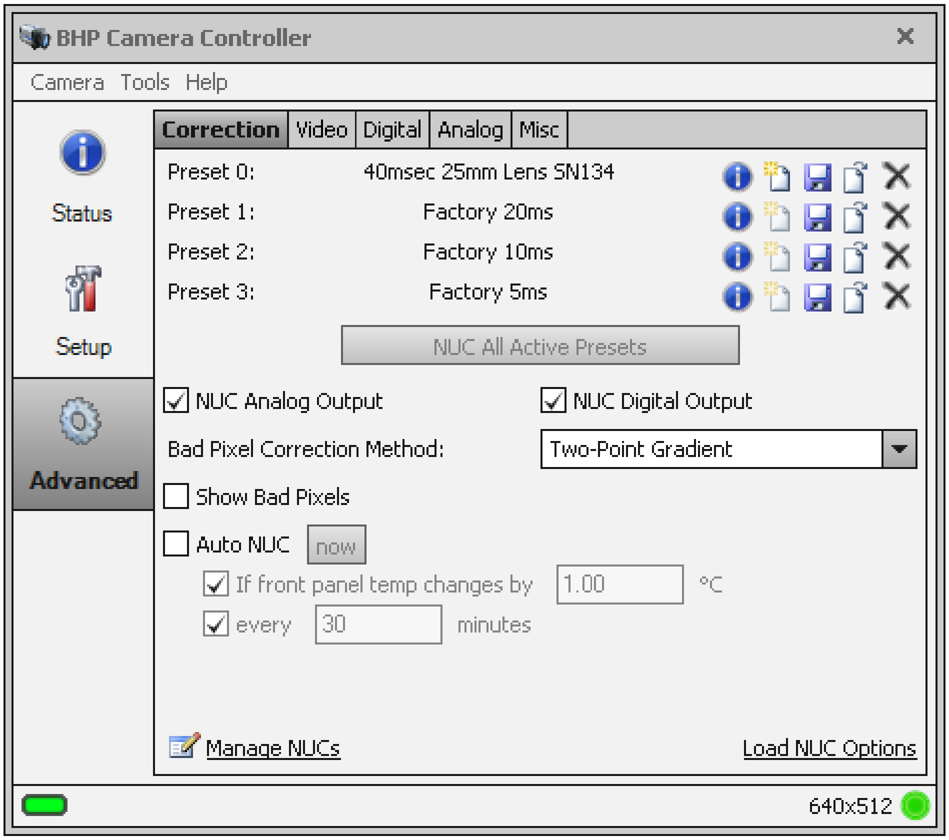


**Figure 47. Check the checkbox for NUC Digital Output to see if the NUC looks decent before you accept it. There is no cost to accepting the NUC. If you don't like it, you can redo it.**

Let's look at what the NUC does to the image now. Figure 48 is an image of the 35 °C blackbody used for the second source. The NUC looks very good. There is a little bit of time-varying noise because of convection-driven air currents that cool off the surface in random swirling patterns.

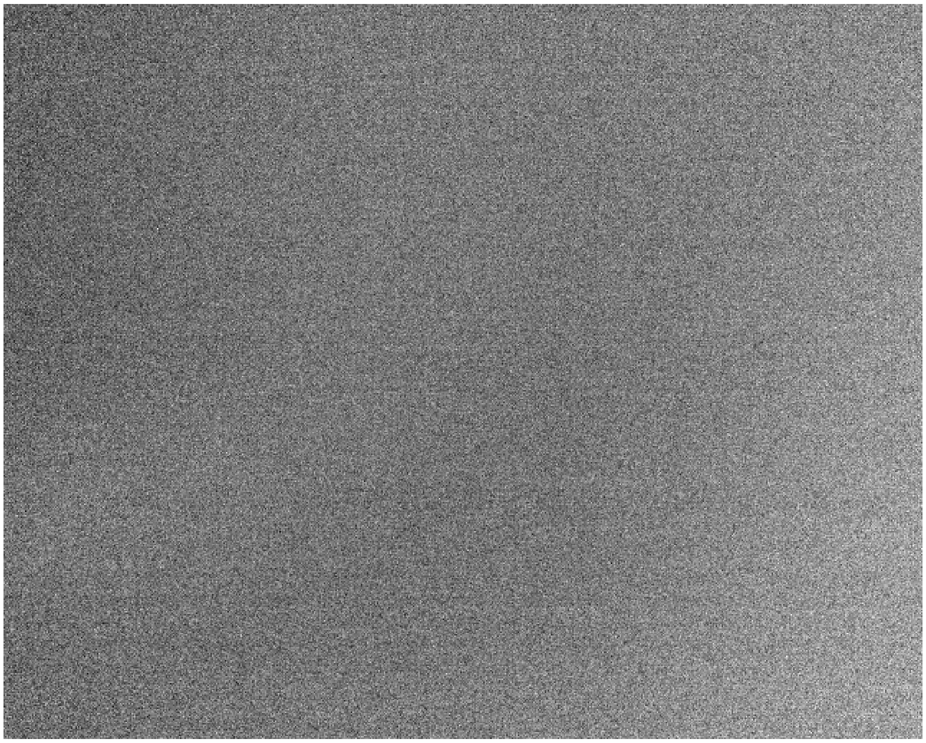

**Figure 48. 35 °C blackbody right after the completion of the two-point NUC. The standard deviation of the image is 6.9 counts. The mean is 12,278 counts. This is a nice NUC. There is some rippling in the image above. This is time-varying low frequency spatial noise that is caused by convection and air currents cooling the face of the blackbody. The NUC looks very good. One can take a long sequence of frames and average them to see the rippling patterns get averaged out.**

Now we will look at a scene with the new NUC. Figure 49 is an image of the bookshelf in my office:

**Figure 49. Bookshelf image with user-created, in-camera NUC applied**

Figure 50 is an image from the camera with the original 40ms factory NUC looking at that same scene. The quality is subjectively similar to the user created NUC.

**Figure 50. Bookshelf image with factory 40ms NUC applied.**

Figure 51 shows the same scene with the NUC turned off. This example really shows that a good NUC is vital to properly correct images of a relatively low-contrast scene like this.

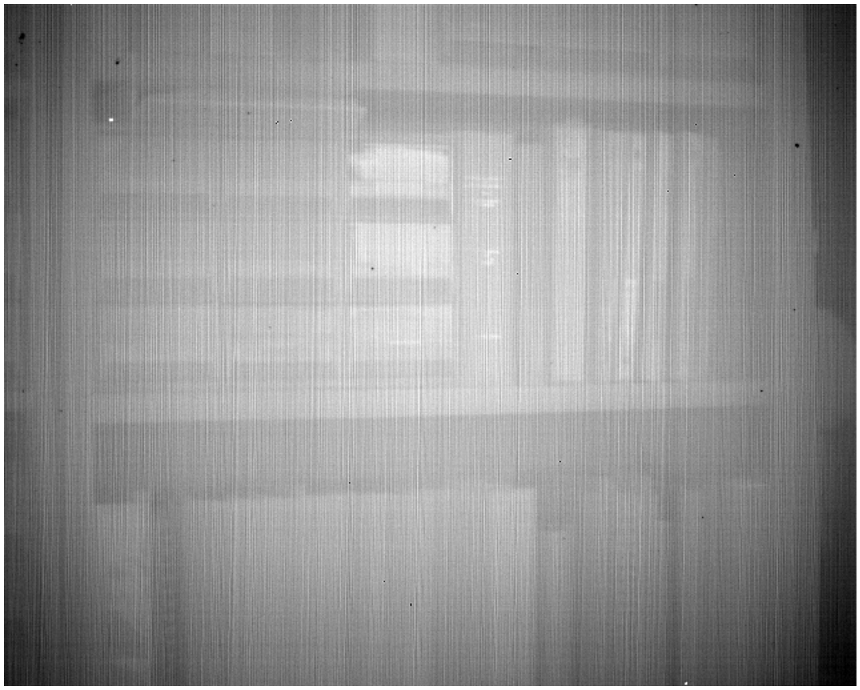

**Figure 51. Bookshelf with the NUC turned off.**

## Saving the Camera State File and the NUC

Now that I have created a new NUC, I want to make sure that I don't lose it! This is not a factory NUC that is backed up at the Niceville plant. It only lives in the camera at the moment. I want to back it up. I also want to save the camera State file so that I can get back to the camera state with the NUC I just created as the active loaded NUC. I showed you how this was done earlier. Now would be a good time to save the State file again. You can download any NUC file off the camera to the host PC to back it up. This works for factory NUCs as well. Go to Manage NUCs, shown below in the red box in Figure 52:

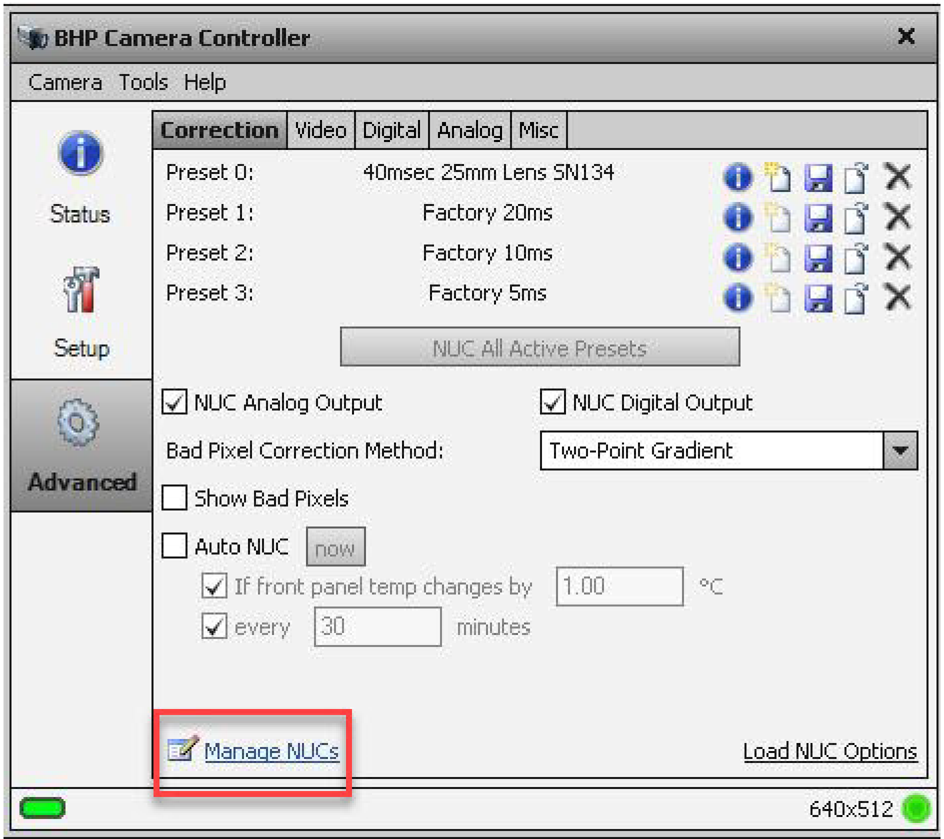


**Figure 52. Managing NUCs link on Correction tab of camera control GUI**

You will get the window shown in Figure 53. You can put focus on the NUC file you want to download and then hit the download button.

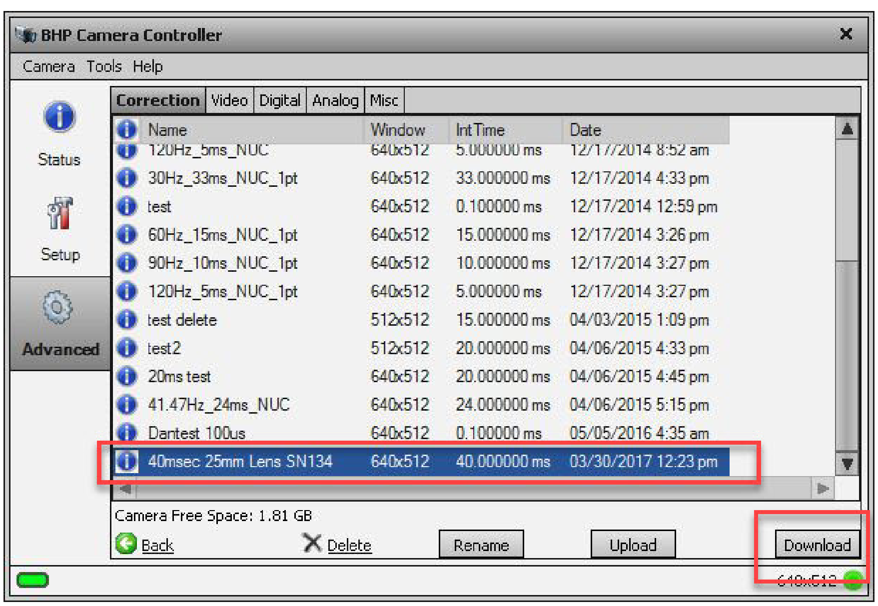


**Figure 53. Downloading a selected NUC with the NUC manager**

Then you can save the NUC to wherever you like, as a *.nuc file, as shown in Figure 54.

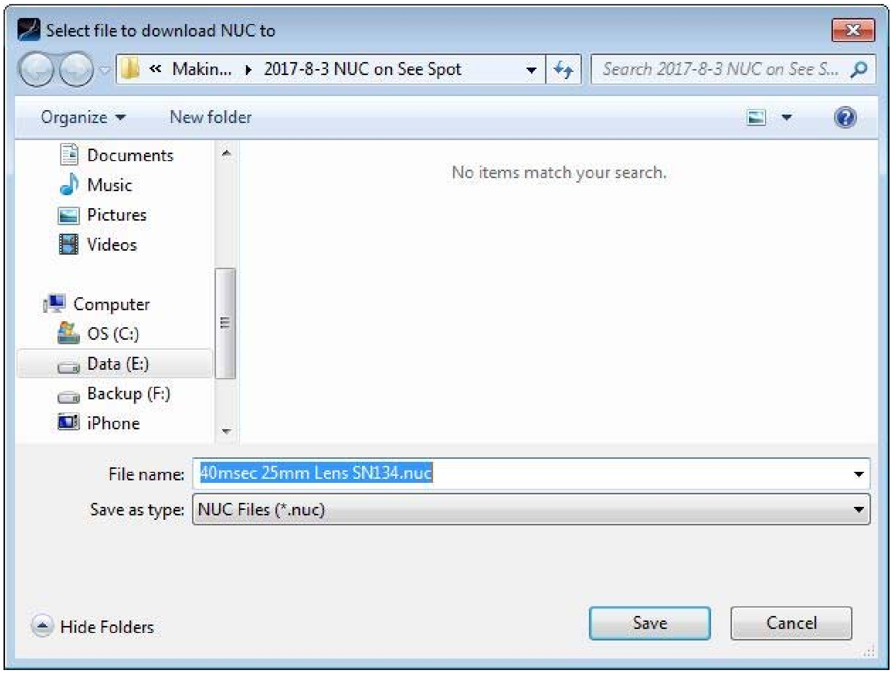


**Figure 54. Saving a camera-side NUC to the host computer**

The *.nuc file is 1281 KB in size. Both the user-created, camera-side NUC and the factory NUCs have the same file size and the same *.nuc file extension (Figure 55).


**Figure 55. NUC file in Windows directory**

## NUC Stability

How long is a NUC good for? Well, it turns out that if you create a two-point NUC on an InSb camera and save it, and then a few days later, you turn on the camera again, load up that NUC and look at a uniform source, you WILL see non-uniform artifacts in the image. There are small changes in the InSb detector offsets that occur every time a cooled camera is power cycled. The gain terms tend to be very stable over time, but not offsets.

Figure 56 is an image of the 35 °C blackbody right after the NUC was created:

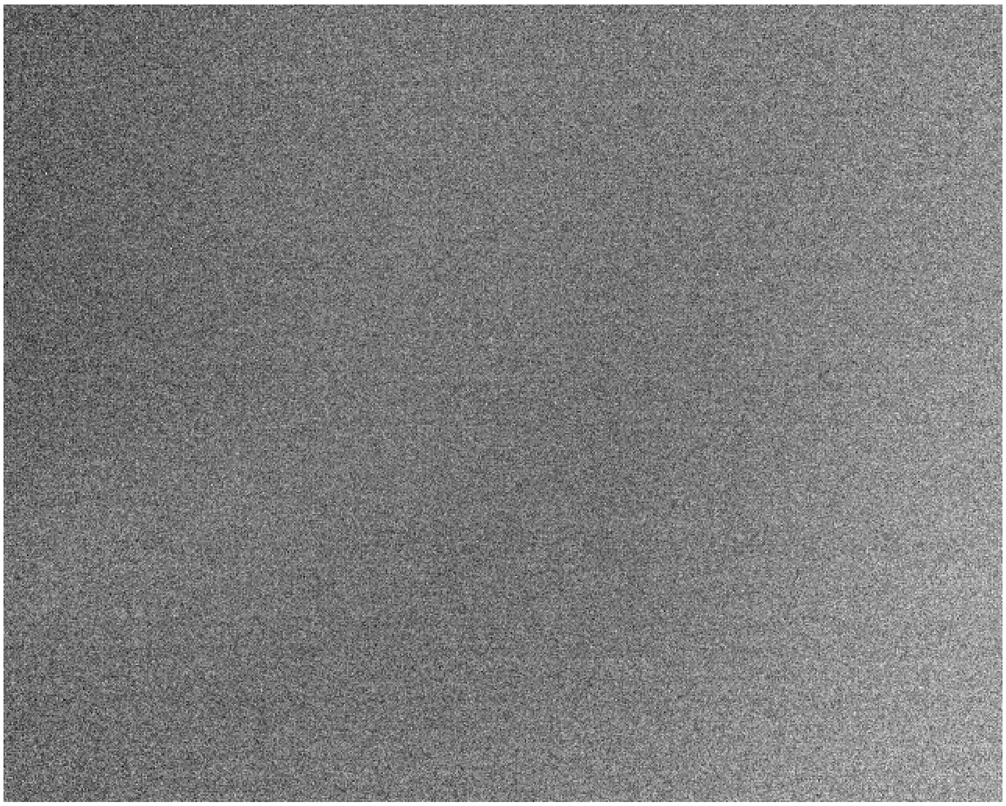

**Figure 56. A single frame of the 35 °C blackbody seconds after the NUC was created. NUCs can't get any fresher than this one. The standard deviation is 6.9 counts and the mean value is 12,278 counts.**

Figure 57 is the same user-created, in-camera NUC applied to the camera while looking at a 35 °C blackbody a few days after the NUC was first created. The camera was turned off for a few days prior to this data collection:

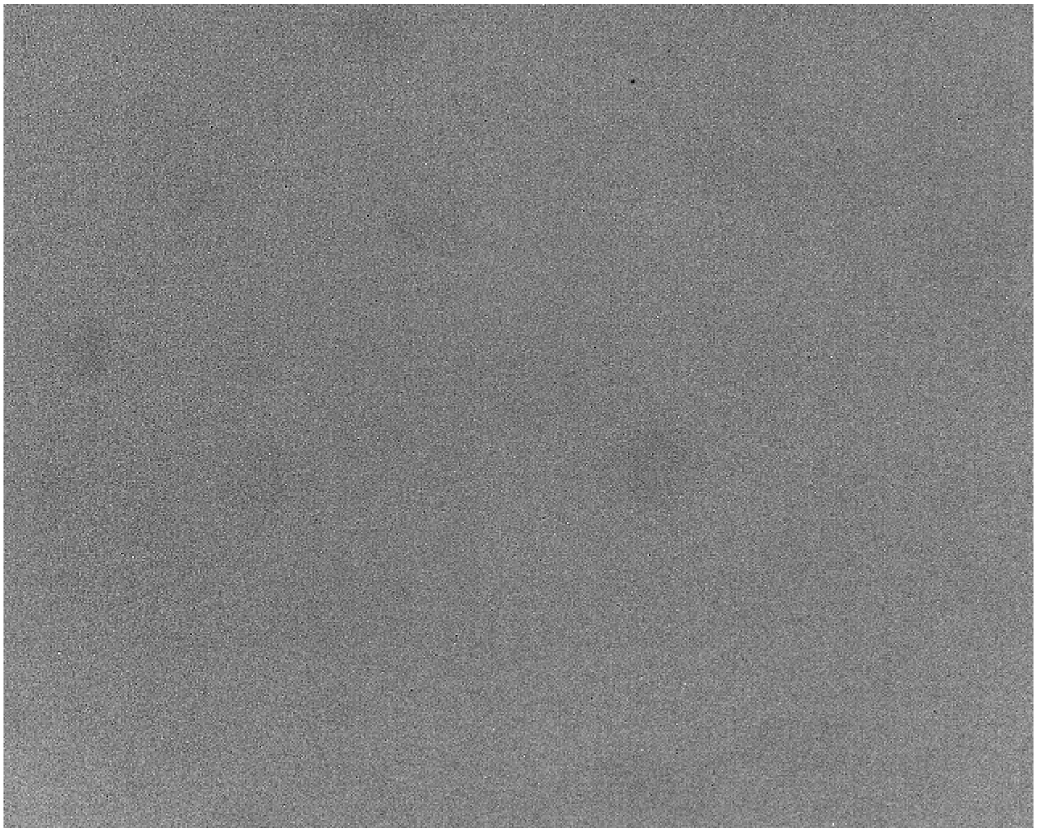

**Figure 57. An image of the same 35 °C blackbody taken a few days after the two-point NUC was created. The camera was turned off for several days, then powered up and allowed to stabilize for several hours. The standard deviation in this image is 11.8 counts. The mean is 12,271 counts. Note that the mean count value is almost identical to what it was when the same blackbody was imaged several days earlier. Then it was 12,278 counts. This characteristic is the key to radiometry – a stable and repeatable camera output for a given input radiance!**

Now I put a 25 °C ambient temperature calibration plate in front of the lens, which I will use to do an external offset update. There is some low spatial frequency spatial noise – the blobs you see here in Figure 58:

**Figure 58.  25 °C ambient calibration plate imaged using a NUC that is several days old.  The standard deviation is 13.5 counts, and the mean is 9534 counts.**

Now I perform an external offset update (Figure 59) which really cleans up the image. Figure 60 was made while the camera was still looking at the 25 °C offset update source:

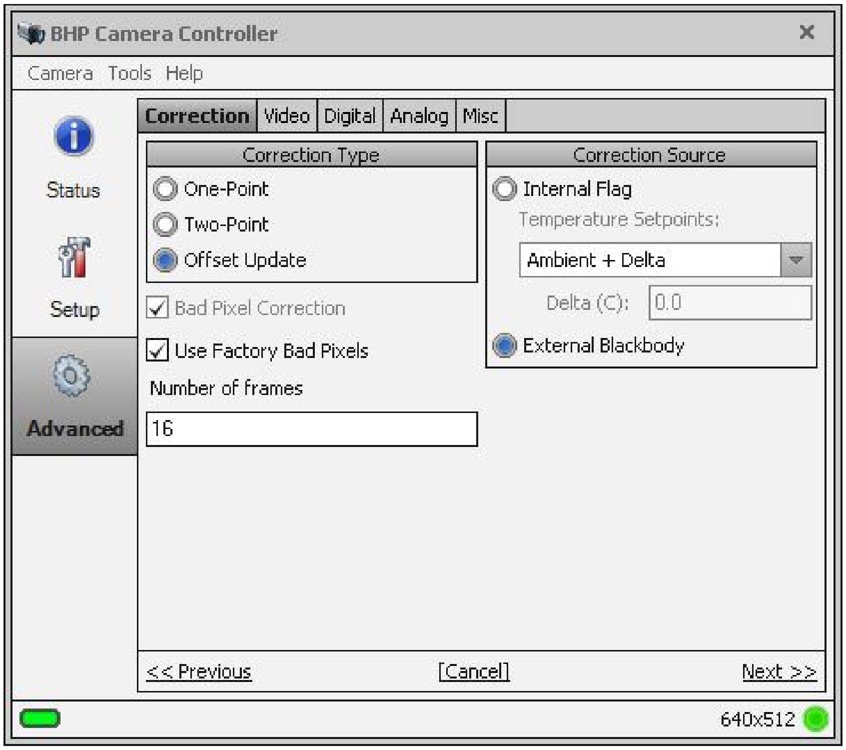


**Figure 59. Commanding an offset update with an external NUC source**

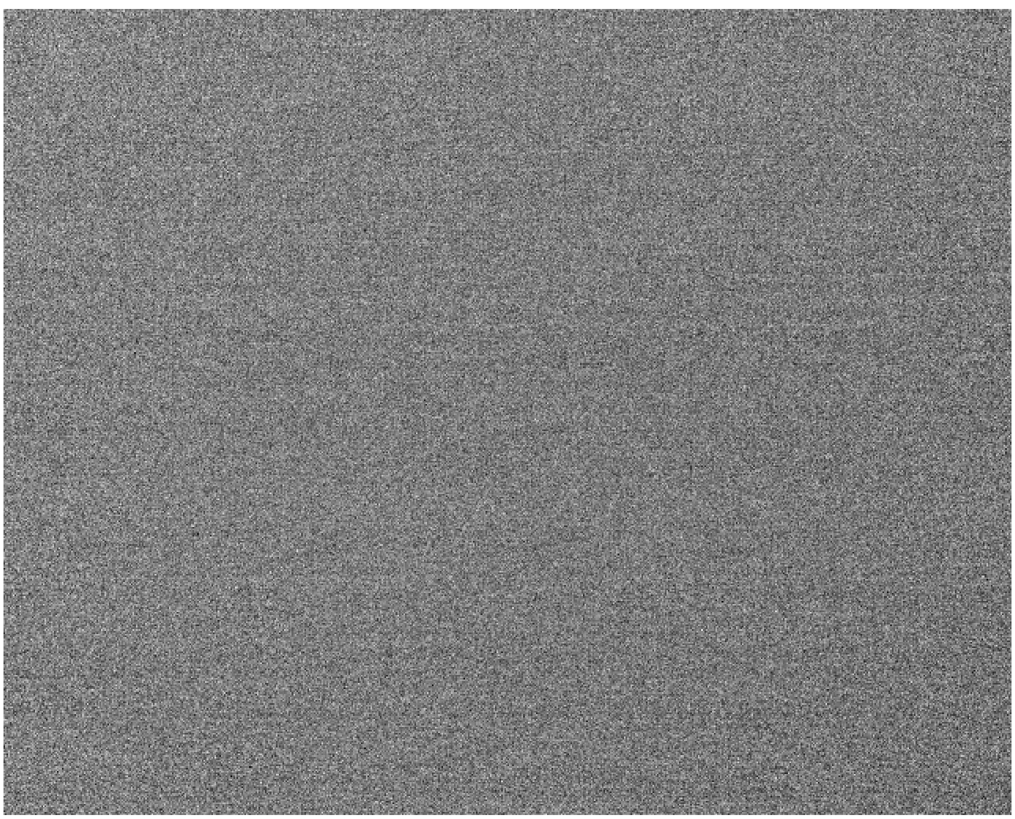

**Figure 60. The 25 °C blackbody right after an external offset update using the blackbody. Standard deviation is 5.4 counts and the mean is 9690 counts. The image is as good as it can get. Notice that there are no convection patterns in the image, which makes sense, since the air temperature is also 25 °C, so there is no temperature differential to drive convection. If I change the temperature of the uniform scene, then non-uniformity will creep back in.**

Now I look at the 35 °C blackbody, and there is some non-uniformity, as shown in Figure 61. My offset update was done with a 25 °C source. When I move off of the 25 °C radiance level used for the offset update, there is some non-uniformity creeping back in. If I wanted to maximize the uniformity for a 35 °C scene, I would do another offset update using the 35 °C blackbody.

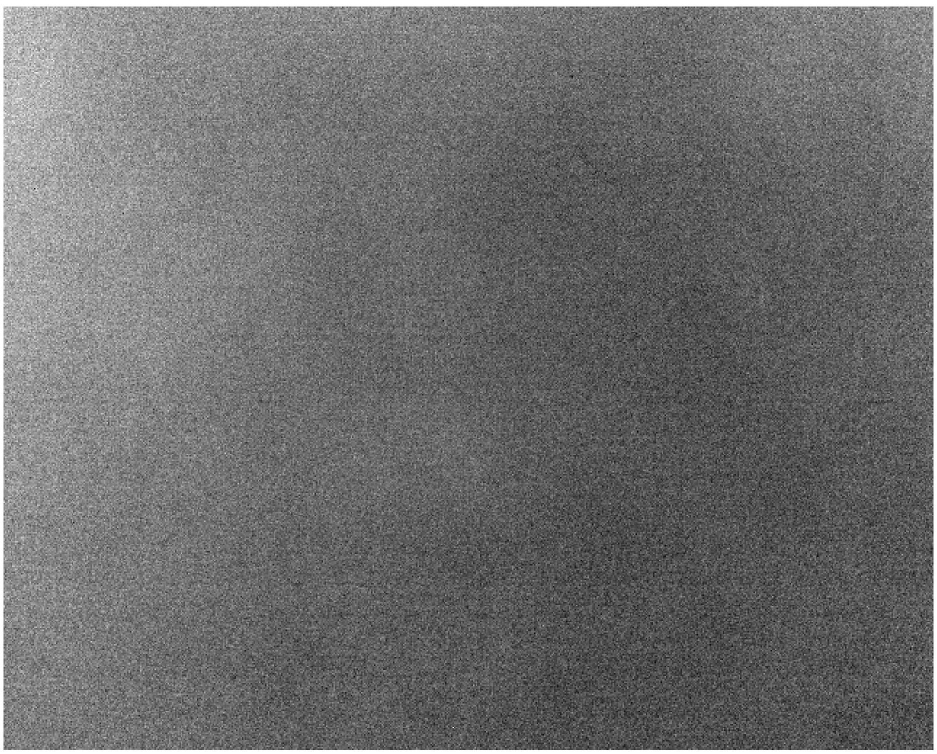

**Figure 61. A single frame of the 35 °C blackbody after the 25 °C ambient plate external offset update. Standard deviation is 10.4 counts, and the mean is 12,278 counts. The image is somewhat improved over what we saw when we first turned on the camera and looked at a 35 °C source.**

These are very stringent test cases, since users rarely ever look at scenes that are anywhere near as uniform as the out-of-focus surface of a blackbody! If one is trying to see tiny details in a very uniform target, it is imperative to do an offset update on a blackbody set to the same radiance as the target. Frame averaging will really help as well to reduce temporal noise.

## Summary of NUC Longevity

1) A two-point correction should be done with the lens in the desired focus position if possible.

2) A two-point NUC will work into the future, but an offset update is highly recommended for maximum uniformity. The gain table should not change over time. If you change to a different focal length lens, then you should do a new two-point NUC just for it.

3) The external offset update works best, because the external source is outside the lens assembly and corrects for non-uniformities created by the lens itself. If an external source is used, then the offset update should be done with the lens in the desired focus position. If the internal shutter is used for an offset update, then the focus position of the lens does not matter. In a laboratory setting, it is usually fairly easy to place an external calibration source in front of the lens. In a range application, the camera may be on a gimbal or tower and then there may be no choice but to use the internal shutter.

4) Any offset update should be done after the camera has been on long enough to stabilize thermally. If you need to start using the camera right away and can't wait for 2 hours until everything has stabilized, then you can just do periodic offset updates right before you collect a video.

5) Sometimes you can point the camera at an area of the scene that is very uniform, like the clear sky at zenith (straight overhead) and collect a sequence of 16 frames that can later be used to apply an external offset update to image data. This is done by a lot of range customers since they are already looking at the sky. If you take a series of images at lower elevation angles, you will see a vertical gradient in the sky because the optical thickness of the atmosphere changes with view angle – the sky looks warmer closer to the horizon. A cloudy sky works too, as long as the clouds are out of focus and the field of view is fully filled. The MWIR image of a clear sky at a 45-degree elevation angle shows the image and line profiles of Figure 62.

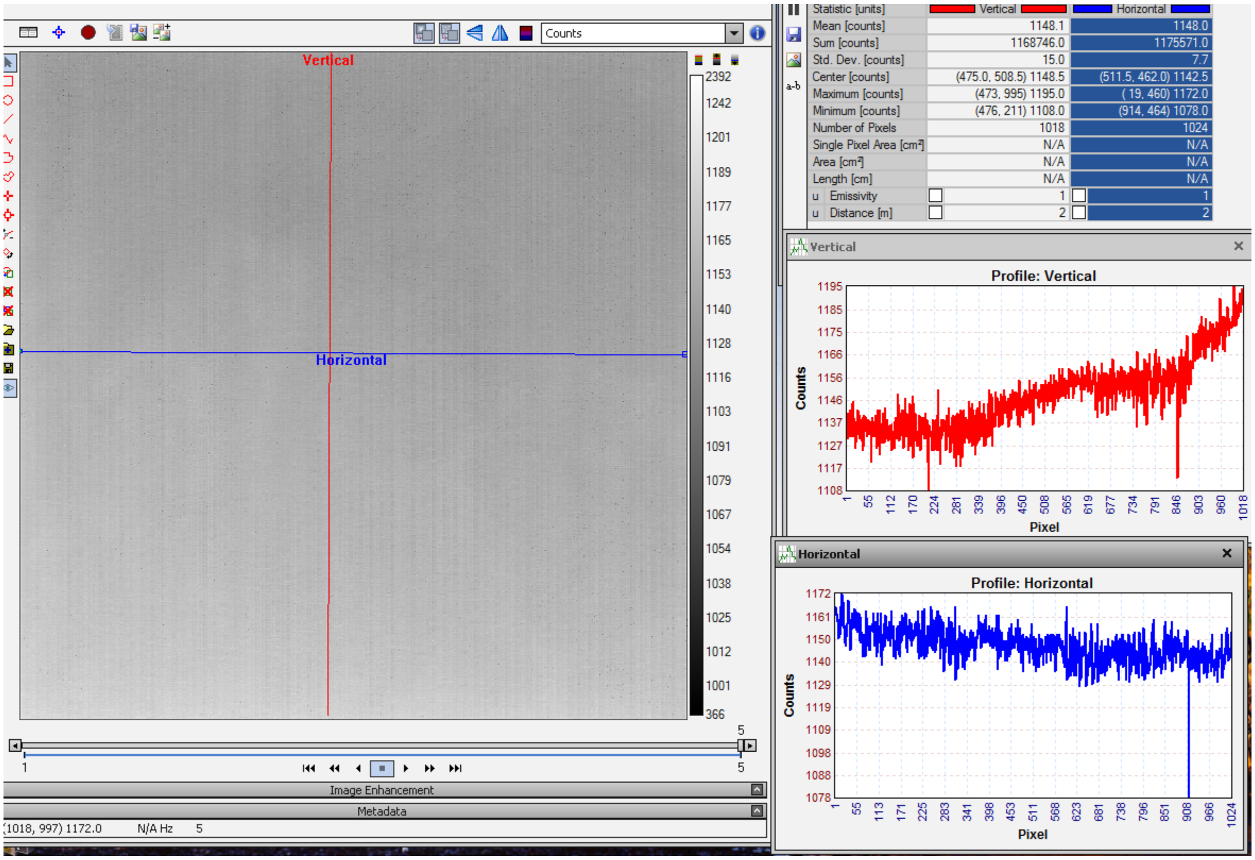


**Figure 62. Clear sky image at 45-degree elevation angle. There is a vertical gradient in the sky radiance showing how the optical thickness of the atmosphere varies with angle.**

## PC Side Correction NUCs

You can also create NUCs in ResearchIR, instead of using NUCs that live in the camera itself. The process of creating this type of NUC is called "PC Side Correction". Users do this when they do not want to "bake in" the NUC coefficients to the image data stream. Instead, they want to record the rawest data they can, which is uncorrected. Then they can apply the NUC in ResearchIR, and if they decide later that the NUC needs to be done better, or they want to add bad pixels to the bad pixel map so they get replaced, then they can make new NUC files and apply them. The utility called REdit is used to "point" the SFMOV file at the new NUC files.

There are a few things one should do before embarking on the use of PC-Side correction.

1) We highly recommend that you turn off the correction in the camera before you do a PC Side correction. You can either uncheck the box in the camera control Advanced tab that says NUC Digital Output (Figure 63):

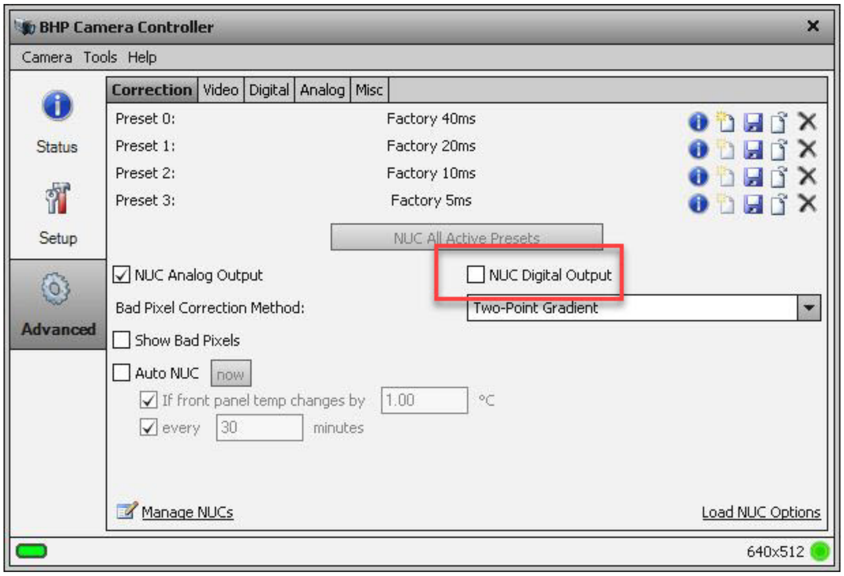


**Figure 63. Turning off the NUC on the digital output by unchecking the checkbox**

Or you can unload the correction for the active preset by using the X button (Figure 64):

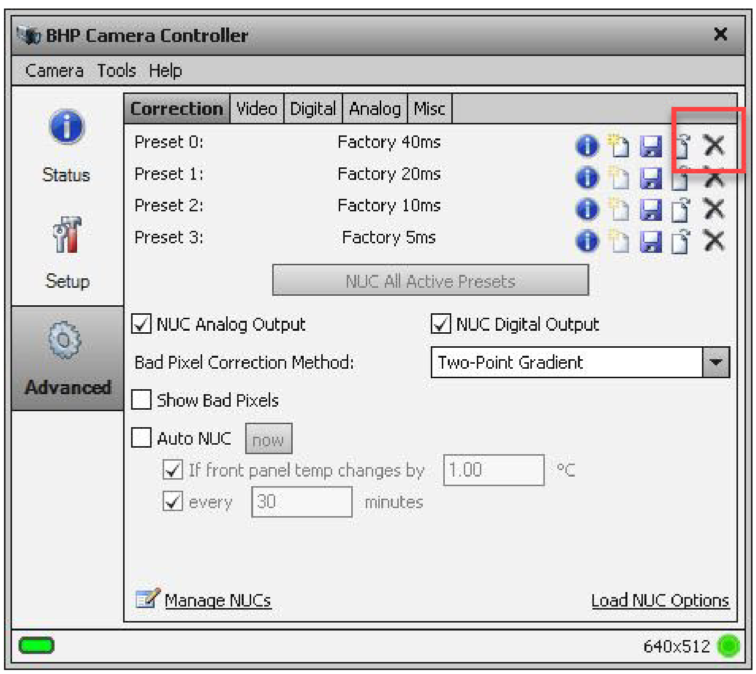


**Figure 64. Press the "X" button to unload the loaded NUC from a preset.**

Now the camera will have no correction applied. Until you create a "PC Side correction", the live images in ResearchIR will be un-NUCed and will look pretty bad. I suggest you uncheck the NUC Digital Output checkbox, rather than unload the factory NUC. I like to leave the Factory NUC loaded, and to leave the NUC Analog Output checkbox checked. That way, my analog video (Composite or HD-SDI) is NUCed. This is very nice to have for troubleshooting or focusing and composing the shot, using a tracker, etc.  NUC the analog video but do not NUC the digital video on the camera side.

It is an **extremely** good idea to save the camera's state at this point, so that when the camera power is cycled, it will come back in a state with the checkboxes checked or unchecked appropriately.  When you do "PC Side Correction", you have to manage the state of the camera and software more closely to make sure you are applying the NUC correctly and not double-NUCing!

It is also highly recommended to NOT use the *.ats file format for saved files.  Go to the Preferences menu in ResearchIR and make sure the checkbox shown in Figure 65 is unchecked.  The install default is that it is checked.  Otherwise, when you record imagery, you will get a single file called an ATS file, with the file extension *.ats.  The ATS file format will contain all the SAF files inside but they aren't accessible to do things like change the three NUC files in post-processing.

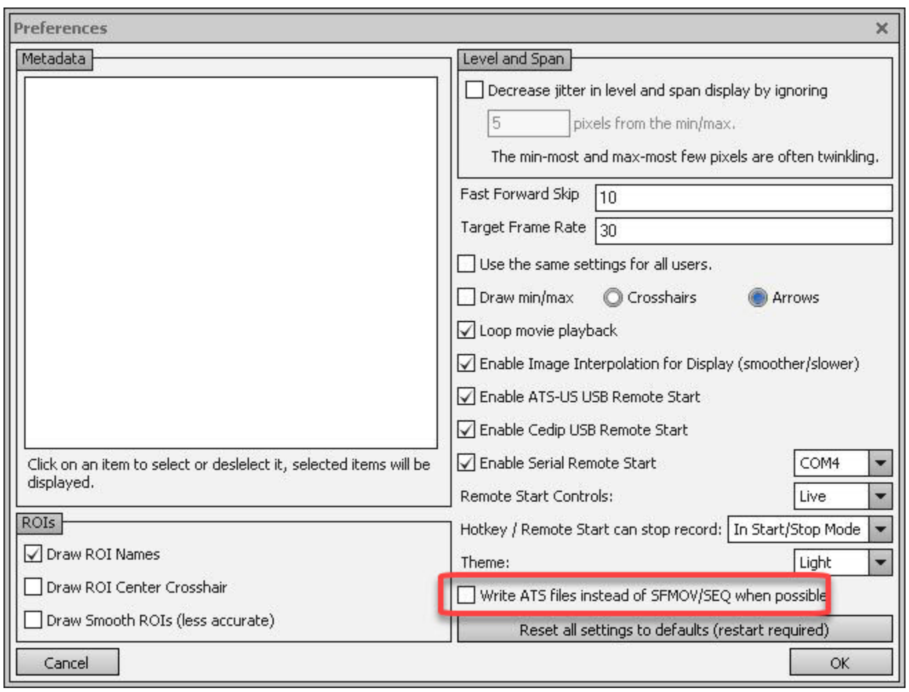


**Figure 65. Uncheck this checkbox to save images in SFMOV format.  The ATS format is nice when you have a factory NUC and calibration and are never going to want to change them, because you get a single image file with everything in it.  But you don't want to use it if you are doing PC Side NUCs.**

## Creating a PC Side Correction

To perform a PC Side correction, go to Camera/PC Side Correction/Perform (Figure 66):

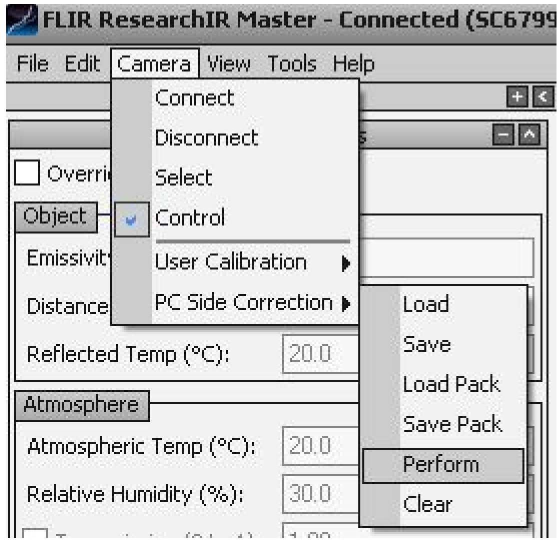


**Figure 66. Perform PC Side Correction**

If the NUC is turned off in the camera, then you will want to perform a Two Point Correction. See the Selection window for PC-side corrections in Figure 67. It says Update Offset, not Offset Update as it does in the camera controller, but it is the same thing, except the NUC source option is external only.

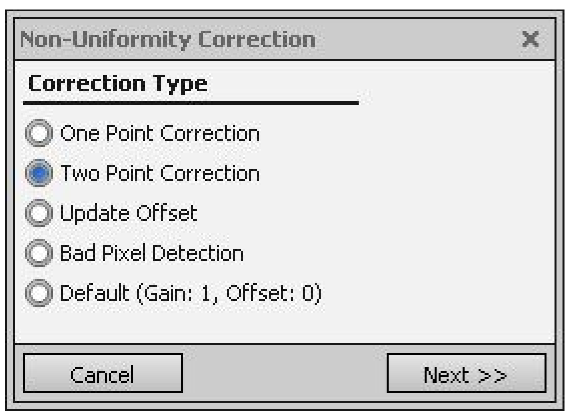


**Figure 67. Selection window for PC-side corrections. It says Update Offset, not Offset Update as it does in the camera controller, but it is the same thing, except the source option is external only.**

For starters, you can accept the default values in the bad pixel window in Figure 68. The tolerance is the tolerance in gain for a pixel. The gain limits are actually 1 / (1 +/- tolerance). The tolerance is checked on an intermediate value, and the gain is calculated as (1 / that intermediate value). If there are still twinkling pixels after you create the NUC, you can reduce the Twinkling pixel delta default value of 100 to something smaller to hopefully catch the twinkling pixel on the next attempt. You can also select the Bad Pixel Detection option and recalculate the bad pixel map without having to repeat the NUC. The Default (Gain : 1, Offset : 0) option creates a set of NUC files which are essentially pass-throughs – they do not affect the image data, which stays uncorrected.

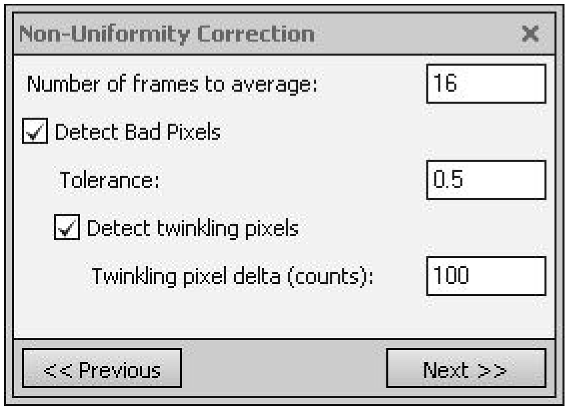


**Figure 68. Bad pixel detection criteria are selected here**

Put either the cold or the hot source in the field of view first, then hit Next, then follow the instructions, as shown in Figure 69. It does not matter which source you put in first, because ResearchIR figures out which is which and does the math correctly to get positive gain table values.

MAKE SURE THAT:

1. *The source is spatially uniform in radiance.*
2. *The source surface is out of focus. Usually the source is placed very close to the lens and this is not an issue unless you have very short focal length optics.*
3. *You are not getting a reflection of the coldstop. Surface reflectivity can create a blurry cold spot in the center of the image. This is why NUC sources have to have very diffuse (rough or matte) surfaces.*

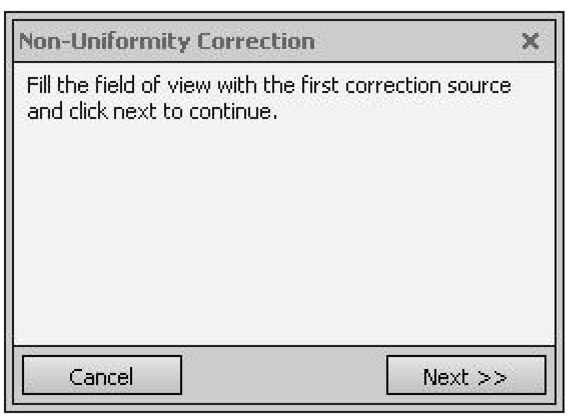


**Figure 69. Follow the instructions**

The second source you place in front of the lens should be like the first, just hotter or colder. The key ingredients for a good NUC source are:

- A high degree of uniformity
- A very high emissivity and diffuse (matte) surface finish

- The surface must be out of focus.

The second source should ideally give you at least 3000 counts of signal more than the first source. Put that source in place and then follow the instructions in the window (Figure 70):

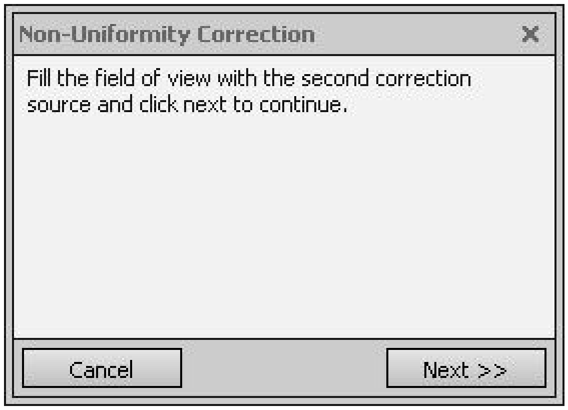


**Figure 70. Follow the instructions**

The statistics (Figure 71) suggest that this should be a good NUC. There are only 26 bad pixels. If there are thousands of them, then the NUC is likely to have been a failure. We will have to look at the images to evaluate if the bad pixel filter missed some. They can be added by hand but that can be time consuming, so it is often useful to redo the NUC while setting tighter limits on what constitutes a bad pixel.

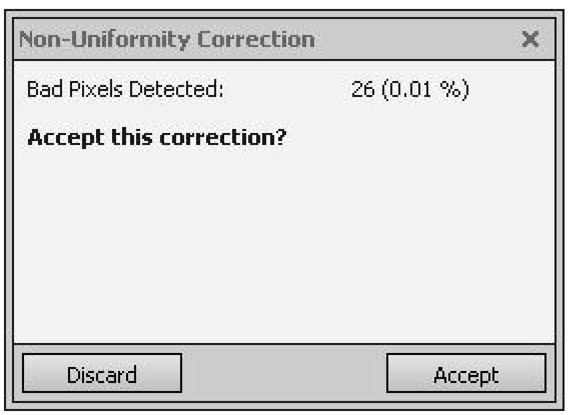


**Figure 71. Bad pixel statistics. Accept the correction if you think it will be ok.**

The NUC will now be applied to the image stream.

You can and SHOULD immediately save the three generated NUC files to the host PC, which is a very good idea, as they are not automatically saved in an obvious directory. ResearchIR handles the PC Side Correction in an interesting way. The corrections are saved in an Applications Data directory, and the software remembers the model and serial number of the camera and associates the PC Side NUC files with it, as well as the Preset that was active when the NUC was created. When the camera is reconnected to the camera in the future, the last saved

PC side correction is loaded into whatever that preset was. The integration time that was used is not embedded in the PC Side Correction files, so you have to make sure it is correct!! This process is not always intuitive, and it could cause problems if you have multiple PC Side NUCs created for different integration times. There is no way to know what PC Side NUC file set is loaded unless you manually load the set from files you saved to the host. It is always much better and safer to save the NUC files to a directory on the host PC when you create the PC Side Corrections, and manually reload them in subsequent sessions. Then you can even have multiple corrections for different lenses or integration times for the same presets. Go to Camera/PC Side Correction/Save, as shown in Figure 72-73:

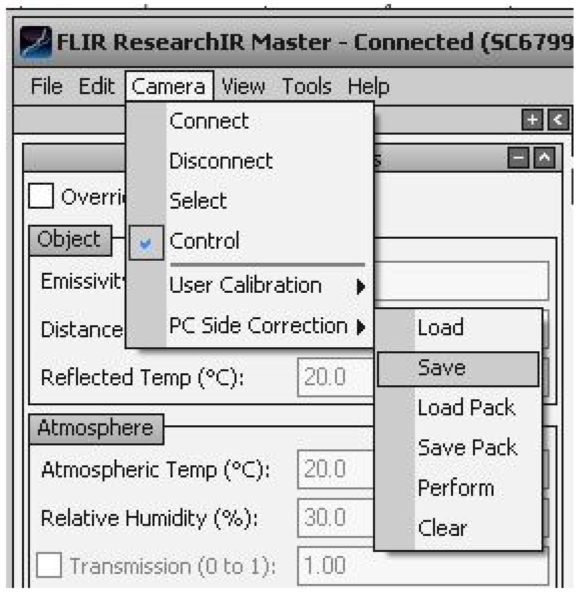


**Figure 72. Saving PC Side Correction to the host. This will be greyed out unless there is a PC Side Correction loaded in ResearchIR.**

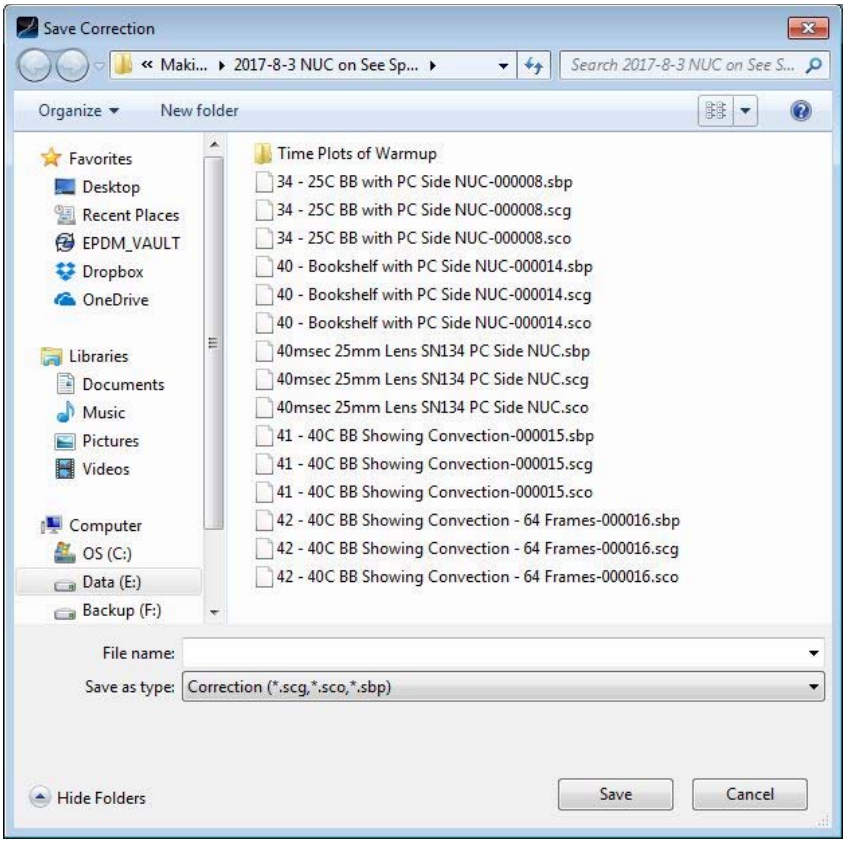


**Figure 73. Select the directory in which you want to save the three NUC files**

You need to type in a File name. There will be three files created, all with the same name and three different file extensions. I give the files a name which is descriptive: "40ms 25mm Lens SN134 PC Side NUC", as shown in Figure 74:

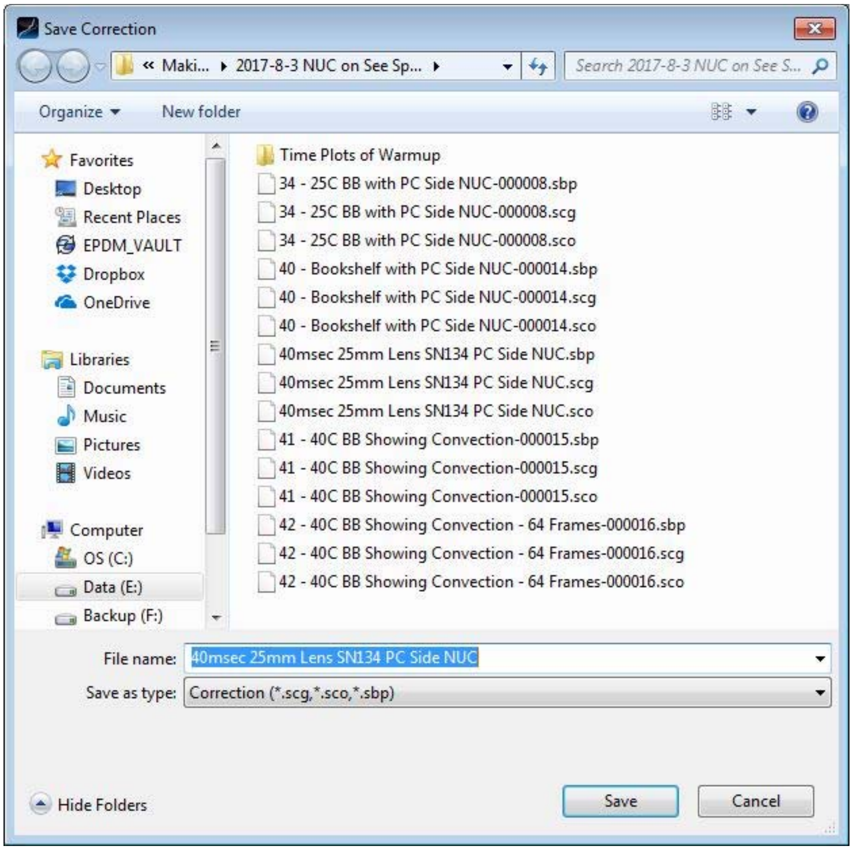


**Figure 74. Naming the NUC**

When I hit return, three files will be saved: a *.sbp, *.scg and *.sco. They are the bad pixel map, the gain table and the offset table respectively. These three files should be archived somewhere safe on the computer! It is a good idea to put the files in a folder along with a text document that describes how they were made, and a copy of the camera state file that was used at the time the NUC file trio was created.

## PC-Side Correction Pack:

If you do PC-Side corrections for more than one preset, then you can use Save Pack and all the files needed will be in there. Here in Figure 75, I was using the Save option, but you can see the Save Pack option below it in the list.

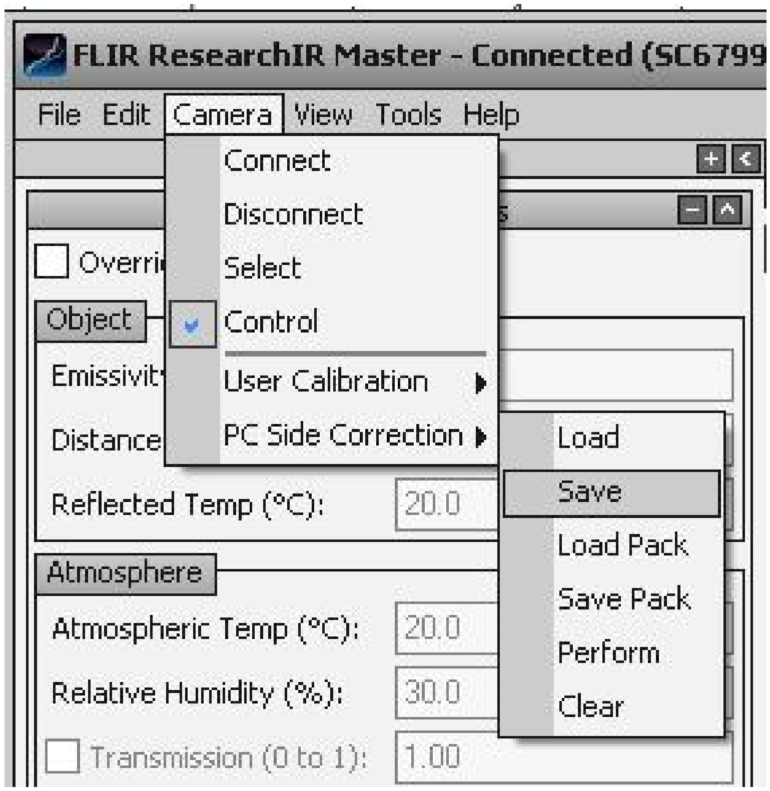


**Figure 75. You can Save or Save Pack in the PC Side Correction selection**

If you load the pack later and want to do superframing, or just have all the NUCS loaded, then RIR will apply the correction files to the presets in the same order. This is yet another thing to keep track of with copious notes and screen shots. Doing and using your own PC Side Corrections is powerful, but you must maintain meticulous records of exactly what everything was set to!

Figure 76 shows the three files that will be created in the directory if you use Save, as opposed to Save Pack:

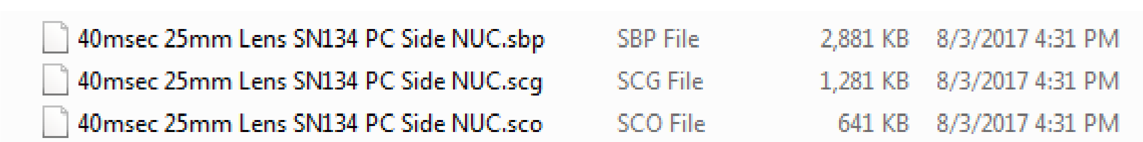

| | | | |
|---|---|---|---|
| 40msec 25mm Lens SN134 PC Side NUC.sbp | SBP File | 2,881 KB | 8/3/2017 4:31 PM |
| 40msec 25mm Lens SN134 PC Side NUC.scg | SCG File | 1,281 KB | 8/3/2017 4:31 PM |
| 40msec 25mm Lens SN134 PC Side NUC.sco | SCO File | 641 KB | 8/3/2017 4:31 PM |

**Figure 76. Three associated NUC files in a Windows directory**

The user is able to apply these three files to an SFMOV file even if the NUC files were not loaded into the PC at the time the SFMOV was collected. This is the power of the PC Side NUC technique – you are saving raw image data with no correction applied, and then you apply the correction in the PC. If the NUC was either not there or of low quality during the image data collection, it can be redone after and it has not been "baked into" the data. The other big advantage is that you can edit the bad pixel map in ResearchIR after the fact, and you can turn the correction and bad pixel replacement on and off to see its effect on the image. A customer may want to know exactly where the bad pixels or clusters of bad pixels are.

I can take some data on a 25 °C blackbody to evaluate the quality of the PC Side Correction I just created. I will get seven files if I also have a user calibration loaded, as shown in Figure 77:

| | | | |
|---|---|---|---|
| 34 - 25C BB with PC Side NUC-000008.cal | CAL File | 4 KB | 8/3/2017 4:29 PM |
| 34 - 25C BB with PC Side NUC-000008.inc | INC File | 1 KB | 8/3/2017 4:29 PM |
| 34 - 25C BB with PC Side NUC-000008.pod | POD File | 7 KB | 8/3/2017 4:29 PM |
| 34 - 25C BB with PC Side NUC-000008.sbp | SBP File | 2,881 KB | 8/3/2017 4:29 PM |
| 34 - 25C BB with PC Side NUC-000008.scg | SCG File | 1,281 KB | 8/3/2017 4:29 PM |
| 34 - 25C BB with PC Side NUC-000008.sco | SCO File | 641 KB | 8/3/2017 4:29 PM |
| 34 - 25C BB with PC Side NUC-000008.sfmov | SFMOV File | 10,243 KB | 8/3/2017 4:29 PM |

**Figure 77. Seven files captured with the use of a PC-side NUC in a Windows directory**

1. The *.cal and the *.inc file have the radiometric calibration data in them. The include file has a pointer to the *.cal file in it – one of the lines has a tag called CAFILE followed by the name of the file.
2. The *.pod file has parametric object data, or POD in it, most notably the time stamp and metadata from the image header.
3. The *.sbp, *.scg and *.sco files have the NUC and bad pixel map data.
4. The SFMOV has the raw digital counts values in it and metadata pointers to the include, POD and NUC files in it.

Let's take a look at the results of this NUC applied to a 16-frame sequence of a 25C blackbody, as shown in Figure 78. We can calculate the corrected uniformity from the image statistics. I will apply a 16-frame average to knock down the temporal noise.

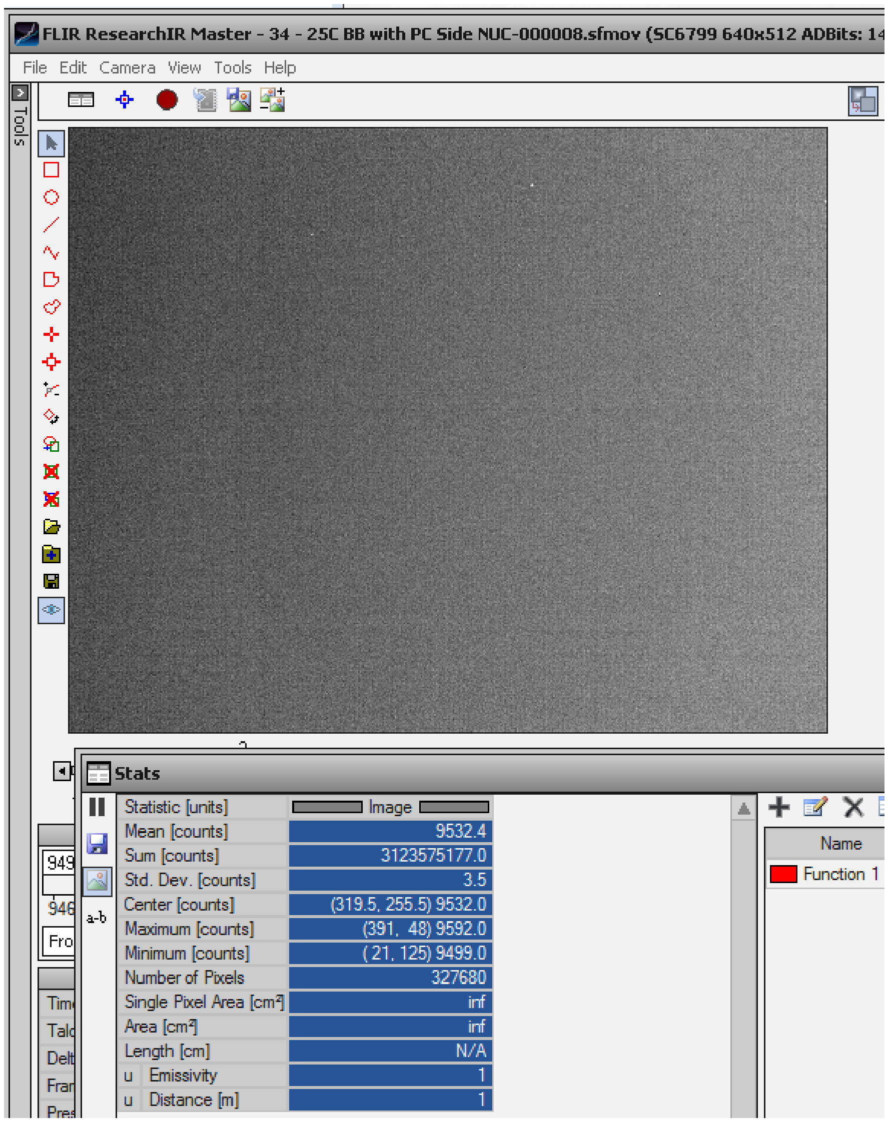


**Figure 78. PC-side correction applied to an image of a 25C blackbody with image statistics**

The mean is 9532 counts and the standard deviation is 3.5 counts. This is very uniform – the corrected uniformity is 3.5/9532 or 0.04%! If I turn off the frame averaging and look at a single frame, the standard deviation goes to 6.1 counts, so the corrected uniformity almost doubles. Notice the statistics – there is one pixel that is at 9592 counts, and is located at coordinates (391, 48). It is the white pixel shown in the red box in Figure 79:

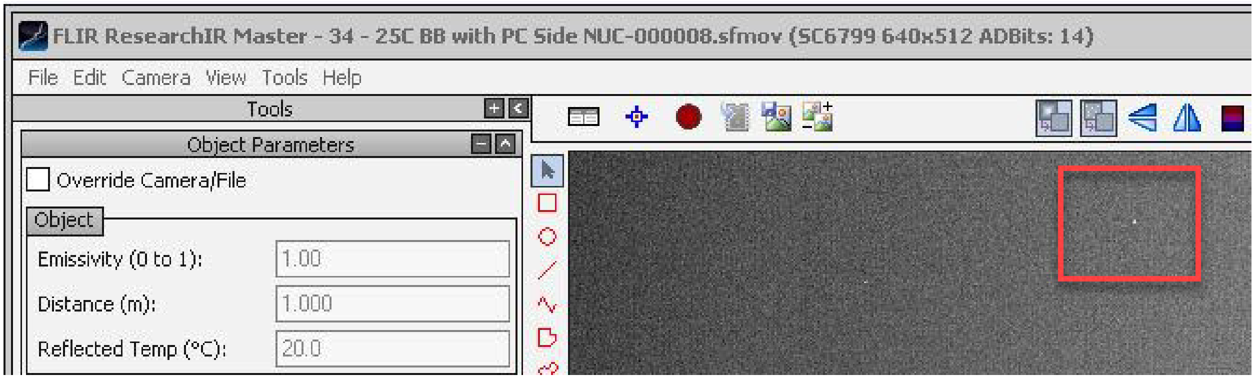


**Figure 79. Location of a "hot" pixel that needs to be added to the bad pixel map**

I can use the pixel picking tool to add it to the bad pixel map – then it will be replaced by a nearest neighbor. The operation of that tool is described in more detail in the ResearchIR user manual. It is a good idea to do the bad pixel replacement with the frame averaging filter on, with a 16-frame depth to it. You can see the bad pixels much more easily that way. I also recommend looking at a higher contrast scene (around the room you are in works well) and panning the camera around on the scene to find bad pixels. When you do this, the AGC can really make the bad pixels pop out.

I can go back and edit the bad pixel map in ResearchIR on a live camera with a loaded PC-side correction, or I can edit the map on existing data files made with a PC-side correction using the bad pixel picker tool, BUT it does not permanently save the new bad pixel map as you are editing it, until you select File/Save Bad Pixel Map.

You can also manually save the PC Side correction files again to save the edited bad pixel map for posterity – this is recommended to do after you finish editing the bad pixel map. In the example above, the bad pixel is skewing the standard deviation statistics, and it should have been replaced when the camera was live, right after the NUC was created, but it is easy to miss them!

Remember, if you add some new bad pixels with the picker, then save the PC Side NUC again, then RIR will incorporate the newly picked bad pixels into the *.sbp file that gets saved.

## Offset Update

Any infrared camera NUC will drift over time, and the drift shows up in the offsets. The image will slowly get more non-uniform. For a midwave InSb camera that is fully warmed up and stabilized, a fresh two-point NUC can "look good" for tens of minutes. However, the nonuniformities will start to creep back in, and they eventually become noticeable when looking at a uniform scene. Then it is time to perform an offset update, which is a temporary update to the NUCs offset table.

Many of the science cameras have an internal shutter that will put a uniform temperature black painted piece of metal in front of the warm window on the camera's dewar. This will present a uniform scene to the sensor. The only disadvantage of this shutter is that it is BEHIND the lens, so it does not correct out the non-uniformities that are introduced by the lens itself. The lens can introduce rings into the image, particularly on older lenses that were designed without as much care in placing internal glare stops.

The offset update feature is available both for NUCs that are stored in the camera, and for PC-side NUCs. For the in-camera NUCs, there is an Auto NUC feature which is usually on by default. The checkbox for Auto NUC in the advanced tab is checked. When it is checked, then the camera will perform an automatic shutter offset update every 30 minutes (default value that can be changed) or every time the front panel temperature changes by 1 degree C (default value which is also changeable). Figure 80 shows a screen shot of this with the checkbox unchecked.

I will always **make sure** that this box is **unchecked** before I do any critical data collection especially one that might take a while, because without warning, the camera will perform this operation and the video stream will have a short interval where the camera sees the internal shutter and not the scene!

**MAKE SURE** you save the camera state with this unchecked auto NUC configuration if you are about to do an important data collection that cannot be repeated. Suppose the camera somehow loses power five minutes before a rocket launch window. You power the camera back up but forget to un-check the checkbox, and it comes up checked because it was checked the last time the State file was saved. Now the camera might do an auto NUC with the shutter during the rocket flight and cost you a fraction of a second of flight image data! This is not good! If the camera was just restarted, then the front panel temperature can change rapidly, and then you can get multiple Auto NUC events within a few minutes.

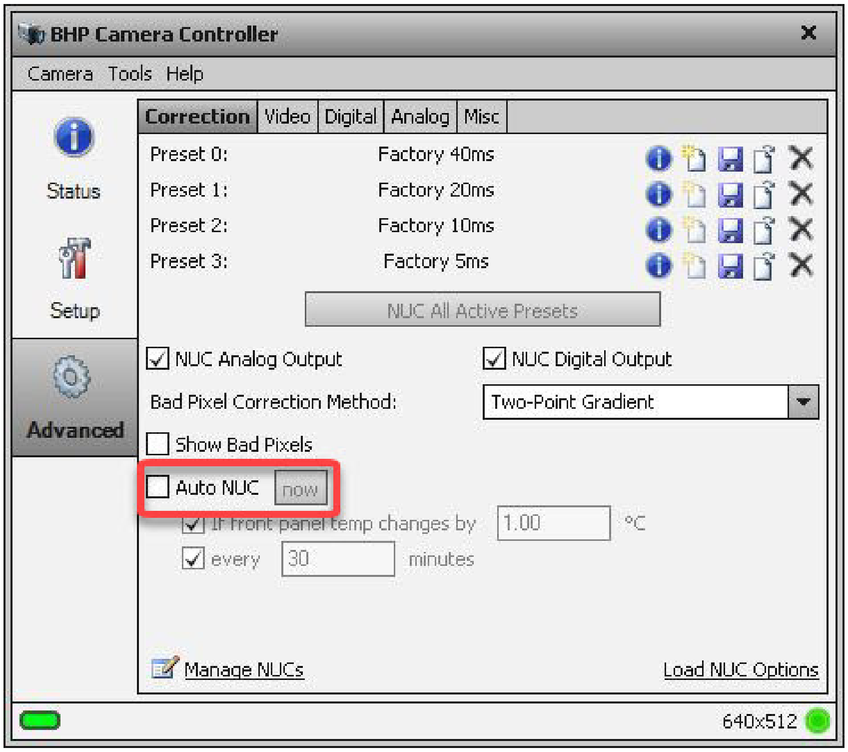


**Figure 80. Camera controller with Auto NUC checkbox unchecked. Having this box unchecked is a very good idea if you cannot afford to have any interruption in your image stream during an important collection!**

I often use the Offset Update option when I am conducting an experiment where I am taking data. I will perform it using an external calibration source right before I collect imagery. You can also use the internal flag option and then the camera will drop in the shutter (internal flag) behind the lens, but then you do not get the benefit of doing the offset update through the lens. There are options for the temperature of the internal flag in the dropdown menu shown in Figure 81. You will have to use this option if you cannot get to the camera physically to hold a calibration plate in front of the lens, which could be the case for a camera on a gimbal atop a tower. You can heat up the flag to some extent, but if you are operating at short integration times, then a 40C flag won't generate any significant signal on the sensor.

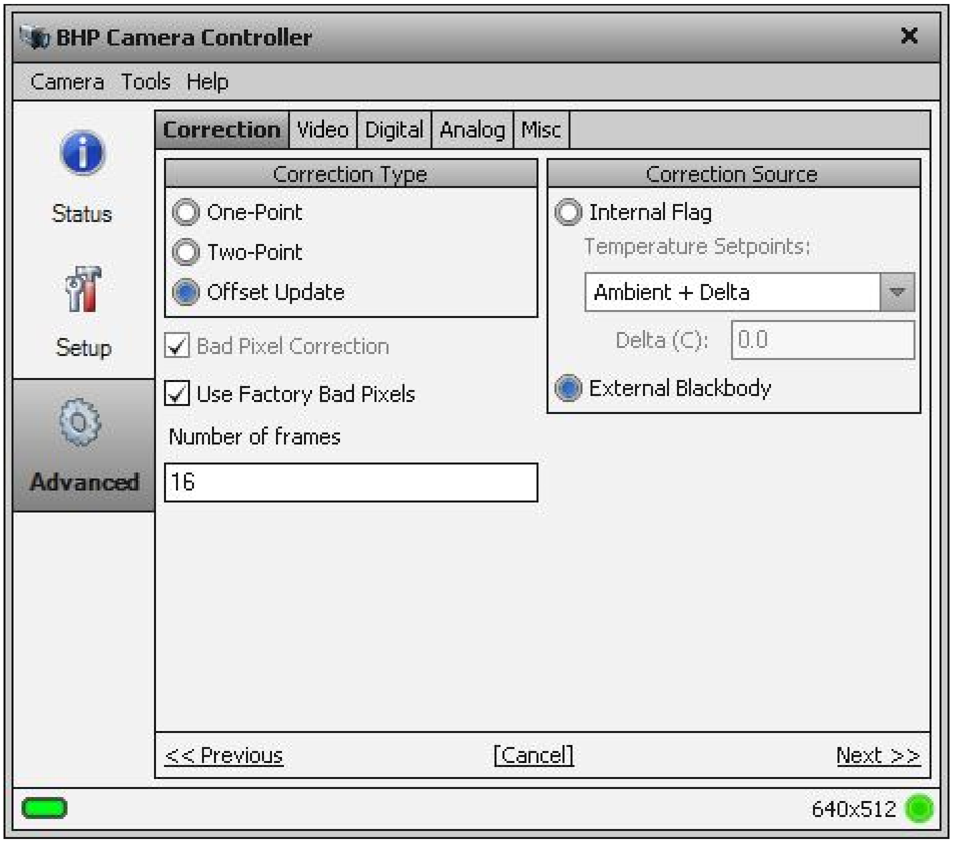


**Figure 81. Options for the correction source in a camera-side Offset Update include the internal flag and an external blackbody, which is the default option and the most preferred, assuming you can put an external source in front of the lens.**

I recommend you always do an offset update with an External blackbody source if you can, not the internal flag. The flag in the camera is behind the lens, so non-uniformities caused by the lens itself are not corrected out! You can see this in Figure 82. This is a particularly big deal for long focal length lenses that are out in the sun and are getting warmed up on one side. We have a motorized lens cover for the RS8300 series of cameras. This allows the user to do an external offset update on that camera, which has a 10-inch diameter objective lens and many elements in a 19-inch lens assembly.

One customer built their own high temperature area blackbody using an 18-inch aluminum square plate that is 0.5 inches thick. They mounted handles on it so that it could be held up in front of a telescope. The plate had a roughened surface on the front (done with a belt sander running 400 grit sandpaper) and high temperature black spray paint, the kind used for barbeque grills. The back of the plate has film heaters that could get the temperature up to 200 °C within a few minutes. This design is a great NUC source for an external offset update on a system running with a short integration time to look at very hot targets like missile launches. It should give better uniformity than an ambient temperature external source.

I always try to record a 16-frame image sequence of my external NUC source held close to the lens right before a data collection. It is straightforward to later use those 16 frames to apply an offset update to subsequent data files in an image processing environment like MATLAB. The advantage of this when using a camera-side NUC is that it is simpler and faster than using the camera controller to go through the offset update NUC process. If one is doing a long duration recording, the external source can be held in front of the lens for a second right at the beginning of the collection, then taken off after some frames of it have been collected. Then the offset update is embedded in the file, like a leader on a reel of a movie on 16mm film, and it can be extracted later and used to post process the rest of the movie. In fact, one can put the external source back in front of the lens at the end of the collection and grab more frames that can be used to generate a second offset update. If the collection is a long one, say 10 minutes, you can use the first 16 frames of the external source to update the offsets on the first 5 minutes of the collection, and the second 16 frames at the end for the last 5 minutes of the collection. It works very well.

Figure 82. Note the concentric circular ring patterns on the lower part of the image. This camera and lens system should have had an external offset update applied to it. The internal shutter will NOT make these rings go away. With this older lens (circa 2000 manufacture) , the focus should be set correctly for the object distance first, then the external offset correction performed, and then the focus should NOT be adjusted. If it is adjusted, the rings will come back and a new offset is required. This phenomenon of the rings has been much less of a problem with the newer HDC lenses.

**Here are some important points to keep in mind when performing an external offset update:**

- Use a flat metal plate with rough surface painted with flat black paint like Krylon 1602
- Preset the lens focus setting before you do the offset update.
- Hold at slight angle to the optical axis to prevent narcissus, especially if your plate surface is not a really good diffuse emitter, i.e. it is smooth. See Figure 83.
- The plate surface should be out of focus in the image.
- The plate should have the same in-band radiance as the mean of scene or target you want to image
- At least try to get the plate radiance into the linear region of sensor ROIC by heating it up in a convection oven or with film heaters on the back.
- Newer ROICs are very linear down to low well fill, so ambient-temperature External Offset plates work pretty well.

Figure 83. A circular external offset update correction source. The plate is made of aluminum with a roughened surface and flat black paint on it. Note the plate is held at an angle with respect to the optical axis to prevent any narcissus reflections of the cold stop. The lens is focused to infinity and thus the plate surface is way out of focus, as it should be. The lens in the figure is a 200mm unit, so you do not have to hold the plate too close for it to be out of focus.

## Changing the Focus Setting

When you change the focus setting of a lens, it does change the uniformity of the image. Here is a camera image of a 25 °C blackbody right after doing an external offset update using the same blackbody. The lens was focused at a range of about 0.5m right before the external offset update was performed and the focus ring was not touched. The result shown in Figure 84 is a very, very uniform image, but it is a little meaningless since I am looking at the same target that I just used to do an offset update. Of course, it will be extremely uniform.

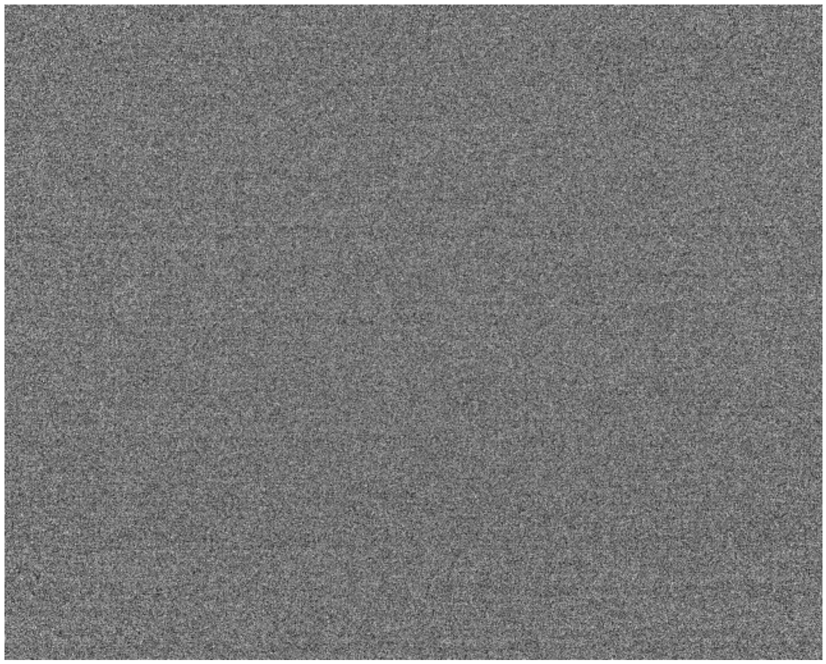

**Figure 84. Image of a 25 °C blackbody right after using said blackbody to perform an external offset update**

But now I rotated the focus ring to focus at a distance of 2m, while still looking at the same blackbody source. Look at the non-uniformity introduced by the lens itself (Figure 85):

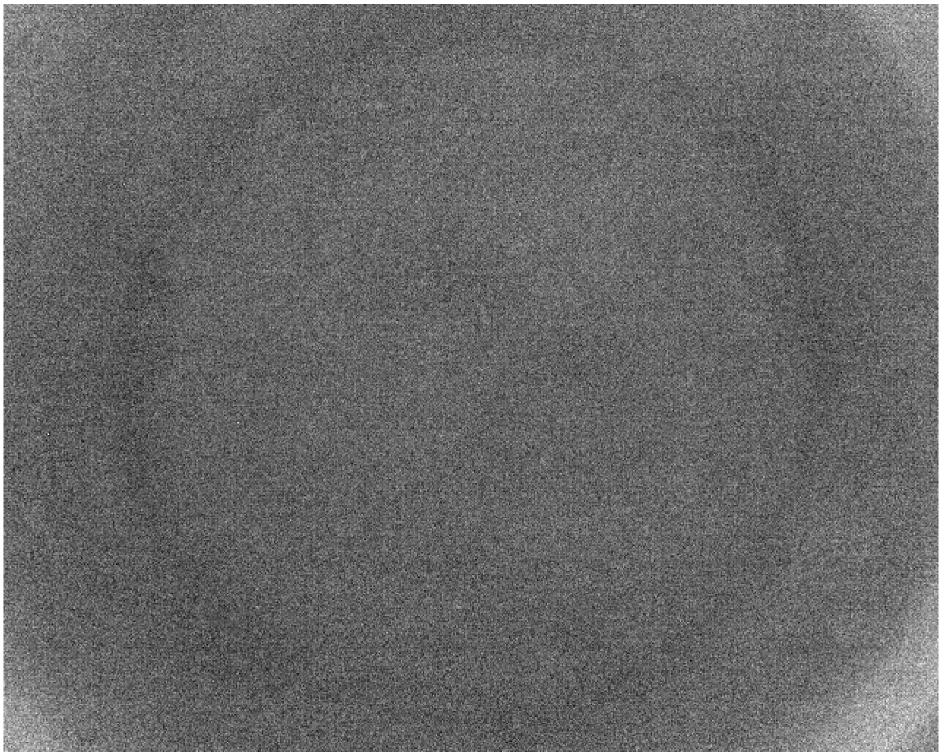

**Figure 85. Nice NUC spoiled by changing the lens focus**

This is why you always want to do an external offset update with the focus set to whatever you want it to be during data collection! Sometimes you might be changing the focus around while data collecting, in which case, do an external offset update before each collection with the focus set correctly, if you can.

Figure 86 is the image after a fresh external offset update done with the new focus setting:

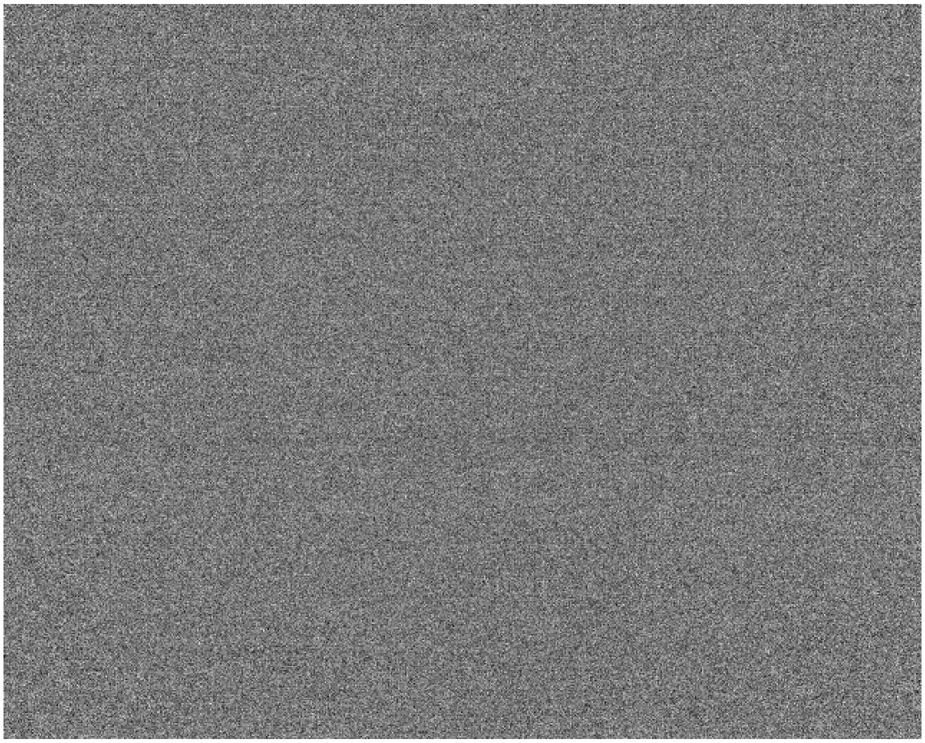

**Figure 86. Fresh external offset update done after focus was changed**

And here (Figure 87) is an image taken with this fresh NUC of a scene about 2 meters away:

**Figure 87. A fresh offset update makes images look great**

## Why NUCs are Made with Multiple Image Frames

When a thermal IR camera looks at a blackbody target above a certain temperature, usually about 15C above ambient, it may start to see convection patterns that swirl across the image. These patterns introduce low frequency spatial noise into the image. The best way to prevent this spatial noise from getting "baked into" the NUC is to take multiple frames and average them so that the "swirls" get averaged out.

Figure 88 is a pair of sequential frames from a 16-frame sequence of a 40 °C blackbody:

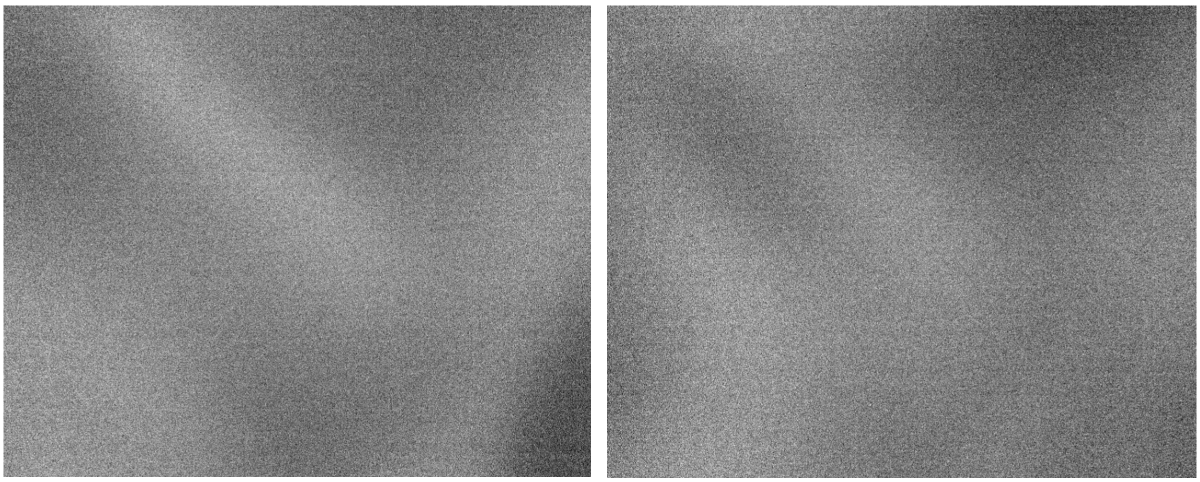

**Figure 88a-b. Two frames from a 16-frame image sequence showing convection patterns (swirls)**

Here (Figure 89) is the average of the 16 frames. There is still some low-frequency image noise remaining.

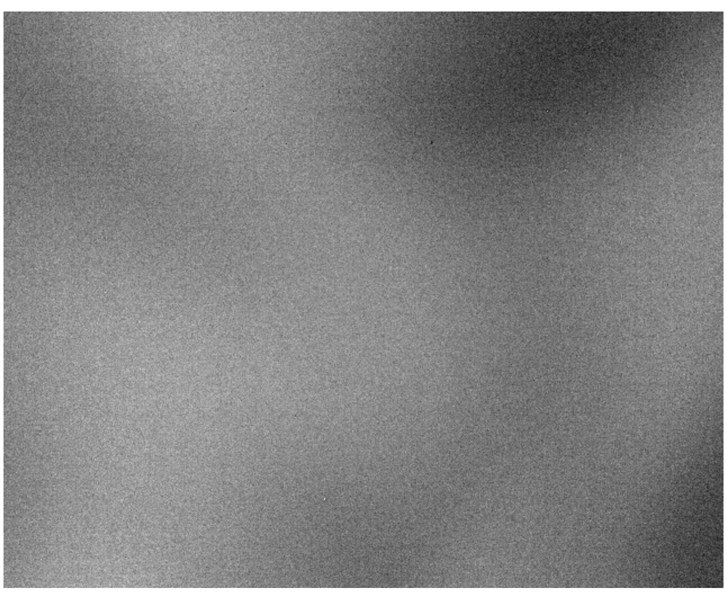

**Figure 89. Average of 16 frames**

If I average all 64 frames, the image gets even more uniform in terms of the convection swirl patterns, as shown in Figure 90. Note that the sprinkling of anomalous white pixels in the image really stand out when you average many frames, because the temporal noise has been reduced.

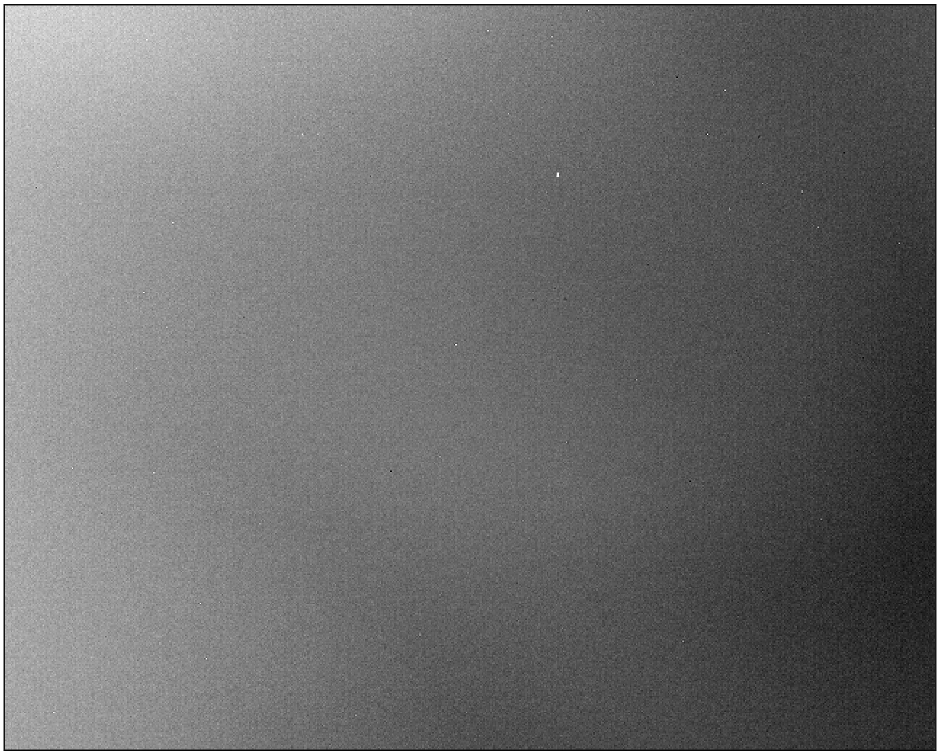

**Figure 90. 64 frame average that has less convection swirl patterns compared to a single frame from the sequence or an average of the first 16 frames from the 64 frames.**

It should be advantageous to take more frames in the averaging, like 64 frames or 128 frames. That is an option in the NUC wizard, as shown below in Figure 91. I do not know what the limit on the number of frames is. I would probably do 200 at most, but you can try it yourself!

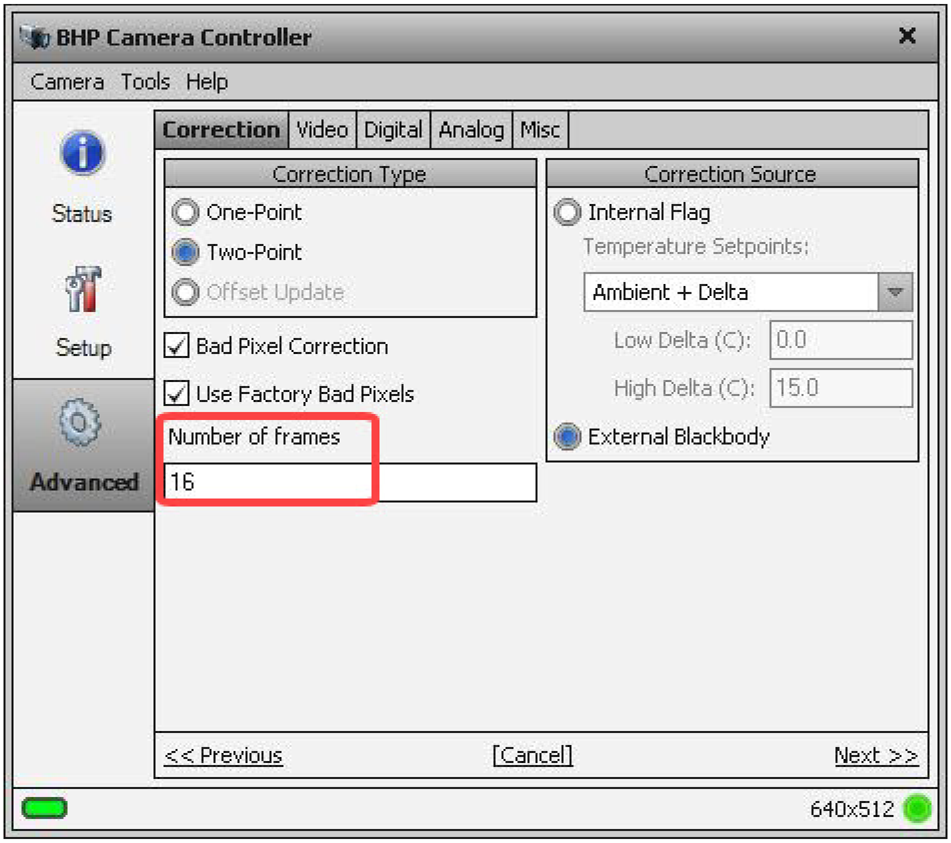


**Figure 91. Correction tab with number of frames field highlighted**

## Commercial Blackbody Sources

Blackbodies are a critical part of any IR camera power user's equipment inventory. They are required to perform user NUCs in camera, or PC-side corrections. For customers looking at hot targets, it is important to be able to present a hot blackbody to the camera just to perform an external offset update. Sometimes, it is not enough to do an external offset update with an ambient temperature calibration plate. Blackbodies are expensive, heavy and require substantial amperage to operate, especially the high temperature area blackbody units. Users need to consider the acquisition of blackbodies as part of the procurement of IR cameras.

This is a list of blackbody suppliers that we have used:

- Lumasense (formerly Mikron) - Very good, well characterized emissive coating. 2-inch circular emitter is a typical size. The M340 unit that goes from -20 °C to 150 °C is about US$9000. They make a cavity system that goes to 3000 °C.
- SBIR – very high cost (~$16k), very high performance, Mil qualified, very long lead times. SBIR also has sophisticated testing software called IR Windows that can be used with their blackbodies for automated testing tasks.
- CI Systems - high cost (~$12k), high performance. Great support and they know a lot about their blackbodies. They can also do radiometric calibrations of their blackbodies, so that when one sets the temperature to 100 °C, for example, the in-band radiance that is emitted is exactly what one would get from a perfect blackbody at 100 °C. This calibration is done for particular camera spectral responses.
- EO Industries - Expensive, Mil qualified, high quality controller. Nicely made!
- Infrared Systems Development Corp. – ISDC BBs are less expensive and quite decent quality ($8k price range). They make some really large area blackbody units.
- Omega – low cost, medium performance, not very uniform emitter. Designed more for spotmeter calibration.

### Facts about Blackbodies

- Very expensive devices that are often not considered at the time of camera procurement, though they should be!
- Can be long, long lead time items, as in 16 weeks! Repairs can take a long time as well.
- Expert users should get large area blackbodies for "test range" work - they fill the field of view of long focal length lenses that have large front apertures.
- Blackbody heads can be very heavy and awkward to manipulate in front of camera lenses in the field.

- 4" and 6" square are typical sizes for area blackbodies, 12-inch and bigger are available
- Area blackbodies up to 6 inches square go to a maximum temperature of 600C, unless there are special units that I don't know about.
- Cavities go to 1000 °C or more. SBIR cavities go to 1000 °C, CI Systems go to 1200 °C, Lumasense makes cavities that go to 3000 °C.
- Cavity blackbodies can be quite fragile and hard to ship reliably. I have had two cavity blackbodies break in shipment. The cavity heads should be packed extremely well in oversized hard cases with proper foam inserts. It is not enough to wrap them in some bubble wrap and put them in a cardboard box! When my SBIR cavity broke, it cost $5000 to have it fixed and recalibrated and the repair and recalibration took several months!
- Blackbodies should be recalibrated at least once per year. This is not cheap and not fast to have done, especially on cavity blackbodies. Blackbodies can drift in terms of absolute emitter temperature accuracy, but also spatial uniformity. Figure 92 shows a blackbody that is two years out of calibration, and the Peltier heating/cooling modules are not all tracking each other. There is a 3.9 °C variation across the emitter surface.

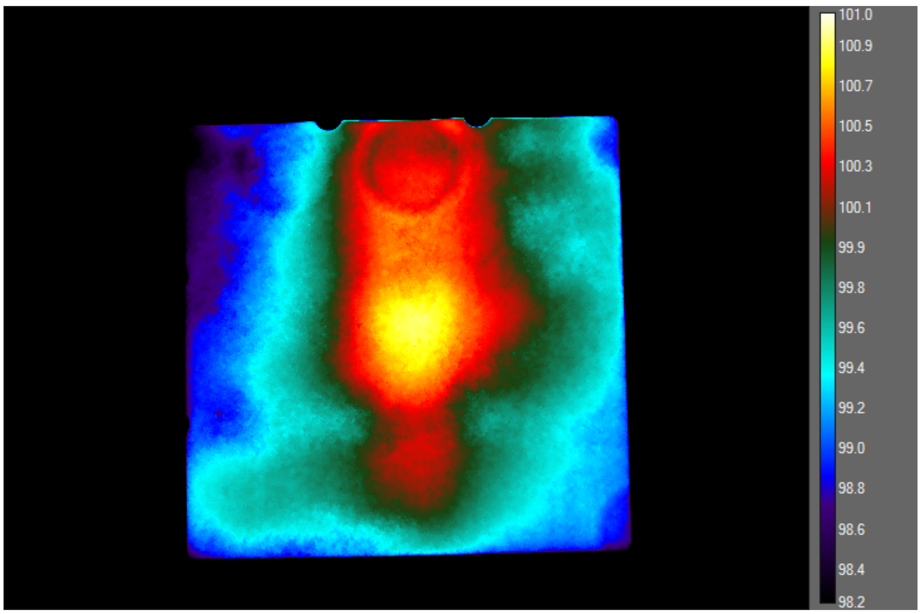


**Figure 92. Blackbody that needs service and is two years out of calibration.**

**There is significant non- uniformity apparent.**

SBIR has a new Infinity blackbody product line with a special coating on the emitter surface. The coating is called Vantablack and has emissivities approaching 0.999 in the midwave IR band as shown in Figure 93. The Infinity units (shown in Figure 94) have a shroud in front of the emitter to reduce the chances of physical contact with the emitter surface, which is very fragile.

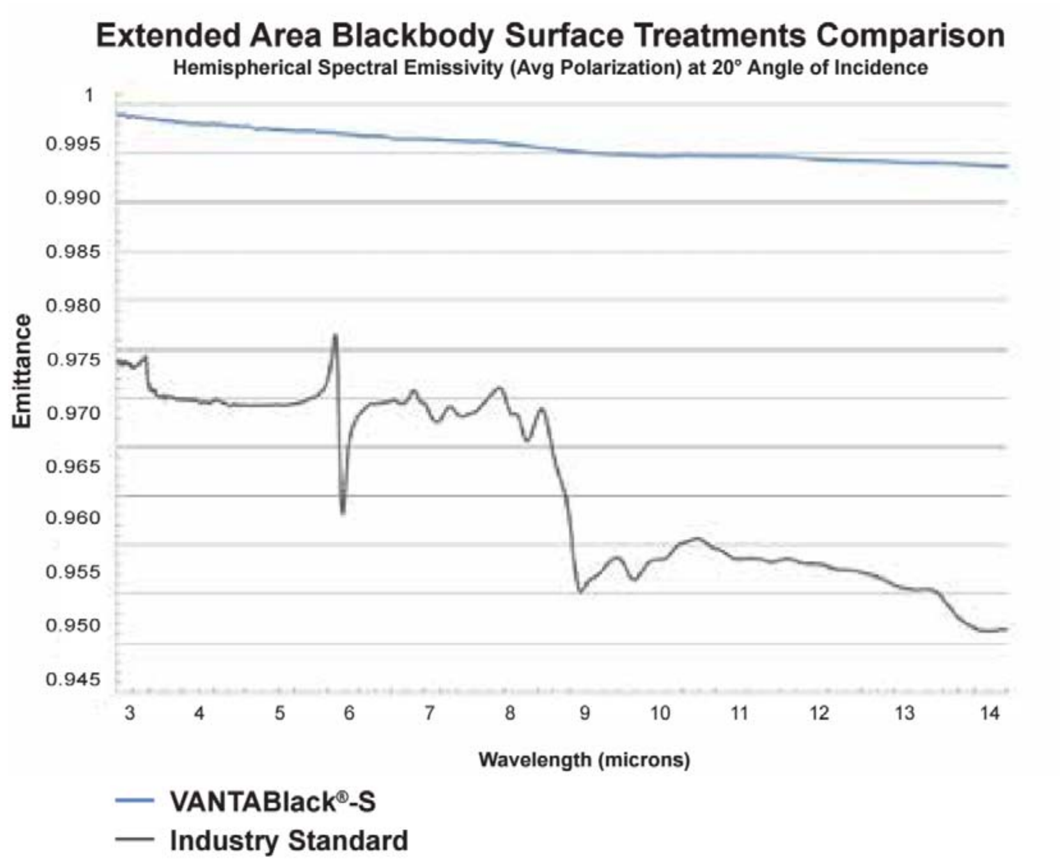


**Figure 93. Vantablack emittance curve compared to a typical blackbody black paint coating**

**Figure 94. SBIR Infinity series blackbody with protective emitter shroud**

## High-Temperature Area Blackbodies

These units have a fast slew rate up to high temperatures but are slow to cool. I use a muffin fan to blow air onto the surface if I want to cool it faster. They tend to have 600 °C temperature limits, and in-band emissivity values around 0.95. The emitter is copper block painted with special black paint and special roughened metal finish. The surface is heated to a precise temperature by heaters on the back of the emitter. An SBIR high temperature area blackbody is shown in Figure 95. The emitter surface emissivity will drop over time since the paint gets "cooked". The emitter surface starts to look more "glazed" and brown colored. These blackbodies are used for non-uniformity calibrations at short integration times.

**Figure 95. A 6-inch square high temperature blackbody from Santa Barbara Infrared. This will go to 600 °C**


**Figure 96. Infrared Systems Development Corp. 12-inch blackbody that goes to 230 °C.**

Figure 96 shows a 12-inch square blackbody made by ISDC.  The temperature range is from ambient to 230 °C.  The aluminum sheet metal shroud keeps air currents from perturbing the uniformity.  This supplier also has a 12-inch model that goes to 500 °C.  This kind of blackbody is very good for radiometric calibration on a telescope with a large entrance aperture.

### Cavity Blackbodies

Cavity blackbodies are used as absolute radiometric sources.  They have intrinsically high emissivities due to their design, which is typically a long tapered conical cavity.  Figure 97 shows an SBIR cavity blackbody system with a one-inch aperture and an aperture wheel.


**Figure 97. SBIR 4100 series 1000C cavity blackbody head and controller**

Cavity blackbodies are normally not used for NUCing since the emitter is usually a one-inch circular area, though they can be used to NUC IR microscopes. The aperture wheels can be equipped with apertures which are laser cut into thin brass sheet.  The typical circular aperture diameters for SBIR blackbodies are as follows:

0.600”, 0.400”, 0.200”, 0.100”, 0.050”, 0.025”. 0.0125”

Cavities can be purchased with various useful accessories, including an aperture wheel, a filter wheel, a chopper wheel and a shutter. The aperture wheel can hold precision apertures which range down to sizes as small as a few thousands of an inch.  The filter wheels are designed for 1-inch diameter bandpass filters.

Cavities that can go to 1000 °C or hotter are designed to be controlled down to 50 °C, typically. Cavities can take a long time to stabilize at the setpoint temperature.  I have watched the mean counts in an ROI on the cavity emitter aperture continue to climb for many minutes after the controller indicated that the unit was ready! It is a good idea to allow plenty of time for the blackbody to stabilize at the setpoint temperature, as in 45 minutes to an hour if one is ramping from a cold turn on to 1000 °C, for example.  Because they take a long time to cool off, it is advisable to do the calibration data collection series in ever-increasing temperatures. The more cavities one can get access to, the faster the calibration process goes.

The long time to stabilization exception is the Lumasense ultra-high temperature series of blackbodies with 10kW of heating power, shown in Figure 98. They stabilize in under a minute and can only be operated by a total of about 30 minutes at their full temperature. They require water cooling and argon gas purging to prevent the cavity itself from burning up in operation. The lifetime of the cavity increases dramatically as the setpoint temperature is lowered.


**Figure 98. Lumasense (formerly Mikron) cavity blackbodies capable of very high temperature operation.**

## Factory NUC Corrections and Radiometric Calibrations

FLIR has several facilities that have a whole range of blackbodies for NUCing and calibration. The Nashua, New Hampshire repair facility is shown in Figure 99. The blackbodies are all turned on and stabilized to a variety of increasing temperature values, making it possible to do many calibrations in a short period of time with no waiting for blackbodies to stabilize. We can do factory NUCs with radiometric calibrations for a wide variety of cameras and temperature ranges, from -20 °C to 3000 °C. The blackbodies are periodically calibrated to maintain accuracy and uniformity in the case of area blackbodies. The cost of ownership of the ultra-high temperature cavities is significant – they require industrial levels of electrical power, water cooling and argon gas purging during operation, and have finite and relatively short cavity lifetimes due to the intense heating. Many customers outside of government laboratories and research facilities cannot afford to maintain a very expensive array of blackbody assets and turn instead to FLIR for calibrations.

For special calibrations that are different from our standard factory NUCs and calibrations, there is a special engineering request form that can be submitted by a FLIR sales representative.  If the SER is approved, then special calibration laboratory work can be done at an additional cost premium over standard NUC and radiometric calibration offerings.  There is a longer lead time associated with that work, and we do not guarantee that the camera system will meet any factory specification or the customer's requirements – this is a "best effort only" process.


**Figure 99.  Calibration "Arc" with over twenty blackbodies used for radiometric calibration at FLIR's Nashua New Hampshire facility. The camera is arranged on a radial arm that moves the camera in an arc to access any blackbody which maintaining a constant camera to aperture distance.**

A FLIR InGaAs camera that is being calibrated on a calibration arc in the manufacturing center in Niceville, Florida is shown in Figure 100.

**Figure 100. FLIR A6261sc camera pointed at 1200 °C Lumasense M330 cavity blackbody**

## Troubleshooting a NUC

The most common problem with a user created NUC is that the sources used were not uniform enough for the application. Many customers will try to do a "hillbilly NUC" where they will use the palm of their hand as a hot source. Human hands are just not that uniform across the palm! A better choice if you have to do a NUC quickly in the field is to use a person's bare back or bare stomach. The surface radiance of either the back or the stomach skin is usually more uniform than the palm, and it will make other people laugh to see you pulling up your shirt.

Here are three variations of correction on an image of a man parachuting in Figures 101-103. The first has no correction applied, the next one has a NUC created with the palm of the hand and an ambient temperature calibration plate, and finally the third image has an applied NUC that was created with a laboratory area blackbody source:

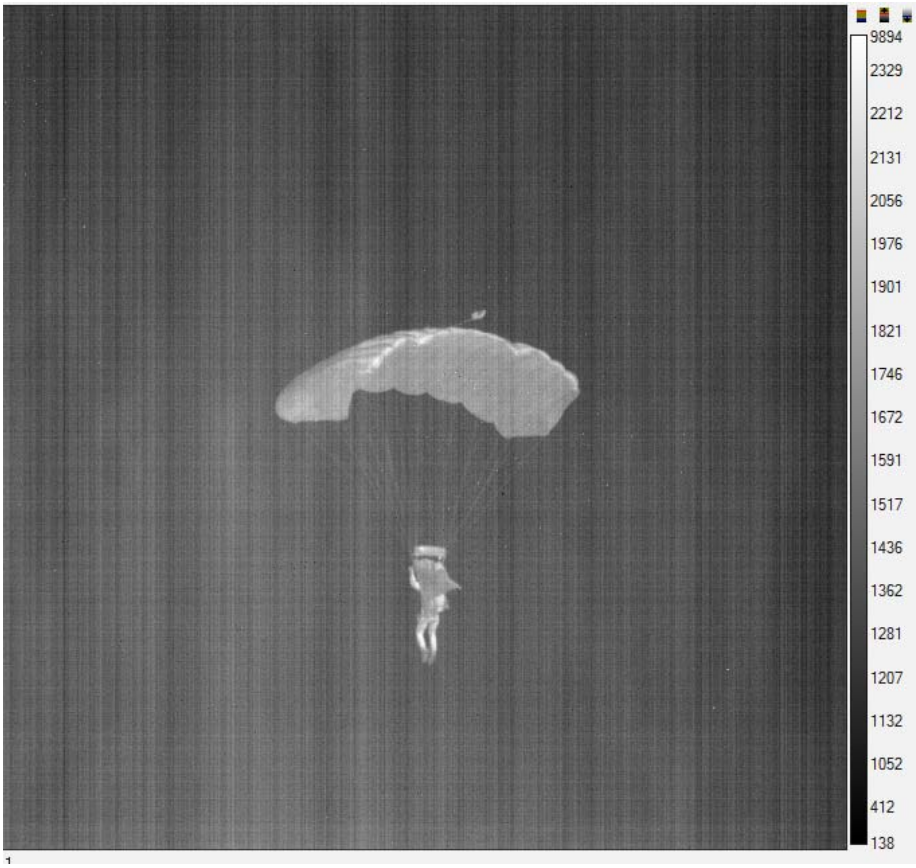


**Figure 101. This image has no NUC applied.**

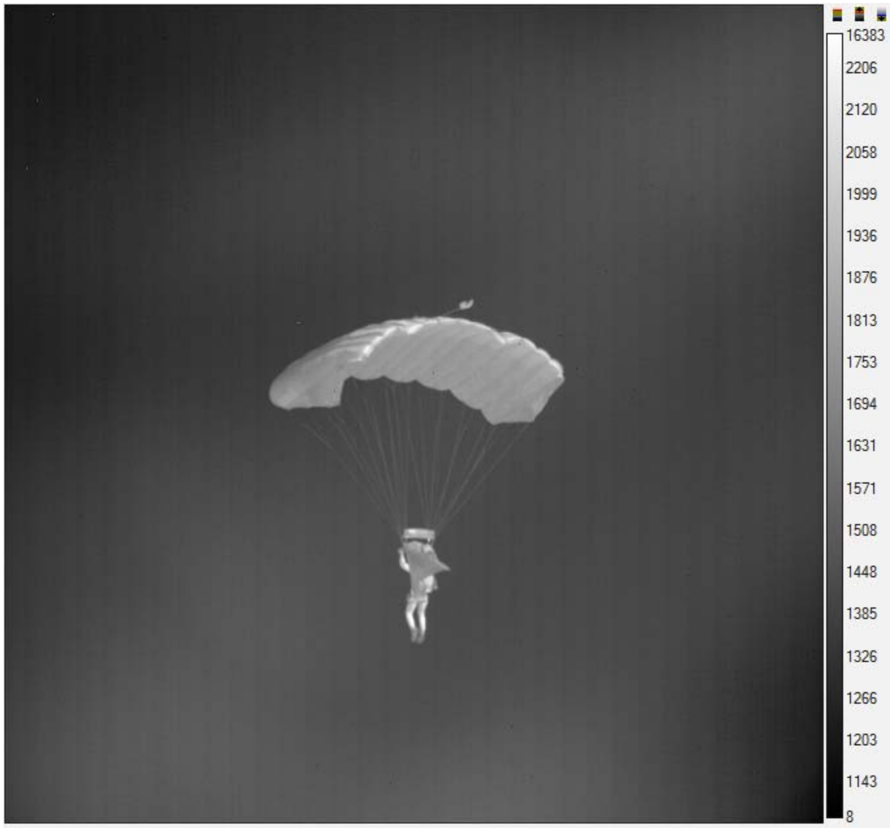


**Figure 102. This image has a NUC that was created with a calibration plate and my palm. It is sometimes called a hillbilly NUC! Note the low frequency shading variations from the top left to the bottom right corners. This NUC was created and then used on the day of collection to get decent enough looking imagery for camera pointing and focusing.**

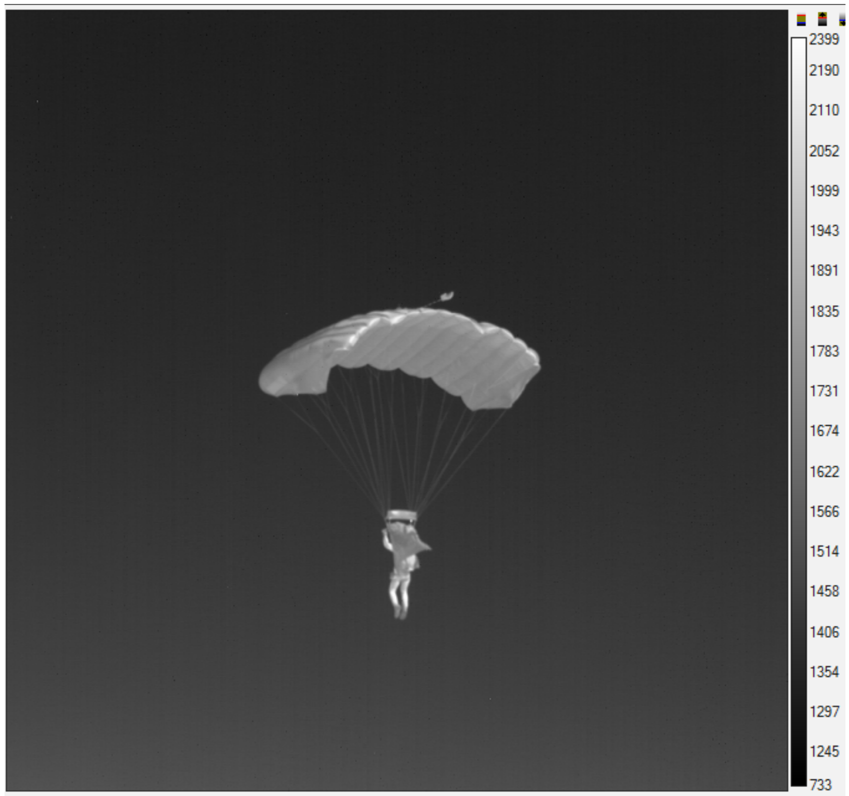


**Figure 103. This image has a laboratory NUC applied to it a week after the image was taken.**

## What happens when the NUC Sources are not in the Camera System's Linear Range?

I created a series of image sequences for a FLIR SC6000 camera with no NUC applied. The camera was operated at 47 microseconds integration time with a 25mm lens. The radiometric transfer function for this camera was shown at the beginning of this chapter and is shown again in Figure 104. I generated NUCs using various pairs of NUC source temperatures and applied them to images taken with the camera to show the right way and the wrong way to make NUCs. The right way is to use the camera within its linear range. If you don't, then the NUCed image will not look good, and you will see column noise. Here is the radiometric transfer function again. Not only does it show you how to calculate radiance from digital counts, it also shows you what the useful linear range of the camera is! It is a really, really good idea to measure this function on your camera so that you know what range of digital counts keeps the camera in its linear range. If you always keep the image pixels in this range, then a properly made NUC should work well and the radiometric calibration will as well. If the image goes outside the linear range, then you will at least understand the cause of the loss of uniformity and the breakdown of the radiometric calibration, and you can set your expectations accordingly, as in excluding those parts of the image from any radiometric analysis.

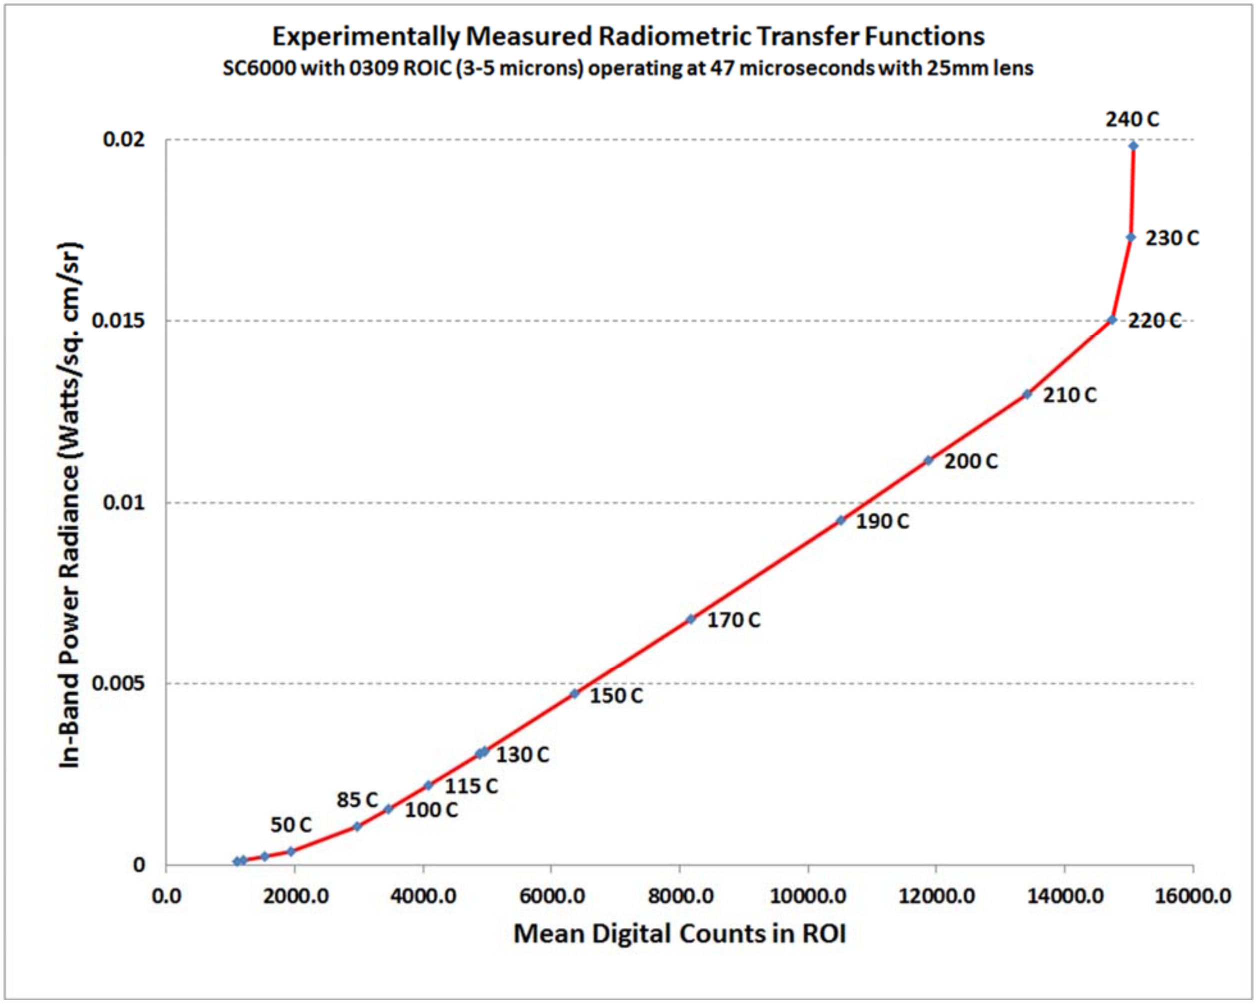


**Figure 104. Radiometric transfer function for an SC6000 camera operating at 47 microseconds integration time**

The useful linear range for this camera at this integration time corresponds to scenes with emitter temperatures between 100 °C and 210 °C. Below the radiance of a 100 °C blackbody, the ROIC is not very linear, as shown in Figure 105.

I will show you an image of a 150 °C cavity blackbody that was NUCed in two ways. The first way is the wrong way to do it. I used 50 °C and 130 °C for the two NUC source temperatures. You can see that 50 °C and 130 °C are NOT good NUC points because the 50 °C point does not lie on a straight line with the 130 °C and the 150 °C data points. This NUC will not do a good job of correcting any temperature scene!

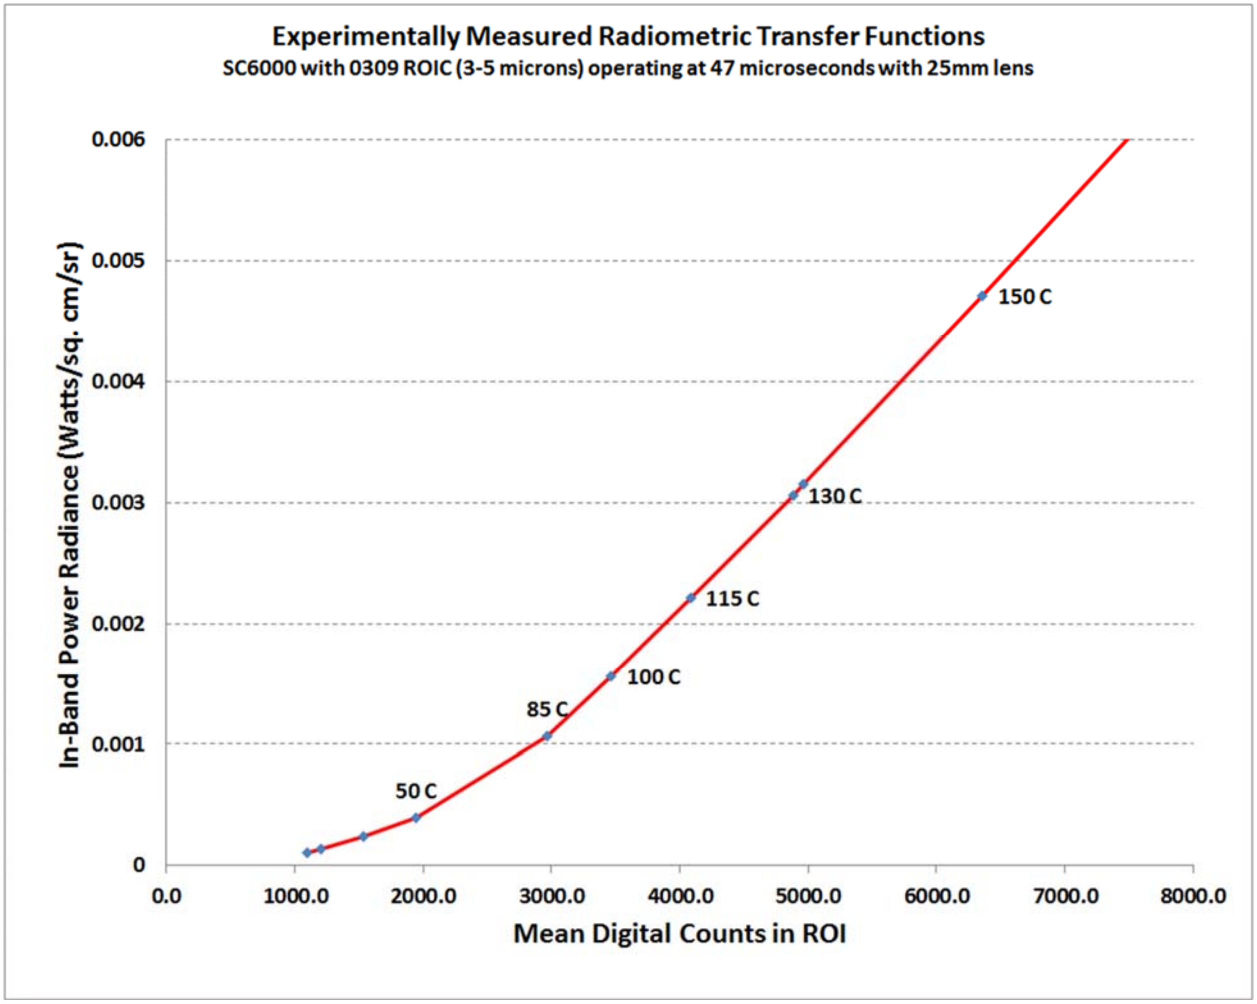


**Figure 105. Zoomed in radiometric transfer function showing that 50 °C, 130 °C and 150 °C radiance points do NOT lie on a straight line. The 50 °C radiance point is down in the non-linear portion of the ROIC.**

Figure 106 is an image of the 150 °C cavity with a 50 °C/130 °C NUC applied. You can see that the image of the cavity itself is poorly corrected – there is a lot of column noise. The ambient 23C surfaces around the cavity are definitely not being corrected well either:

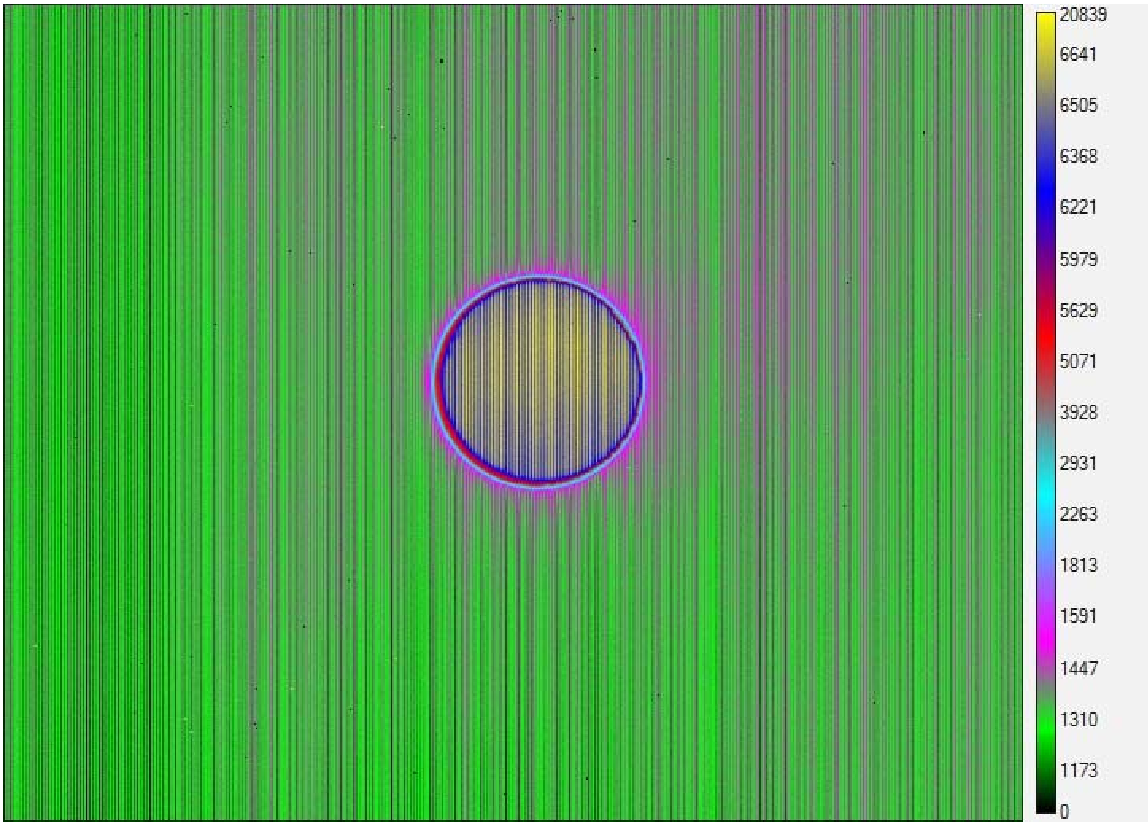


**Figure 106. 150 °C cavity image corrected with 50 °C and 130 °C NUC sources. The image of the cavity is poorly corrected.**

Now I will apply a NUC made with 100 °C and 130 °C blackbody sources. The image of the cavity is much more uniform, as seen in Figure 107. The 23 °C ambient surfaces look bad, because the correction does not work down at that low radiance level:

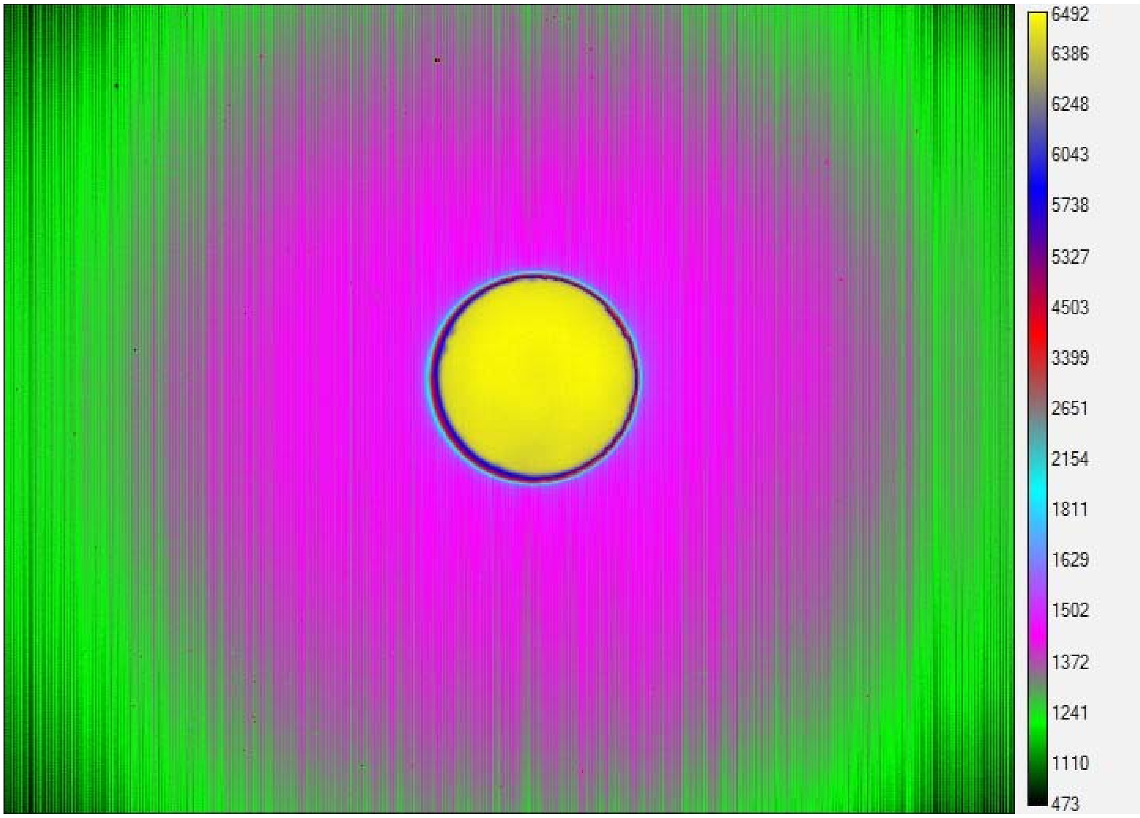


**Figure 107. The same 150 °C cavity image data corrected with a NUC that was made with 100 °C and 130 °C NUC sources. The cavity is much more uniform. The ambient temperature parts of the scene still look very poor.**

### NUCing over a Smaller Range

It is still possible to NUC a camera like this SC6000 for lower temperature scenes. Even though the ROIC is non-linear for radiance values below about 100 °C scene temperatures, we can still do a NUC over a narrower range of temperatures and use it, as long as the scene itself has a similarly narrow range of temperature. **You are doing a piecewise linear fit to a non-linear function, and in the limit of a narrow radiance span, the fit becomes better and better.** To illustrate this, I created some NUCs in the lower region of the radiometric transfer function, still operating at 47 microseconds integration time. First, I made a NUC using 25 °C and 35 °C area blackbodies. These two sources only give a difference of 330 digital counts between them, which meets the definition of a narrow span in the ROIC dynamic range, since this is only 2% of the 16383-count digitizer range. I then pulled back the camera and imaged the same two blackbodies and the room around them, as shown in Figure 108. The image quality is quite good, and this NUC could be used, as long as the temperature of the scene stays in this range.


**Figure 108. 20 °C blackbody on the left and 35 °C blackbody on the right. The NUC that is applied was generated with the two blackbodies at the same 20 °C and 35 °C temperatures.**

Now look what happens if I use the same 20 °C/35 °C NUC on a scene with 15 °C and 25 °C blackbodies. The uniformity on the 15 °C blackbody (the left one in Figure 109) looks worse than when that same blackbody was at 20 °C. The piecewise-linear NUC works, but only in a narrow temperature range!

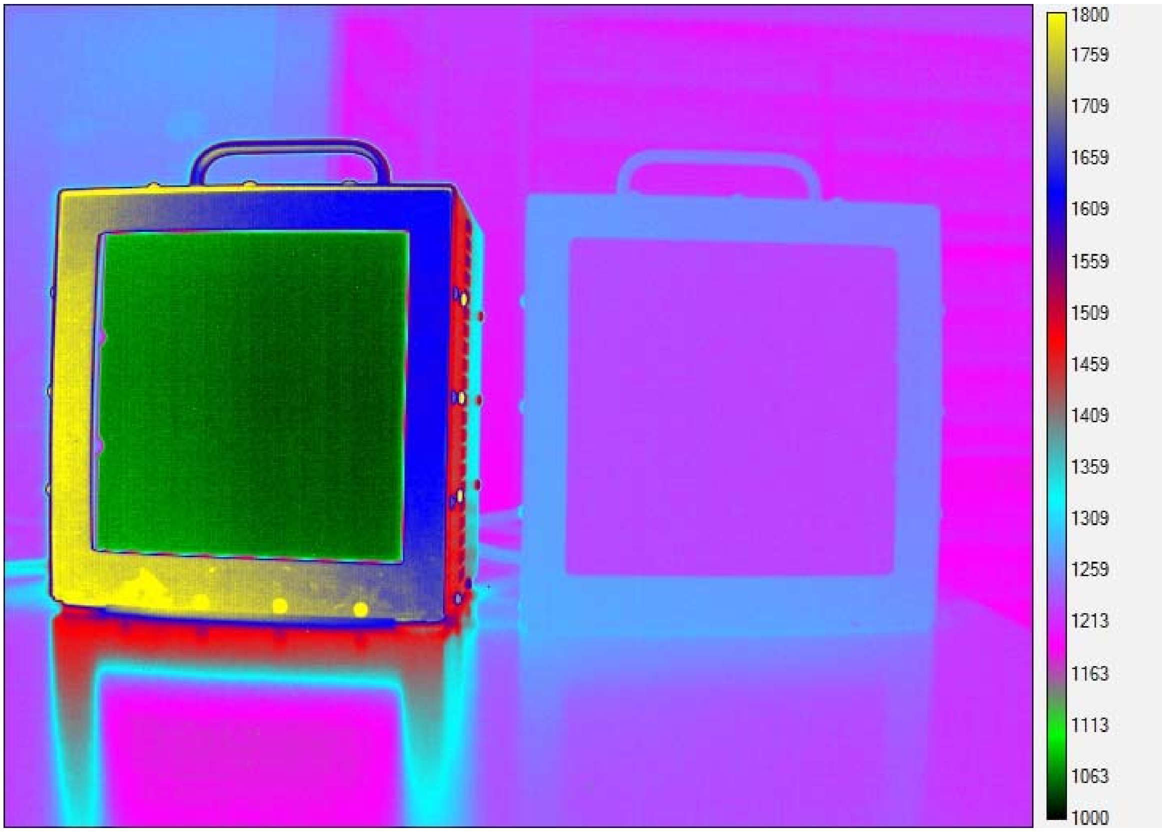


**Figure 109. 15 °C (left) and 25 °C (right) blackbodies. The NUC was generated from 20 °C and 35 °C blackbodies. There is a clear reduction in the uniformity of the 15 °C blackbody compared to the same blackbody at 20 °C in Figure 108.**

I can "clean up" the 15 °C blackbody by using a NUC generated from 12 °C and 20 °C NUC source temperatures, as shown in Figure 110, but now the warmer parts of the image like the yellow-colored bezel around the left blackbody have bad correction. There is really no way to get a two-point NUC to perform well down in this non-linear region of the ROIC with any kind of decent dynamic range. This suggests a maxim:

***Two Point NUCs are linear transformations, and as such, will only work well within a linear range of a camera sensor***

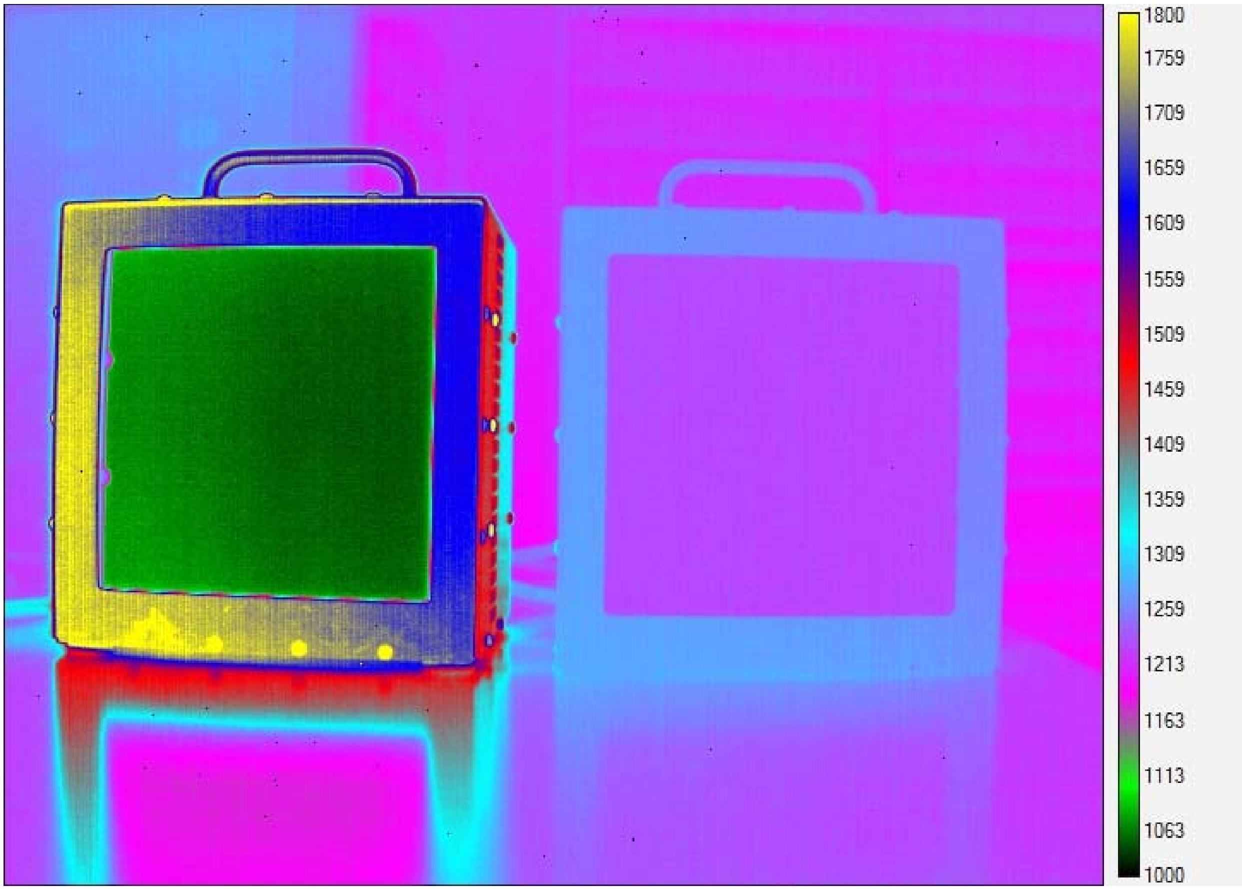


**Figure 110. 15 °C (left) and 25 °C (right) blackbodies. The NUC was generated from 12 °C and 20 °C blackbodies. There is an improvement in the uniformity of the 15 °C blackbody, at the expense of making the warmer parts of the image look worse.**

## Managing NUCs

There are several useful tools that can be used to manage NUCs. We have discussed some of them earlier, but here they are again. Consider the icons to the right of the loaded NUC names.

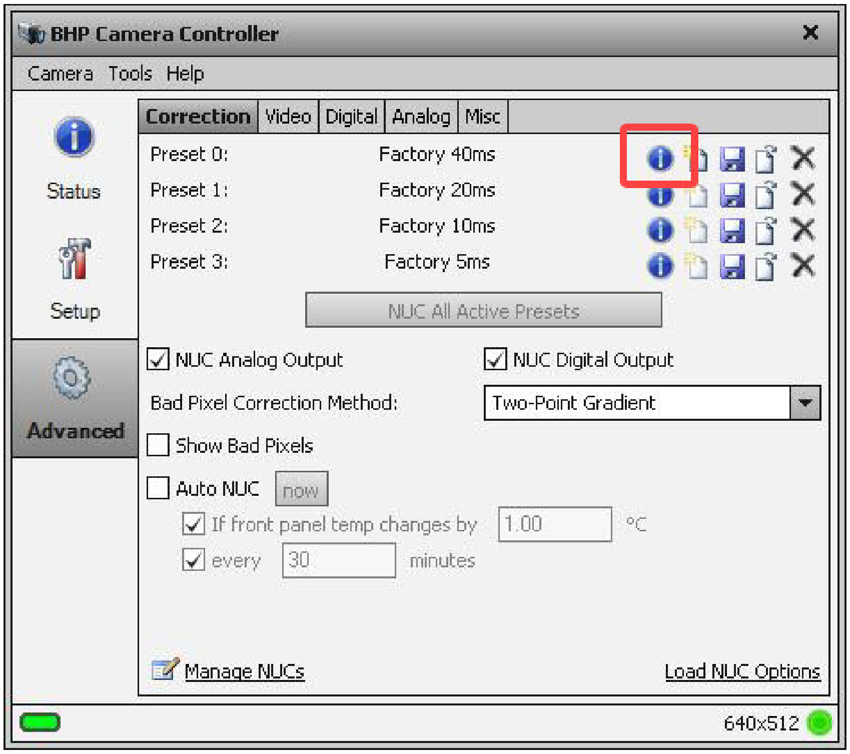


**Figure 111. NUC information icon shown in the red outlined box**

See Figure 111. If you click on the little icon that looks like the letter "I" in a blue circle, you will get this screen with NUC information in it (Figure 112):

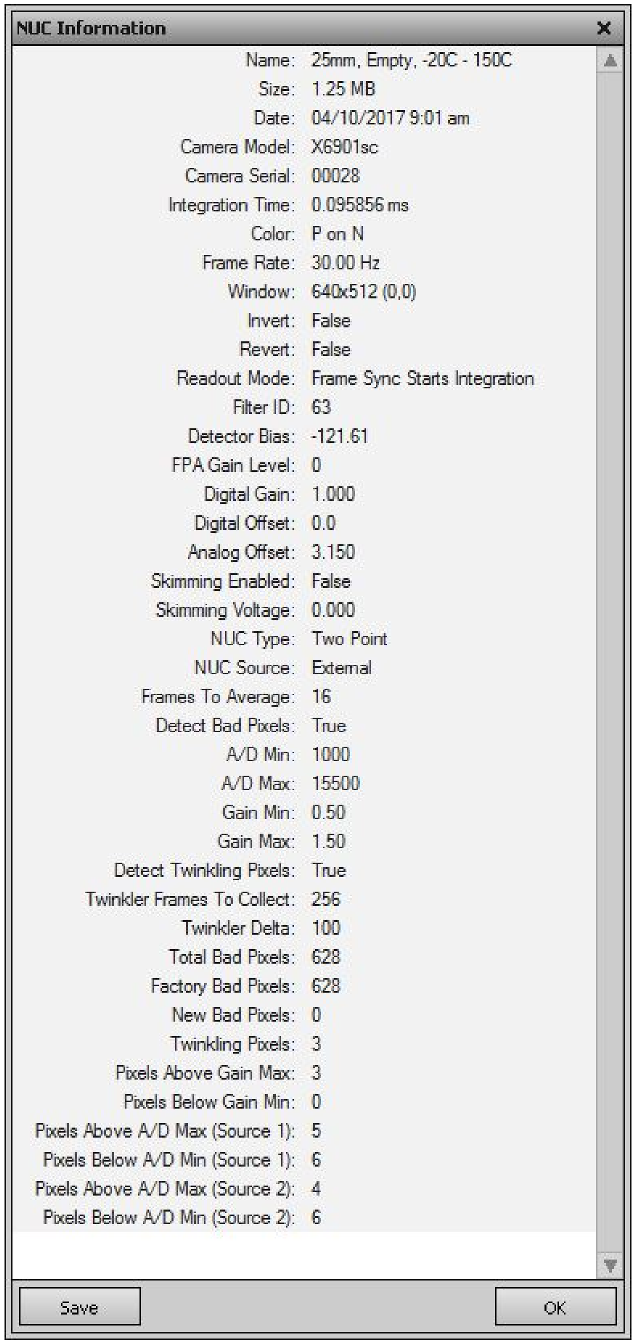


**Figure 112. NUC information**

The Save button in the lower left corner will save out the information to a text file. If you want to see what NUCs are available in the camera besides the ones that may be loaded into the presets, click on the Manage NUCs link shown in Figure 113:

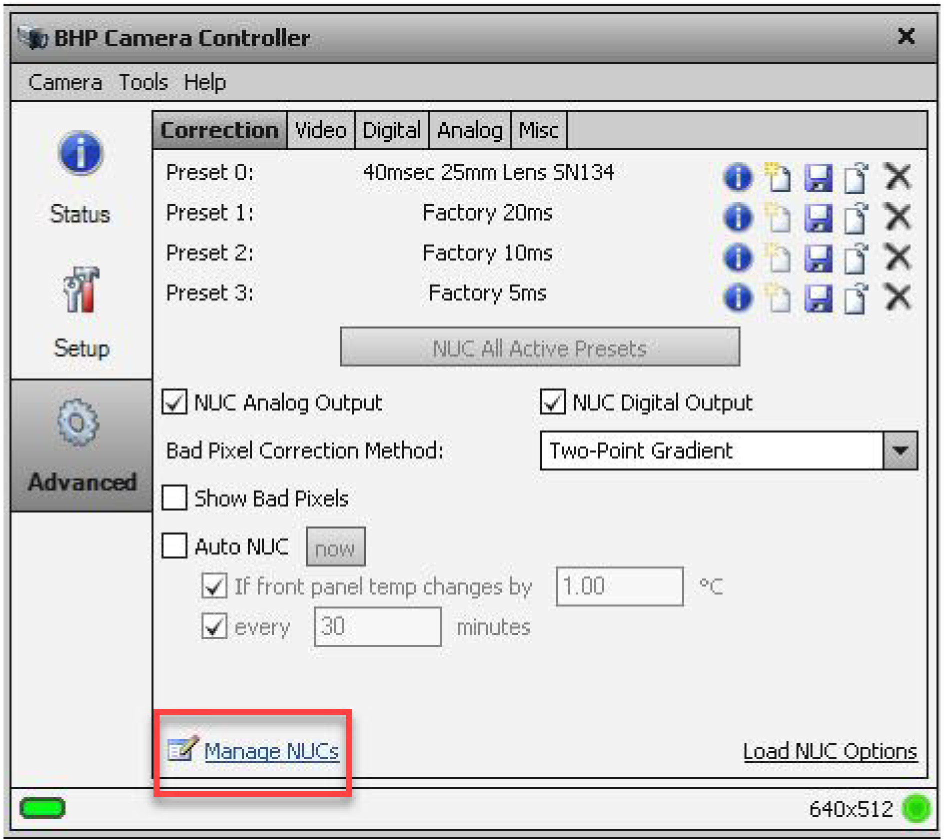


**Figure 113. Manage NUCs link on Correction tab of Advanced window**

When you click on this link, you get a list of all the NUCs that are in the camera memory. It also shows you much NUC memory is available to add more NUCs. The blue "i" button will show you the NUC information. You can delete, rename NUCs, upload more NUCs from the host PC, or download to the host PC using the icons on the bottom of the window, as shown in Figure 114.

**Downloading NUCS to the host PC and saving them in a backed-up folder is a very good idea!**

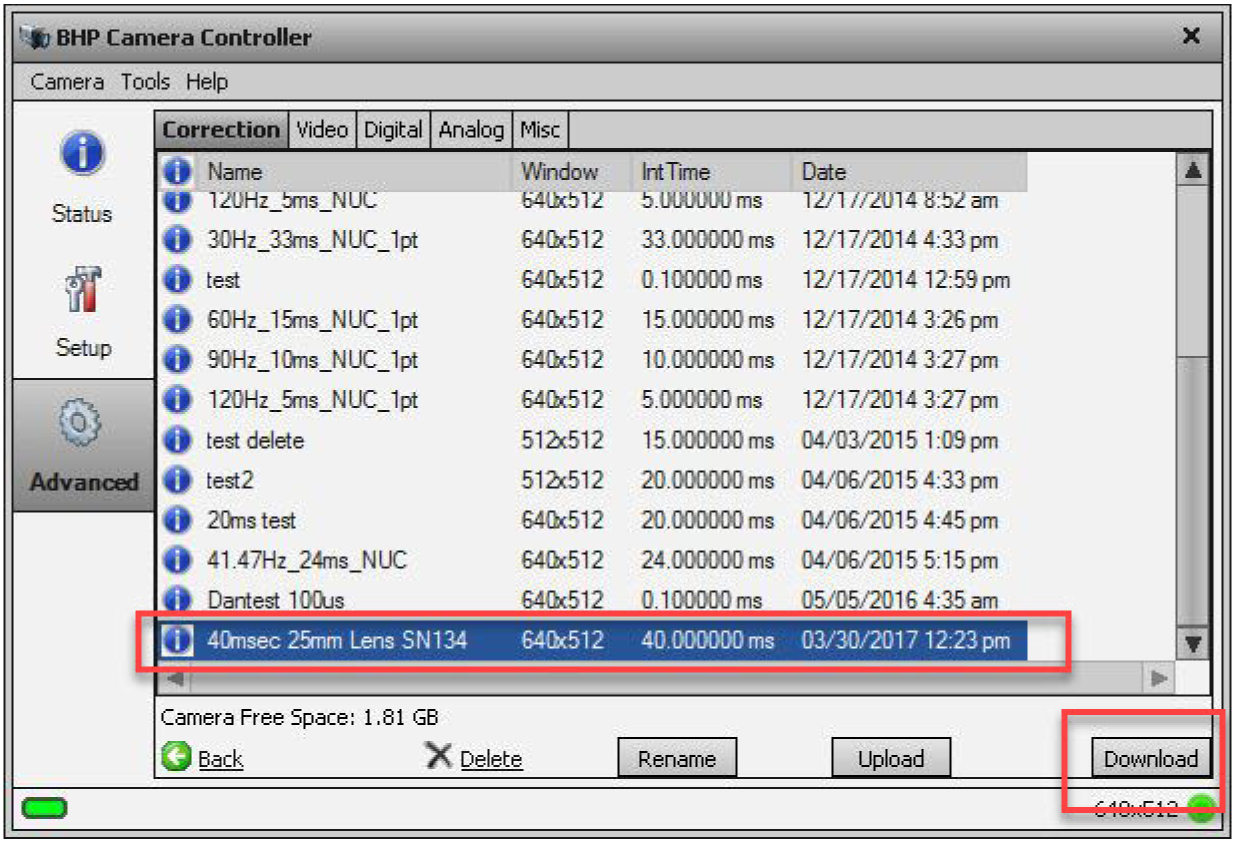


**Figure 114. NUC management window with a particular user created NUC highlighted**

When you hit the Download button, you get the window in Figure 115:

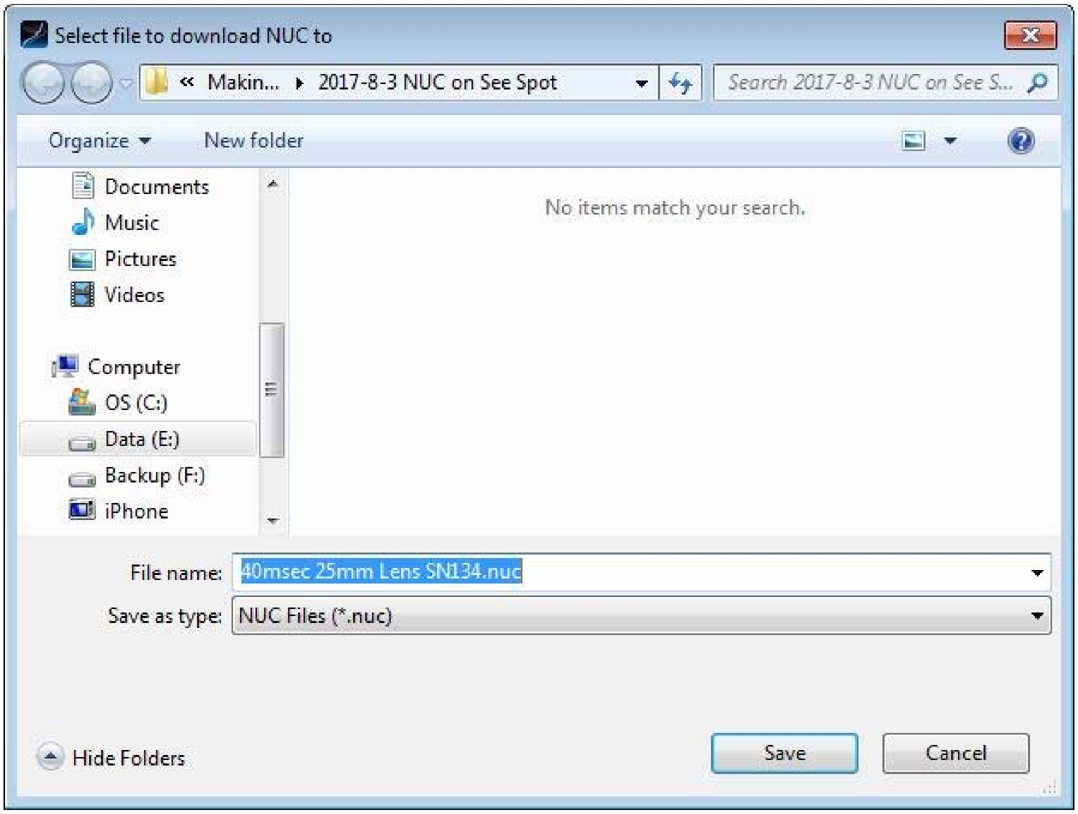


**Figure 115. Downloaded camera-side NUC file. The extension is *.nuc. It is a good idea to use the same name as the NUC name in the camera controller.**

The NUC file is about 1.3 MB for a 640x512 camera. Figure 116 is the file as it appears in its Window directory on the host PC:


**Figure 116. Camera-side NUC file in the host PC folder**

## Load NUC Options, Digital Gain and Offset

Typically, all of the camera configuration parameters are derived from the current Camera State. When the camera is powered up, it loads the last saved camera state. The names of the NUCs that were loaded when the state file was saved are stored as part of the state file. Normally the NUC is performed with the settings that are eventually going to be part of the state. If a NUC is loaded that has a setting that differs from the camera state, the state will override the NUC. If the user wants the NUC setting to override the state, then "Load NUC Options" can be set.

There is a link button called Load NUC Options which opens this Load NUC Options window, shown in Figure 117. This window controls what parameters in a NUC are used to set the camera settings when that NUC is loaded. The install default for the Load NUC Options is to load the table only, but that can be overridden and saved to the State file.

When a user loads either a factory NUC or a user-created in-camera NUC, some or all the camera settings used at the time the NUC was created can be forced into the camera, depending on the selections in this control window. If those parameters are the same as what is called for in the currently loaded State file, then nothing changes.


**Figure 117. Load NUC Options window**

It is a very good idea to have all the same settings applied so that the camera is in the same state as it was when the NUC was created. It is better to manage this via State files, rather than via the NUC itself.  Getting all the parameters back to what they were during the NUC creation is something that should be noted carefully for user NUCs.  I like to manage this by first creating a NUC or set of NUCs, then without changing any global camera parameters, saving a State file that ensures that all the settings on the camera are preserved and can be recalled when I need to use the NUC or set of NUCS in the future.

# Chapter 3 : Radiometric User Calibration

Introduction ........ 3—174
Displaying Radiometric Data in ResearchIR ........ 3—175
Measuring Radiance........ 3—177
FLIR Factory Radiometric Calibrations ........ 3—180
Measuring Temperature ........ 3—181
Creating your own radiometric user calibration........ 3—183
Things to consider before you try a user calibration ........ 3—187
Example User Calibration Procedure on SC6700 See Spot Camera........ 3—187
Troubleshooting a user calibration ........ 3—214

## Introduction

A user calibration is a pair of data files that instruct ResearchIR (hereafter known as RIR for the sake of brevity) on how to convert digital counts to power radiance units. The calibration files exist on the host PC that is running RIR, not in the camera itself. Once the conversion is made to radiance units, a further conversion can take place in the RIR software to display apparent temperature units. A very common question from customers: “How can I convert from counts to temperature?” The answer is that there is no universally applicable direct conversion from counts into temperature, except in a very limited set of use cases where the experimental setup and the object of interest itself can be used to generate the radiometric calibration by controlling its temperature precisely. In other words, you use the object you want to measure in your experiment as the calibration source. The emissivity of the target is then “baked into” the calibration. For most scenarios, the camera system’s measured radiance from a target needs to be “corrected back” to the target, where it is further corrected for reflected radiance off the target surface and for the emissivity of the target itself. Once those corrections are made, the corrected radiance is converted to a temperature using a lookup table that is embedded in one of the two calibration data files.

Most customers opt to purchase factory calibrations for their cameras, rather than create their own user calibrations. There are many advantages to this approach, since the factory calibrations are stored in the camera and are associated with high-quality factory NUCs which are also on board the camera. Factory calibrations come with Tdrift correction, which is not available with user calibrations.

As mentioned above, there are two data files that are needed for ResearchIR to convert the digital counts for each pixel from a movie or still image into radiometric units. One is called an Include file with the extension *.inc. The other file, called a calibration file or cal file, has to have the same name as the include file, and has the extension *.cal. These two files must be carried along together to apply calibrations to SFMOV files so that the user can measure an image in apparent temperature units or in-band radiance units. By this I mean that the files must be in the SAME directory as the SFMOV files for RIR to make proper use of them.

These two files are generated by the user using a calibration wizard that one can launch from within RIR but only when a camera is connected, or CalibratIR, a standalone utility that is installed alongside RIR. The include file can be read in a standard text editor (like Notepad) and is full of “SAF tags” or parameters used in the SAF (Standard Archive Format) metadata and datafile management protocols developed by the Air Force Research Laboratory. The calibration file can only be read by CalibratIR or RIR, not a text editor. We will take a close look at the contents of an include file later on.

## Displaying Radiometric Data in ResearchIR

Let's look at a laboratory blackbody that is set to 45 °C with a midwave camera. The camera I used for this experiment is a FLIR SC6700 See-Spot infrared camera that has been calibrated with a user calibration with a scene temperature range that goes from 5 °C to 100 °C. The integration time is 10ms. In Figure 1 below, RIR is computing image statistics for a region of interest (ROI for short) that is subtended by the 45 °C heated section of the blackbody. These are displayed in the Stats window, which can be docked into the main window or can roam free undocked. Here the Stats window is docked to the right of the image window:

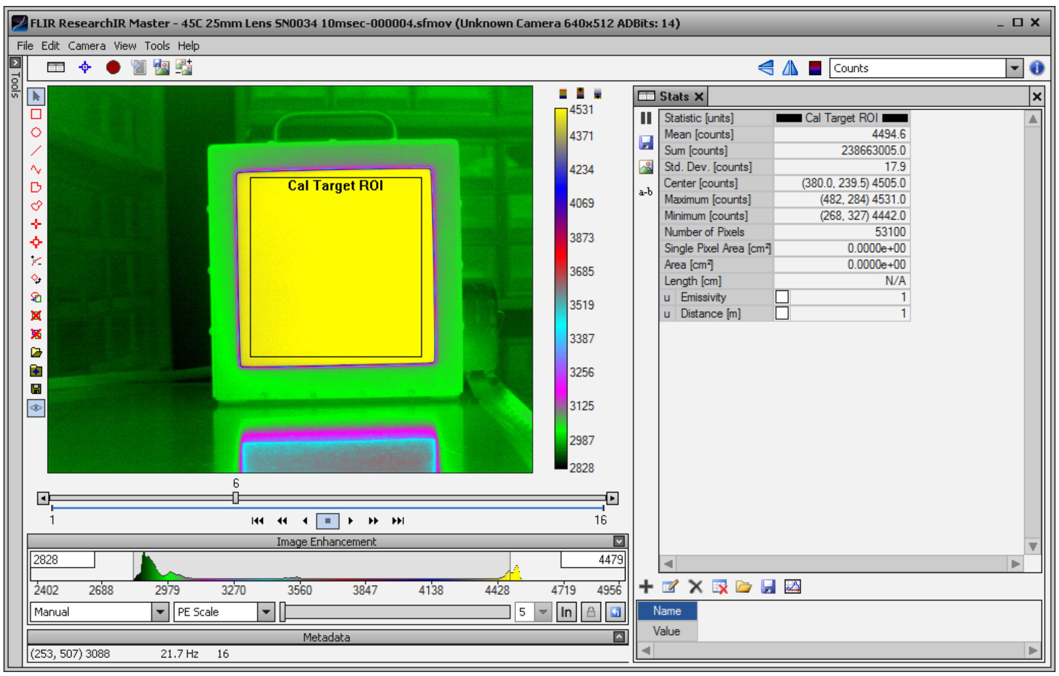


**Figure 1. Digital counts image in the 4.5-4.7μm spectral range of a 45 °C area blackbody**

We get a mean value of 4494.6 digital counts in the ROI. The useful full-scale count range for this camera is about 14,500 counts, so this 45 °C target is generating a signal which is about 1/3 of the full-scale range of the camera at this integration time (10ms). The digitizer is a 14-bit unit, so the counts can actually range up to 16,383, which is $2^{14}$-1. The non-uniformity correction applied to the raw data has the effect of reducing the counts in the center of the image and increasing the counts in the corners of the image. This is because of the gain map that is applied to compensate for the coldshield shading effect. All of this is described in detail in Chapter 2.

**Note: It is highly recommended that radiometric calibrations only be used for measurements on image pixels that are less than 14,000 digital counts in NUCed images.**

This value of 4494.6 counts is not very interesting to the engineer – these image data need to be converted into engineering units to tie the measurements made from the IR images to physically meaningful units of measure. The radiometric calibrations (factory or user) convert the counts into radiometric engineering units:  either in-band radiance units or apparent temperature units. Measuring radiance enables engineers to quantify IR emissions from a target, usually to predict performance of a sensor designed to search and/or track that target.  Measuring temperature enables engineers to study temperature-dependent material properties without physical contact. A nice example of non-contact thermography is Figure 2, an image of a sea lion's face with statistics in Fahrenheit degrees.  The eyes of these animals can get infected and will both get hotter than normal, or one eye can be hotter than the other one.

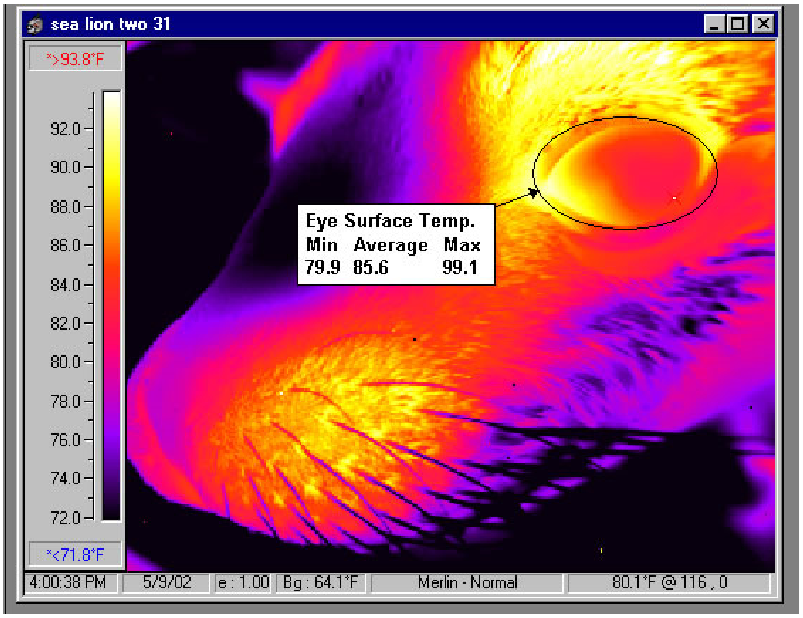


**Figure 2.  Sea lion face imaged with a radiometric camera**

Most customers that want radiometrically-calibrated cameras are just interested in measuring temperature.  A small minority of highly technical users (mostly at military test ranges) also care about measuring the in-band power radiance of targets and scenes.  To that end, these technical customers will be very specific about the spectral band of the camera they are using, because the radiance the camera measures is only in the waveband the camera is sensitive in – it is not the total radiance integrated over all wavelengths.  The purpose of measuring target radiance is to be able to determine the radiant intensity of the target in units of watts/steradian.  The radiant intensity (or just intensity as it is called in ResearchIR) is the area-integrated radiance from a target.  The radiant intensity is then used to determine how much infrared radiation emitted by a target in a particular waveband (and measured in watts units) is collected by a particular sensor, such as a missile seeker, for a given range to the target. The target power collected by the seeker optics is focused onto an infrared sensor, and radiometric models of the seeker sensor are used to predict signal-to-noise ratios, maximum detection range as a function of view angle on the target and other performance metrics.

Measuring the radiance of a target is a necessary step towards measuring its temperature with an infrared camera. Temperature is to some extent "waveband-agnostic". If one uses a calibrated midwave IR camera, measures the temperature of an object and corrects for emissivity and reflected energy, the apparent temperature ***should*** be the same as the measurement one would get with a longwave IR camera that has its own calibration. In other words, **a temperature measurement should be independent of waveband if it is done right**. However, this is predicated on all the associated correction parameters being set correctly in the software. It also depends on the nature of the object and how closely it approximates a Planck radiator, that is, an object that emits continuum radiation in a manner that can be described by the Planck equation. A gas cloud does not usually act like a blackbody emitter and measuring the temperature of an object with no defined surface is problematic anyways.

Even if one does not know the emissivity of the object, it is often very useful to talk about the apparent temperature of an object as though it were a blackbody emitter. For instance, an opaque and very hot gas cloud generated by a sooty, dirty explosion of fuel can behave very much like a blackbody and emit IR radiation from incandescently glowing carbon particles. Suppose this cloud has an apparent temperature of 1500 °C, according to a radiometrically calibrated midwave (3-5µm) camera. Using this apparent temperature value is a much more intuitive way to describe the cloud's IR appearance to an engineer, compared to saying that it has an in-band radiance in the 3-5 µm band of 3.75 Watts/cm$^2$/sr (which is the 3-5µm radiance of a blackbody at 1500 °C). Even if the object is not emitting a continuum of radiation like a blackbody does, it can still be a very useful concept to talk about its apparent temperature. A jet exhaust plume is a selective radiator, and will emit strongly in certain IR bands, and only weakly in other bands. As long as we define what waveband the camera is operating in, we can use an apparent temperature value to describe the scale of the emitted radiation from that plume.

**Note: One can always measure radiance from the direction of a target, but it is not always possible to derive a useful or meaningful temperature value from that radiance.**

## Measuring Radiance

The radiometric calibration (factory or user) initially converts digital count values into in-band radiance in units of Watts/cm$^2$/steradian. Figure 3a is the same image of the 45 °C blackbody from Figure 1, after it has had both a user radiometric calibration and a **spatial calibration** applied to the image data in post processing:

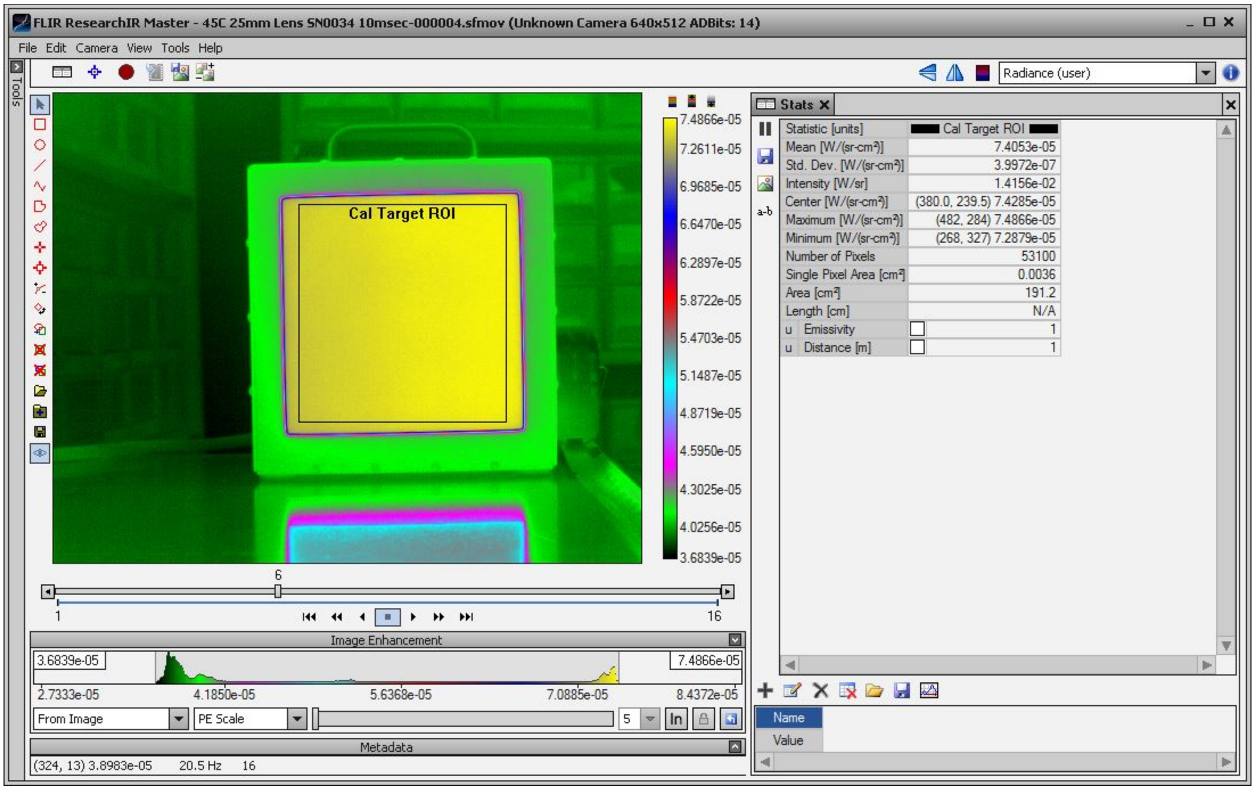


**Figure 3a. In-band radiance image in the 4.5-4.7 microns of a 45C area blackbody**

The mean in-band radiance in the region of interest is 7.4053e-5 Watts/cm$^2$/sr, as shown in Figure 3b. This is the energy in a narrow band of the MWIR spectrum (since this See-Spot camera has a special cold filter on it which passes midwave IR in the 4.5-4.7µm band, as well as a SWIR component which I blocked by using a germanium lens).

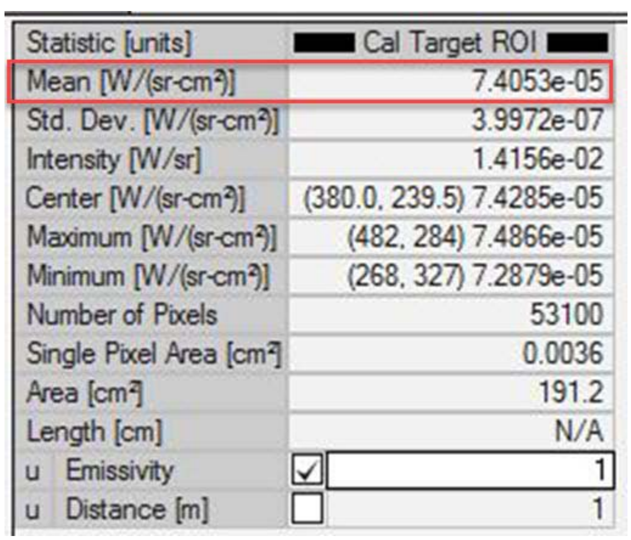

| Statistic [units] | Cal Target ROI |
|---|---|
| Mean [W/(sr-cm²)] | 7.4053e-05 |
| Std. Dev. [W/(sr-cm²)] | 3.9972e-07 |
| Intensity [W/sr] | 1.4156e-02 |
| Center [W/(sr-cm²)] | (380.0, 239.5) 7.4285e-05 |
| Maximum [W/(sr-cm²)] | (482, 284) 7.4866e-05 |
| Minimum [W/(sr-cm²)] | (268, 327) 7.2879e-05 |
| Number of Pixels | 53100 |
| Single Pixel Area [cm²] | 0.0036 |
| Area [cm²] | 191.2 |
| Length [cm] | N/A |
| u Emissivity | ☑ 1 |
| u Distance [m] | ☐ 1 |

**Figure 3b. In-band mean radiance for Cal Target ROI**

Note that RIR is also calculating the intensity, which is an abbreviation of in-band radiant intensity. The intensity, measured in watts/steradian is the in-band radiance of each pixel in the ROI integrated over the emitting area of the target, as shown in Figure 3c below.

| Statistic [units] | Cal Target ROI |
|---|---|
| Mean [W/(sr-cm$^2$)] | 7.4053e-05 |
| Std. Dev. [W/(sr-cm$^2$)] | 3.9972e-07 |
| Intensity [W/sr] | 1.4156e-02 |
| Center [W/(sr-cm$^2$)] | (380.0, 239.5) 7.4285e-05 |
| Maximum [W/(sr-cm$^2$)] | (482, 284) 7.4866e-05 |
| Minimum [W/(sr-cm$^2$)] | (268, 327) 7.2879e-05 |
| Number of Pixels | 53100 |
| Single Pixel Area [cm$^2$] | 0.0036 |
| Area [cm$^2$] | 191.2 |
| Length [cm] | N/A |
| u Emissivity | ✓ 1 |
| u Distance [m] | ☐ 1 |

**Figure 3c. In-band radiant intensity for Cal Target ROI**

There are “tags” in the *.inc file that define the **instantaneous field of view** (IFOV) of a pixel in angular units, in both the horizontal and vertical angular axes. The size of the IFOV is determined by the focal length of the lens and the pixel pitch of the focal plane array (FPA). In this particular camera/lens combination, the IFOVs were 600 microradians, or 0.0006 radians. By specifying the distance (which was set to exactly 1 meter on the lab bench), we have “told” RIR that each pixel has a **ground sample distance** of 0.06cm. This translates to a **footprint** or pixel target area of 0.0036cm$^2$. The ROI has an area of 52,100 pixels, which if multiplied times 0.0036cm$^2$, gives 191.2cm$^2$, as shown in Figure 2d. If this area is then multiplied by the mean radiance of 7.4053e-5 Watts/ cm$^2$/sr, the square centimeter units cancel out and the Intensity is 1.4156e-2 Watts/sr, as shown in Figure 3c. In fact, the intensity of an entire target is most often what range customers want to measure, not the radiance.

| Statistic [units] | Cal Target ROI |
|---|---|
| Mean [W/(sr-cm²)] | 7.4053e-05 |
| Std. Dev. [W/(sr-cm²)] | 3.9972e-07 |
| Intensity [W/sr] | 1.4156e-02 |
| Center [W/(sr-cm²)] | (380.0, 239.5) 7.4285e-05 |
| Maximum [W/(sr-cm²)] | (482, 284) 7.4866e-05 |
| Minimum [W/(sr-cm²)] | (268, 327) 7.2879e-05 |
| Number of Pixels | 53100 |
| Single Pixel Area [cm²] | 0.0036 |
| Area [cm²] | 191.2 |
| Length [cm] | N/A |
| u Emissivity | ☑ 1 |
| u Distance [m] | ☐ 1 |

**Figure 3d. Target area subtended by Cal Target ROI at 1-meter range**

In-band radiance has a linear relationship to digital counts: the radiance is a gain coefficient (known as C1) multiplied by the digital counts, plus a radiance offset (known as C0). Because of this linear relationship between digital counts and in-band radiance, we can easily scale the radiance up or down with additional system responses, such as the in-band atmospheric transmission of the air path between the target and the camera, or the transmission of an infrared window that acts as a viewport into a chamber. An example of this would be that sooty explosion cloud mentioned earlier. If the air path between the camera and the cloud was determined to have a mean in-band transmission of 0.8, then the apparent radiance measured by the camera would be only 80% of the source radiance. One can then scale up the apparent radiance by dividing by 0.8 to get a measurement of the source radiance.[9]

We can also scale the radiance for the emissivity of a target surface, and then we can compensate for reflected radiance off the surface by adjusting the radiance offset. These last two corrections are two critical steps in determining the apparent temperature of a target.

## FLIR Factory Radiometric Calibrations

The FLIR factory calibrations go one step further and use temperature sensors in the lens interface and camera housing to correct for ambient-temperature-induced changes in stray light emitted by the lens and camera housing interior. This is called **$T_{Drift}$ Correction**, and it is specific to factory calibrations. We do not expose how this works to customers who would like to do this for their own user calibrations, and in fact there are no provisions for live $T_{Drift}$

[9] This assumes the air path radiance or self-radiation is negligible compared to the radiance of the target, which is a very good approximation in the case of an explosion cloud that has an apparent temperature of 1500C. Even though the air path is emitting like it has an emissivity of 20%, it is probably only 25-35C air, so it has a very small radiance compared to the explosion.

correction on a user-calibrated camera in RIR. The Tdrift correction was implemented because there is always some so-called **self-radiation** or **parasitic radiation** generated in a thermal IR camera system. Lenses are not perfectly transmitting, and the lens optical materials will always emit a small amount of IR radiation. There can also be parasitic radiation that originates from inside the camera body, reflects off the back surfaces of the lens elements and onto the sensor. A FLIR camera with a factory calibration is calibrated by being pointed at calibration sources in a laboratory with a controlled air temperature of 22 °C. If the camera is operated outdoors on a hot sunny day, the lens and camera will heat up from solar loading and the pixels on the FPA will receive an excess of radiation compared to what they would receive looking at the same target in an air-conditioned calibration lab. Everything in the image will appear slightly hotter because of this effect. If we operated the camera outdoors on a very cold night, we find that the FPA receives less radiation than it would during factory calibration, and we would tend to measure targets as being a little colder than they really are. The Tdrift correction dynamically compensates for changes in the lens and camera interior temperature.

For high temperature calibrations, the self-radiation effect becomes insignificant. Most FLIR science camera users that do user calibrations are looking at hot targets (>100 °C) and are therefore operating in a regime where small changes in optics self-radiation are negligible and can be reliably ignored. In this regime, the Tdrift correction is much less important and a user calibration should give comparable performance to a factory calibration over a wide variation of ambient operating temperature.

## Measuring Temperature

There is a third way to display the 45 °C target image: in apparent temperature units. Apparent temperature means the temperature that the IR camera "thinks" the target is at – this is not always well correlated with the kinetic surface temperature of the target. The kinetic surface temperature is the temperature of the material if we put a contact temperature probe against it – it is the lattice temperature or the true temperature, as compared to the radiation temperature. An object like the laboratory calibration source imaged below has a copper plate that is painted with a black paint that makes the source a very good approximation to a theoretical blackbody, as they are described by the Planck equation. In Figure 4a below, the mean ROI apparent temperature is a very good approximation of the actual kinetic surface temperature, or KST for short, which should be 45 °C, since the emitter should be controlled to a KST of 45 °C by the calibration source controller. The error is only 0.6 C, which is a pretty typical error value for a laboratory benchtop user calibration. Note the calibration source surface emissivity value is set to 0.96 in the Stats window of RIR.[10] This emissivity value has a strong effect on the apparent temperature value, and choosing its value is key to getting apparent temperature measurements that correlate

[10] This emissivity value was specified by the manufacturer as the nominal emissivity of the source, though they (Infrared Systems Development Corp,) do not specify the waveband. We can assume it is a good value for the 3-5μm and 8-13μm bands.

well with the KST, which is the "true" temperature.   In contrast, a polished metal surface will generally always have an apparent temperature that is drastically different from its KST.  The shiny surface is very inefficient at emitting IR radiation.  The camera will see the temperature of whatever is reflected in the surface, NOT the surface itself. For an object like the calibration source, there is a 4% reflectance to it, which means that objects that are in the vicinity of the surface can illuminate it and contribute radiance to the camera FPA.  A common case of this is sunlight reflecting off an object surface, thereby making it look hotter than it actually is.  The temperature values default to Celsius units.[11]

The agreement between the indicated temperature on the blackbody controller and the mean temperature in the ROI with e = 96% or 0.96 is very good, as shown in Figure 4a.  If the user does not know the emissivity of the target, they can use an emissivity value of 0.92 as a pretty good approximation for many materials, with the exception of man-made shiny metal or glass surfaces which have much lower emissivity values.  That value of 0.92 is the default emissivity that is set for factory calibrated cameras in the Object Parameters Tools field.

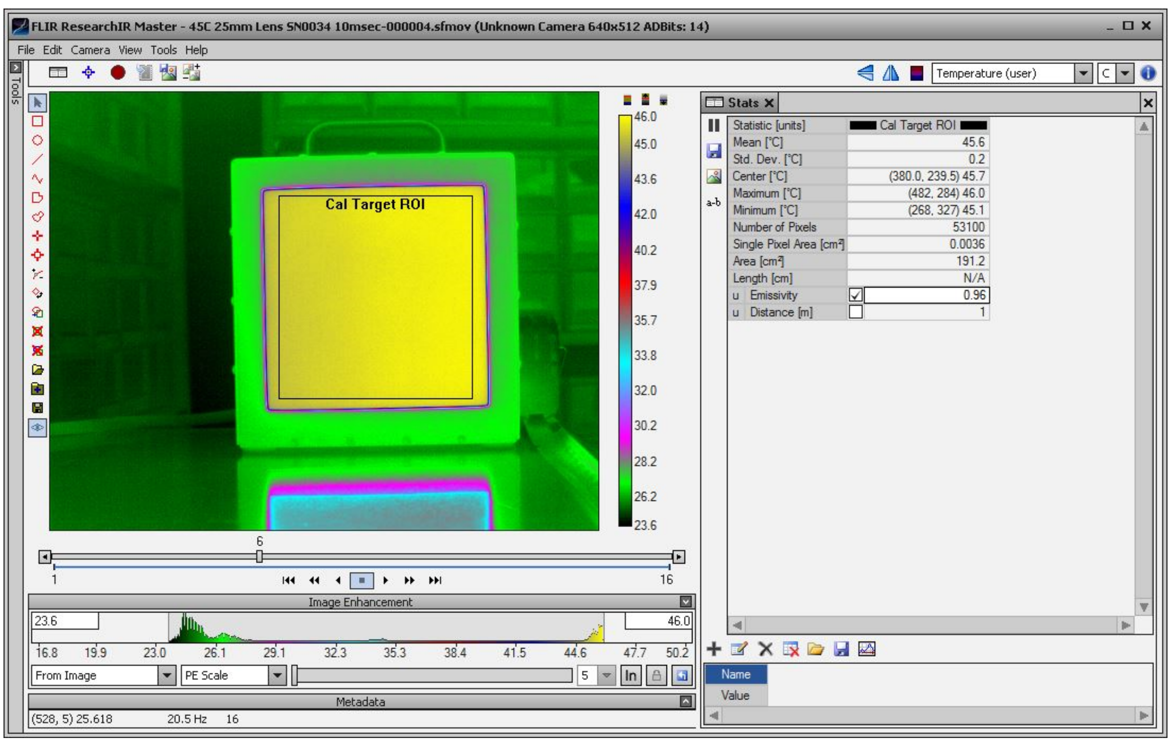


**Figure 4a.  A 45.6C measurement of a blackbody source with an ROI emissivity set to 0.96 to match the manufacturer-specified emissivity. The source temperature was set to 45C, so the measurement accuracy is very good.**

---

[11] We use Celsius as the default for almost everything we do in ResearchIR.  Some science camera customers choose to use the other available temperature units, which include Kelvin, Fahrenheit and Rankin.

If I change the value of the Emissivity field in the Stats window to 0.9, the ROI measures a different temperature value of 47.8 °C, as shown in Figure 4b. For objects at KSTs above the ambient reflected background temperature (as is the case here), the lower the emissivity setting, the higher the measurement of the target temperature. It is always a good idea to determine the emissivity of the target if possible, and if the target emissivity is low, say below 0.9, then you should do what you can to increase it, which might be as simple as painting on some flat black paint or affixing a piece of high emissivity tape to the surface.[12]

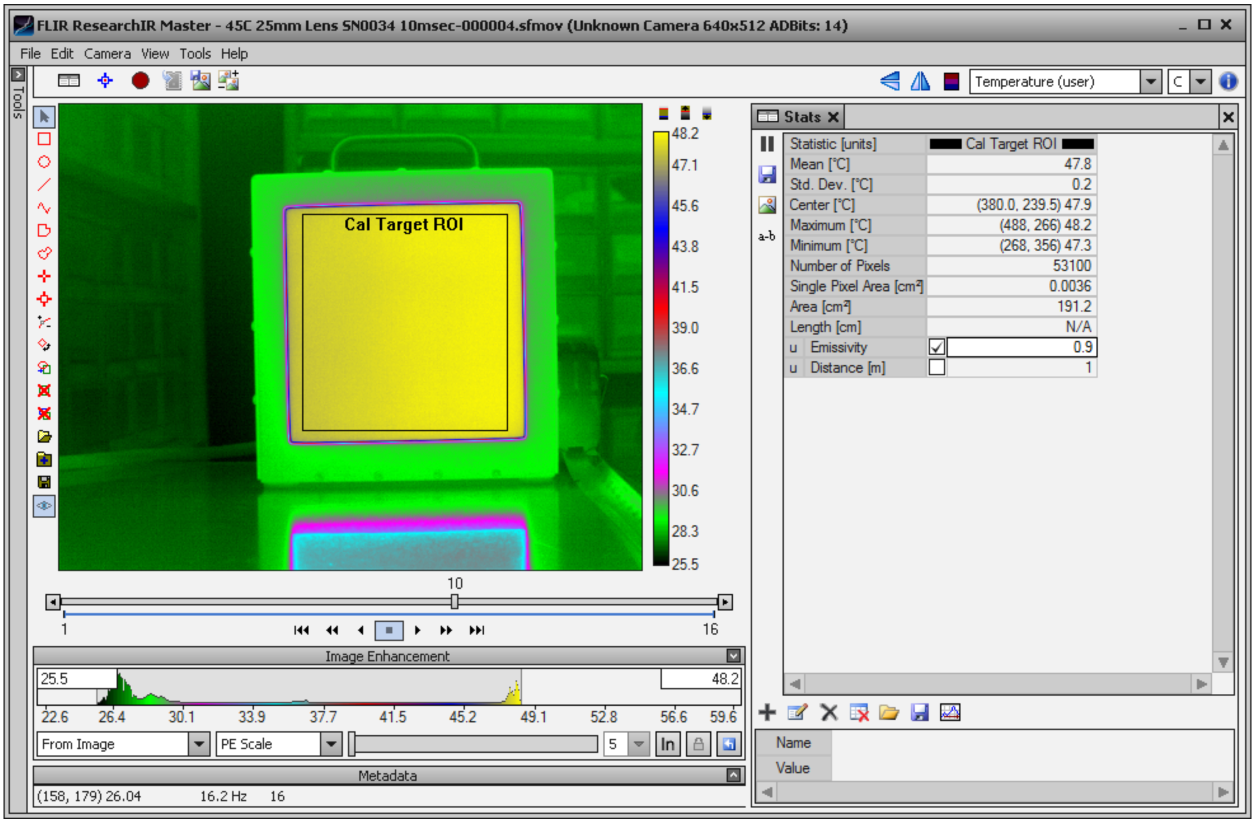


**Figure 4b. The 45C blackbody source shown in Celsius temperature units with an ROI emissivity incorrectly set to 0.9. The temperature measures higher than the known KST of the blackbody by almost 3C.**

## Creating your own radiometric user calibration

ResearchIR allows its users to perform radiometric calibrations on cameras. This is called a user calibration and it produces two files (*.cal and *.inc) that instruct RIR on how to convert digital counts into radiance and temperature. A user calibration is created by taking image data on an infrared source set to at least two different known in-band radiance values. A linear fit to the data yields conversion coefficients (C1, a slope or gain, and C0, an intercept, or offset) that describe

[12] PVC electrical tape can be used in a pinch, but it melts at a pretty low temperature and the adhesive gets gooey and awful. There is a tape called HB-250 made by Optex in Japan that has much better temperature ratings. Here is the website for it: http://www.optex.co.jp/meas/english/potable/pt_80/spec.html

the radiometric transfer function that converts digital counts into in-band radiance. In practice, we recommend taking at least 4 or 5 data points at different ROI mean digital count values using various calibration source temperatures. The mean ROI counts should go up to about 14,000 counts, but not higher, since the center pixels of the camera will go into saturation. Having a number of calibration points improves the quality of the calibration, because errors affecting any given data point tend to get averaged out in the least-squares fit.

The calibration sources are typically commercial laboratory blackbodies that are either area or cavity designs. The area blackbodies are copper or aluminum plates that are heated to precisely known and spatially uniform temperatures. The plates have a rough surface finish and are coated with special black paint that has a high emissivity, so that the plate closely resembles a theoretical blackbody as described by the Planck equation. Area blackbodies like the one shown below in Figure 5a can have emissivity values above 95%.

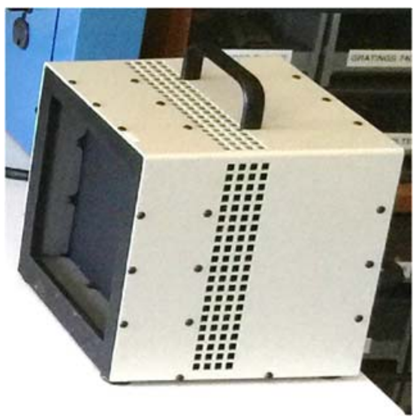

**Figure 5a. An area blackbody. This is an ISDC Model 2106 blackbody source with a 6-inch square emitter with a temperature range of 5 °C to 150 °C and nominal emissivity of 0.96.**

The cavity blackbody is usually a cast ceramic heating element that has a deep conical hole molded into it. The hole is the cavity, and it emits radiation in a manner that is very close to a theoretical blackbody. In fact, emissivities of cavity blackbodies can be higher than 99%, and they can reach temperatures up to 1000 °C for standard models and 3000 °C for special models. Figure 4b shows a trio of CI Systems SR-2 cavities that are being used to calibrate a camera at the FLIR factory in Niceville, Florida. At this facility, we use an "arc" of blackbodies to rapidly calibrate cameras that are mounted on a swing arm. The advantage of this method is that each blackbody is set to its setpoint temperature and allowed to stabilize for at least an hour. Using a single cavity blackbody and changing its setpoint temperature, then waiting for stabilization is very tedious. For a 5-point cavity calibration, it can take 5 hours for the whole process.

**Note: If you plan to do a lot of user calibrations, you should buy at least two area blackbodies for NUCing and two blackbodies as calibration sources.**

**Figure 5b. A set of cavity blackbodies being used for calibration of a FLIR camera**

The radiance values of the calibration sources are computed by ResearchIR using the Planck equation which is numerically integrated over the **spectral response curve** (spectrally-weighted bandpass) of the camera system. A typical camera has a 3-5µm bandpass, so those limits are often used in a user calibration. The spectral response curve of a typical 3-5µm InSb camera system is largely defined by the sensor spectral response, and the cold filter in the camera dewar, and it is a reasonable approximation to a spectral response curve with zero response outside of the band limits, and a response of 1 in band limits. This type of response is called a “top hat” response, because it looks like a men’s top hat, as shown in Figure 6.

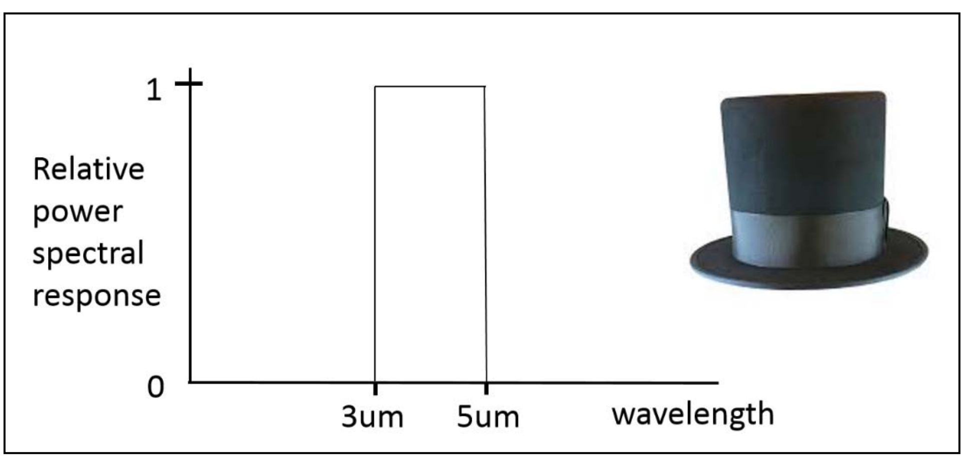

**Figure 6. “Top hat” spectral response function**

A Planck function for a 300 °C blackbody emitter is shown in Figure 7 below:

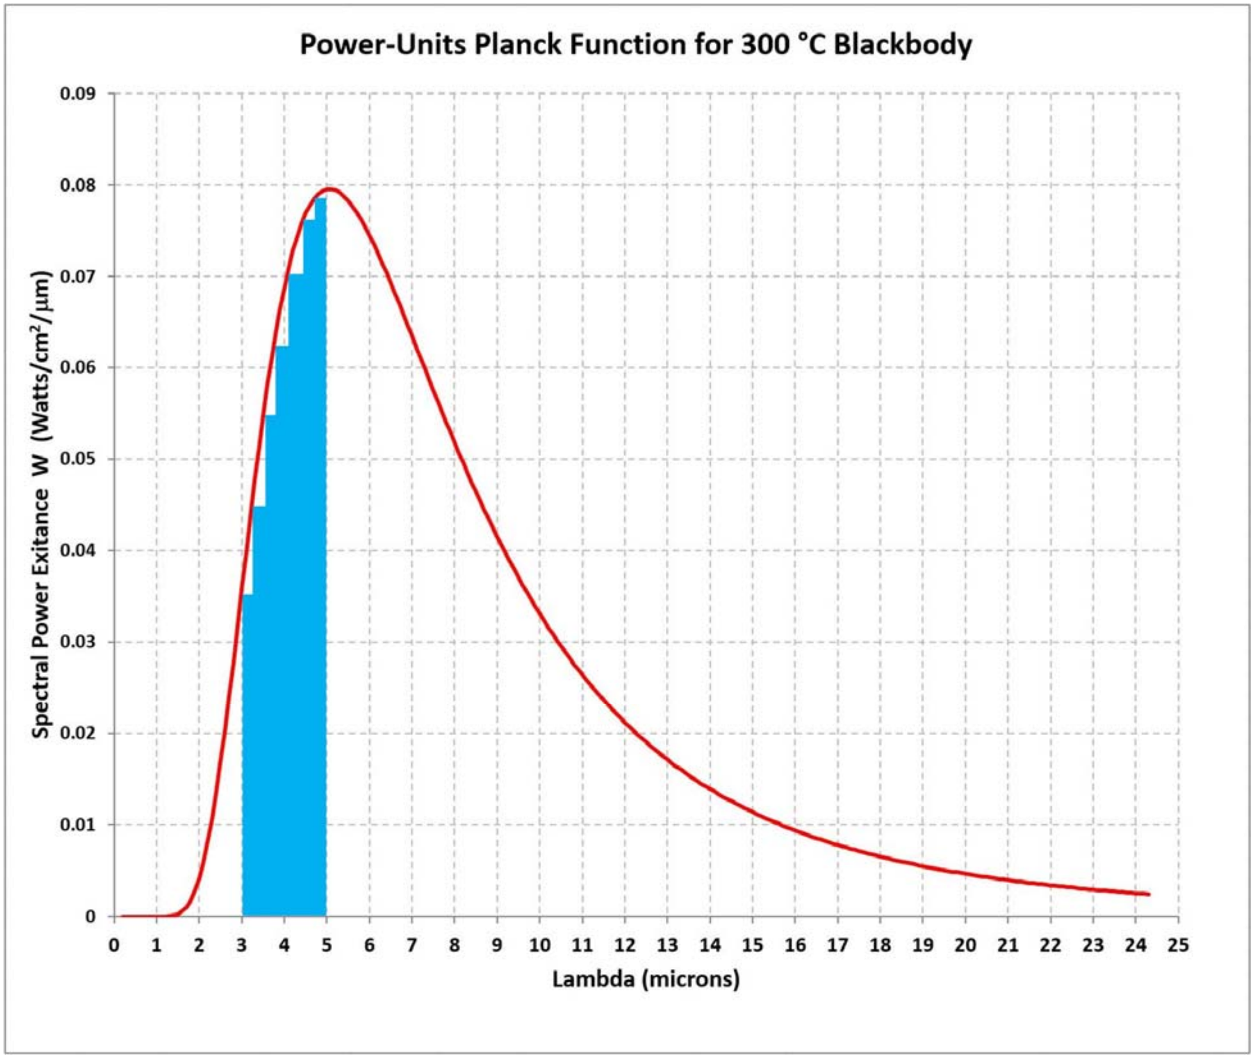


**Figure 7. Spectral exitance curve for 300 °C blackbody**

The vast majority of the user calibrations are done assuming the "top-hat" spectral response of the camera and lens system. The atmospheric transmission is usually set to 100% transmission, since the camera and the calibration source are usually located within a meter or less.[13] A meter of standard laboratory air has little water vapor in it, and does not absorb much infrared radiation in the midwave and longwave IR band.

[13] The exception being when a long focal length lens is used which has a minimum focus distance of several meters or more. In that case, a collimator can be used to project a blackbody target at infinity, and the reflectivity of the mirror or mirrors in the collimator is considered for the calculation of the radiance exiting the collimator after reflection losses.

## Things to consider before you try a user calibration

A. The procedure described below was done on a camera with a factory NUC, but no factory calibration. I used an existing factory NUC which had a 10msec integration time which I kept. It is possible to load an existing two-point NUC made at a particular integration time, and then change the integration time to a different value, as long as the value is not too different, say within a factor of 2 or 3. It is a good idea to then do an offset update with an external ambient calibration plate before acquiring the 16-frame movies.

B. I use the SFMOV file format for my movies, not the ATS format, because the SFMOV format allows the user to modify the calibration after the fact. The SFMOV format gives you three files per movie: the include, the pod and the SFMOV. The ATS format is a wrapper format that puts all three of these files into a single file which is not easily edited. With the SFMOV format, you can use the REdit utility which is included in ResearchIR to go back into an SFMOV and point it at a different include file with a different calibration at any time. This means that you can even create new calibration files after gathering data and apply these new calibrations to old data. I have often taken data in the field and then later recalibrated the data by doing a post-collection calibration and applying the new include file to the SFMOV.

C. You really need to have a blackbody source or sources that can achieve an apparent temperature that is as high as the apparent temperature of the hottest thing you want to measure. For example, if you want to measure jet exhaust plume radiant intensities, you will need a cavity blackbody with an adjustable setpoint temperature that can go up to 1000 °C. It is possible to extrapolate user calibrations to higher temperatures, but there is an increased uncertainty associated with doing that extrapolation. FLIR factory calibrations can be done up to 3000 °C, as we have cavity blackbody sources that can achieve that temperature.

## Example User Calibration Procedure on SC6700 See Spot Camera

I used a camera from our pool of sales demonstration assets to illustrate how a user calibration is done step by step. This special bandpass camera did not come with any factory radiometric calibrations, as it is not generally used for radiometry. The following is the general process I use to calibrate infrared camera with a user calibration to measure temperature, using this FLIR SC6700 See Spot as the test asset.

1) Choose the lens you want to use with the camera. Write down the lens part number and serial number so you can keep track of what lens was on the camera during the calibration. Some users have multiple cameras and the lenses can get mixed up.  While it is true that two examples of the same serial number of 25mm lens probably have very similar transmission values, the best practice for a user calibration is to use the same exact lens for both calibration and subsequent data collection. I always take a smartphone picture of the lens front with the serial number and the camera nameplate and put these pictures in the calibration file folder I create, as shown in Figure 8a:


**Figure 8a.  Capturing camera and lens information**

2) Before starting the calibration, you need to choose the integration time (and for some FPAs, the gain state) for the camera calibration you want to perform.  One way to do this is to get a calibration source that can reach (and ideally go beyond) the temperature of the hottest object you plan to examine.  You can then set the source to this temperature, point the camera at the source, and then adjust the settings of the camera FPA to capture an image of the source at or below the 14,000-count limit, to ensure that no pixels are saturated. If there is saturation, then may be necessary to add a neutral density filter to the optical path to further limit the exposure.  The goal is to have linear camera behavior up to that maximum source temperature.  In this example, I used an existing factory NUC at 10msec and just kept raising the target temperature during the calibration until I reached saturation in the center of the image.  Make sure you write down all the information about the ND filter that you use for the calibration.  Label the holder with a reference number so that you won’t mix up two examples of the filter.

3) When you are ready to calibrate the camera, set the camera on a bench facing the blackbody calibration target.  The distance between the two is generally chosen so that the blackbody fills about 25% or less of the camera’s field of view.  It is not advisable to flood illuminate the camera with the calibration target, because that will increase the levels of stray light and give incorrect readings when one uses the calibrated camera to

look at smaller targets. It is also advisable for the calibration target to subtend at least 50 by 50 pixels on the FPA face, because a really small apparent target size will not be properly measured either, due to diffraction and aberration effects in the optical system. Here is the setup for the user calibration described in this app note. The blackbody is 1 meter from the front of the lens, and this distance is recorded in case the calibration setup needs to be duplicated later. The lens is focused on the emitter surface. The blackbody is set to 35 °C as you can see from the controller in Figure 8b:

**Figure 8b. Camera and calibration blackbody source set at 35 °C**

4) Turn on the camera with the lens you want to use already installed and keep track of the time it takes for the cryocooler to throttle back. This is on the order of 6-7 minutes typically, but it can take longer. It is important to know your camera and how it changes over time, and that information is worth recording in a logbook.

   **Pro Tip: When the time to cool down gets very long, like 15 minutes, it may be time for cooler servicing.**

   While the camera is cooling, connect to it with ResearchIR, launch the camera control window and look at the Status tab, as shown in Figure 8c. You will see the setpoint temperature decreasing as the camera cools down. As an example, the SC6700 See Spot camera reached an FPA temperature of 68.69K, as shown below. This value will change a little over time but is generally very stable by design.

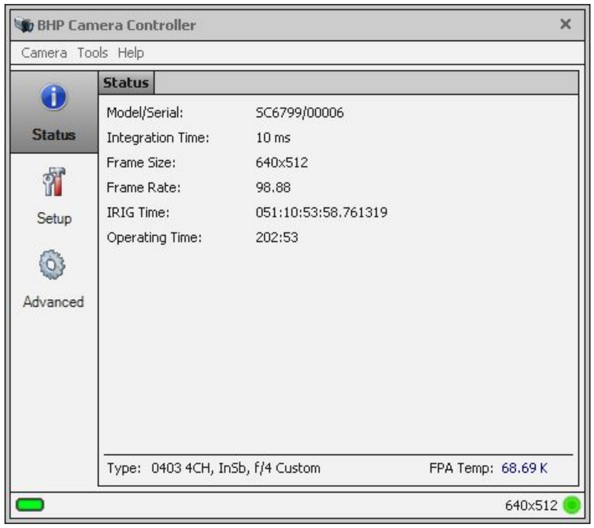


**Figure 8c. Status window in RIR**

5) Leave the camera on for at least an hour before you perform a user calibration. The lens should be installed during this entire warmup period so it can reach a steady state temperature! This stabilization period is particularly important if the calibration is being done with a long integration time with the intention of measuring ambient or cold targets. The lens and camera housing need to stabilize in temperature to reach a stable level of self-radiation, and this can take an hour or more.

6) The blackbody I used for this calibration is a 6-inch square emitter that can operate from 5 °C to 150 °C. It is made by Infrared System Development Corp. and the name plates are shown below in Figure 8d. I always use my smartphone to take pictures of the equipment nameplates and the setup and then I put those photos into a folder associated with the calibration process. These are just snapshots for informational purposes – they don't have to be great photos!


**Figure 8d. Blackbody controller**

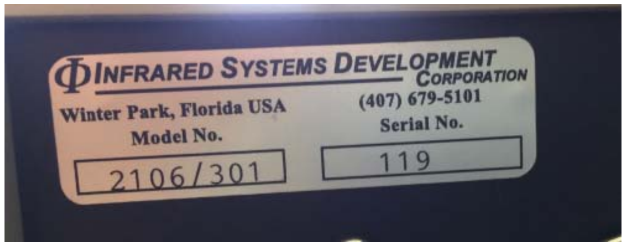


**Figure 8e. Blackbody emitter head nameplate**

7) The camera is still connected to ResearchIR, which is used to collect image data at each calibration target temperature. The particular camera I used for this example is an SC6700 InSb See-Spot camera. It has a bandpass which is not the standard 3-5µm for an InSb camera; instead it is 4.7-4.9µm with a small spike of response at 1.064µm so that it can also see a YAG laser beam. That spike did not affect the calibration I did. We will input the spectral response of the camera and lens combination later. The camera is pointed at the target, which is centered in the image. I always center the calibration target in the field of view, just as I always center the object I want to measure in the field of view. This is where the image quality for any camera is always the best: on axis.

8) Now we have to choose the integration time that will be suitable to produce the best results for the measurements that will be made with the camera. Suppose that you want to calibrate the camera to be able to measure temperature up to 100 °C. The way to determine the calibration settings needed is to set a blackbody to 100 °C, point the camera at it, draw an ROI that is fully subtended by the target, as shown below, and then adjust the integration time up or down until the mean digital counts in the ROI are about 14,000. This value is selected because a typical FLIR US science camera has a 14-bit digitizer in it that converts the FPA signals from each pixel into digital counts, and a 14-bit digitizer goes up to 16,383 counts. When a NUC is applied to the camera, the pixels in the center of the image have their count values decreased to about 14,500 by the NUC gain coefficients. If you adjust the integration time so that the NUCed image has ~14,000 counts on the target in the center of the image, you should be just below saturation. To best see the saturation effect, I turned off the NUC correction, and then I use the segmentation feature in ResearchIR to paint any pixels that go over 16380 counts in a red color. At 10msec integration time, with this lens, we are right at the edge of saturation for a 100 °C blackbody. The mean in the ROI is 15,603 counts but there are a handful of saturated pixels, as shown in Figure 8f below. As a result, the in-band radiance of a 100 °C target with a 0.96 emissivity will be the useful limit of this calibration.

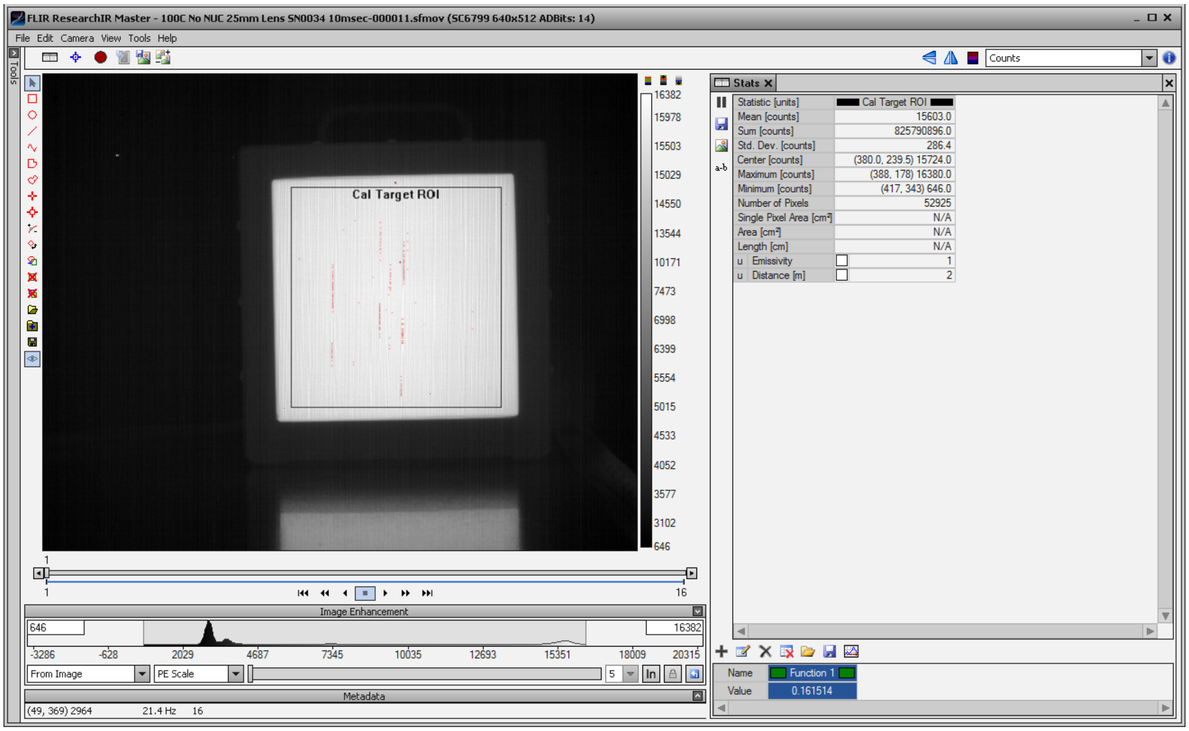


**Figure 8f. Blackbody with a few columns of pixels starting to saturate**

You can save the ROI as a *.sroi file so you can reuse it later. The "Save ROI" button is here on the left side of the image, along with the Load ROI button in Figure 8g:

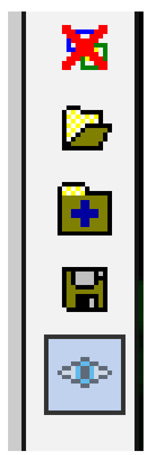

**Figure 8g. ROI tools in RIR**

9) The camera should have a good two-point correction on it, whether that is a factory NUC or a PC-side NUC. The NUC should ideally be done at the integration time you choose based on the desired range of the user cal. While you are choosing the integration time, moving it up or down to settle on a mean count value of 14,000 at the maximum temperature target you want to measure, you can use a NUC made at an integration time that is at least within a factor of 2 or 3 of the integration time you will settle on. There is

an iterative aspect to this process that makes it quite time consuming. We have a set of factory calibrations we do at the Niceville facility with a canned set of NUC source temperatures and approximate integration times.

10) Now you will need to set the controller temperature to an initial value and let it settle to that temperature. The settling time varies depending on what type of blackbody it is. The one used for this calibration is thermoelectrically cooled and can settle within a minute if the temperature change is not excessive. The first temperature point used here is 25 °C. The mean digital count value is 2940.4, which is at the low end of the digitizer range, as shown in Figure 8h below. The range does NOT go to zero for very cold scenes. Part of the reason for that is that there is a built-in voltage offset in the special "flex cable" connecting the FPA to the digitizer input in the camera head hardware. That offset is usually equivalent to around 500-1000 counts. This camera is set to 10msec integration time and has a factory NUC done at 10ms. This calibration will get us measurement capability right up to around 100 °C, while still staying in the linear region of the camera's FPA.

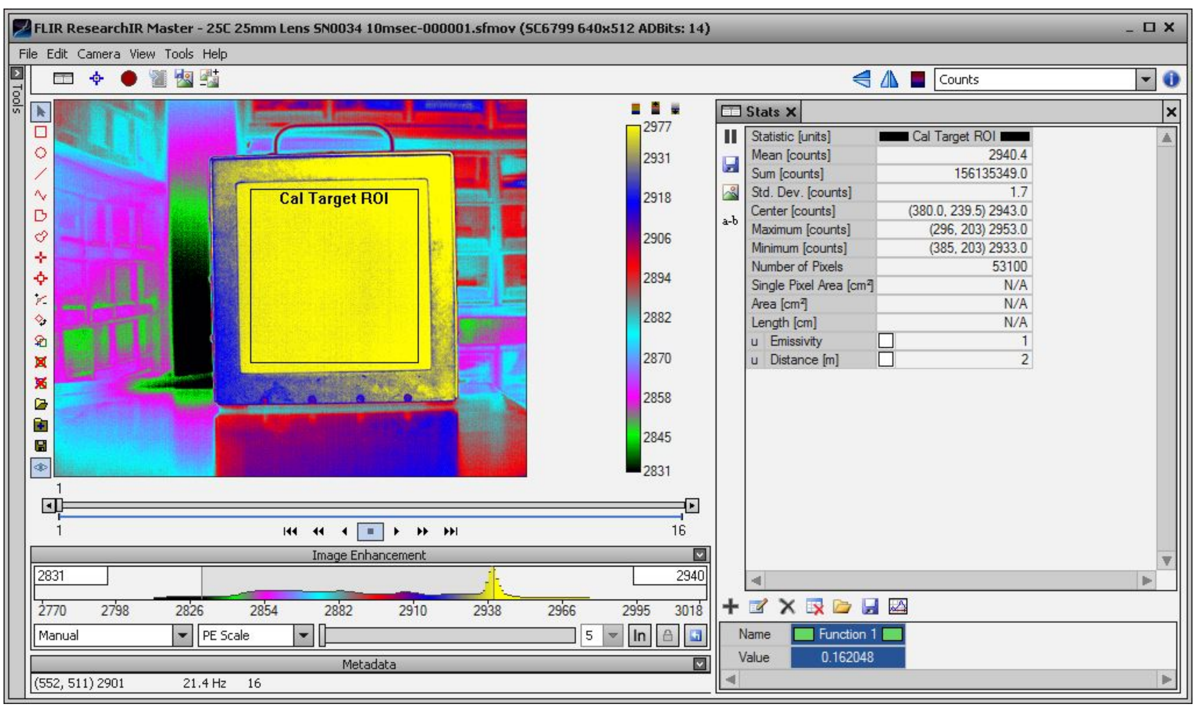

**Figure 8h. 25C blackbody imaged with RIR during calibration**

11) Once the blackbody has appeared to stabilize and the mean digital counts value of the ROI is not changing much, record 16 frames to memory. As mentioned earlier, it is a good idea to do an offset update with an external ambient calibration plate before acquiring the 16-frame movies. I use a movie that is 16 frames in length, because ResearchIR has a frame-averaging filter that can be set to 16 frames maximum. When

that filter is active and you play back the 16-frame movie through at least once, the Stats mean value will settle to a fixed value. Doing this reduces statistical temporal noise by the square root of N frames, or a factor of 4.

I give the files a filename with a sequential prefix that starts at 1 and steps up by 1. The filenames I use are descriptive of the experimental setup factors that are not embedded in the metadata of the movie. The filename for the data file that generated the above still image is “25C 25mm Lens SN0034 10msec-000001.sfmov”. Figure 8i shows the setup I used for that recording:

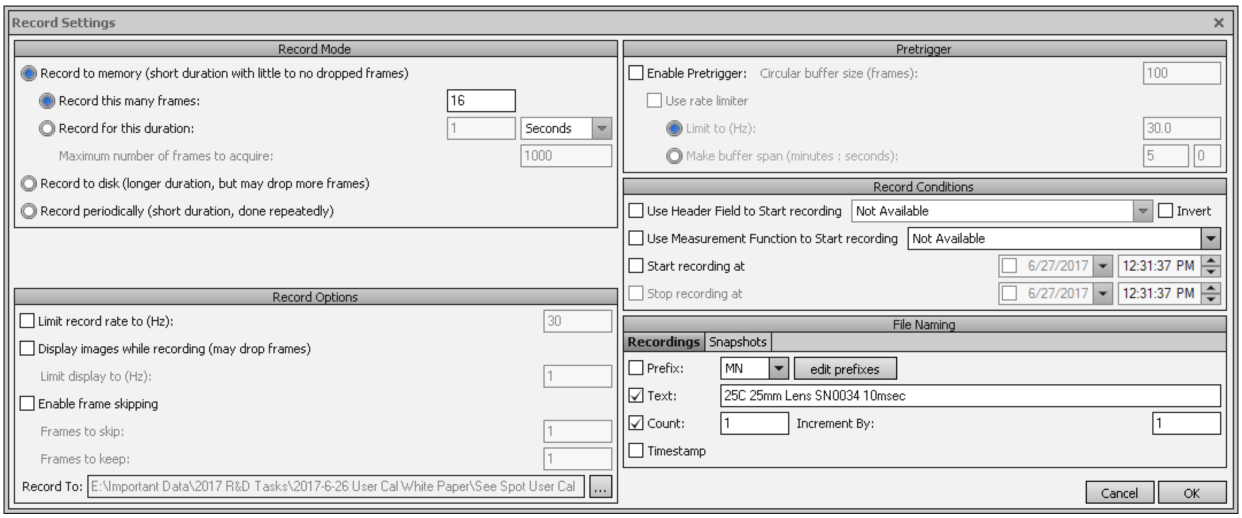


**Figure 8i. Record Settings Setup for recording**

12) The next movie will be made after the blackbody is set to 35 °C and has settled there, so I will change the text to say “35C 25mm Lens SN0034 10msec”. The next suffix will be 2. There are 6 digits to the numeric suffix, so the whole file name will be “35C 25mm Lens SN0034 10msec-000002.sfmov.” Figure 8j shows a screen shot of the RIR window with statistics for that movie. The ROI that I have drawn is the same as the one used for the 25 °C measurement. If I do not move the setup, I will be looking at the same physical portions of the blackbody for all the calibration points I acquire.

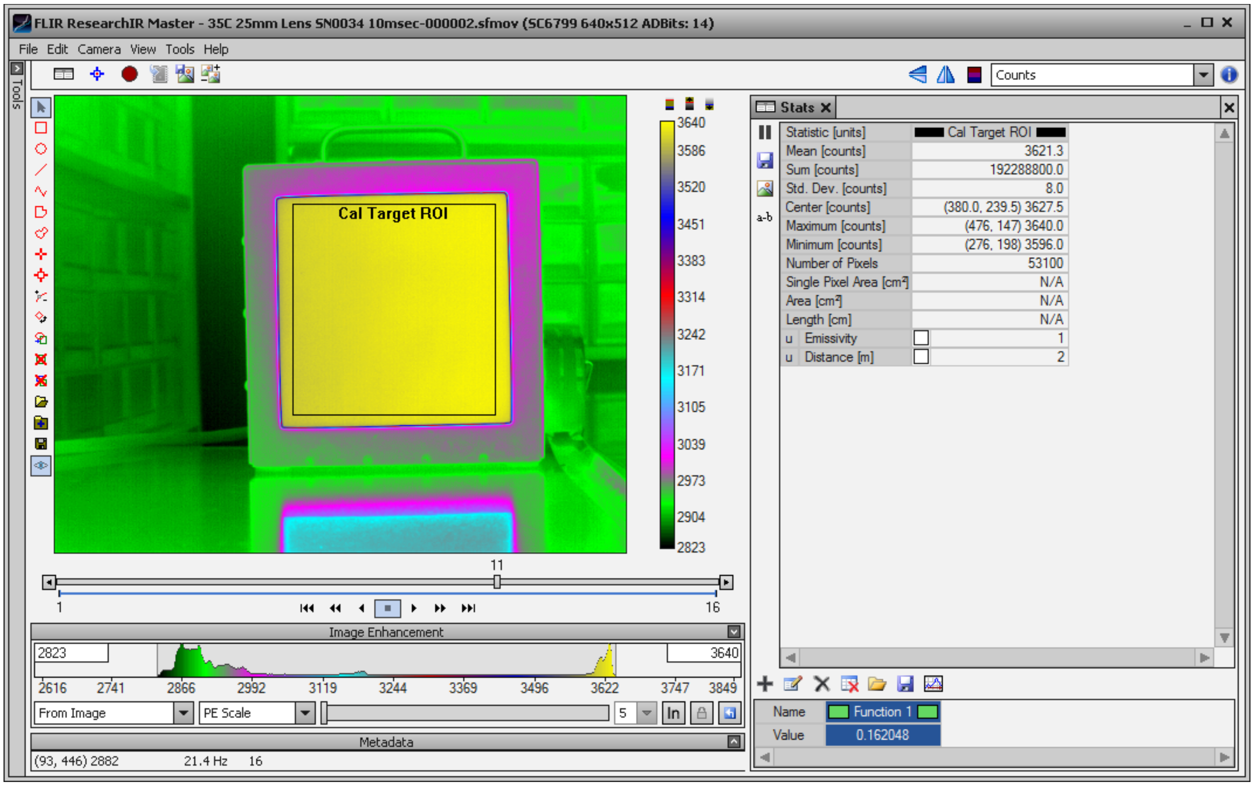


**Figure 8j. A 16-frame movie of a 35 °C blackbody**

13) I repeated this for a set of calibration target temperatures in 10 °C steps up to 105 °C. When the blackbody was at 100 °C, the counts were so close to saturation that I made the next step 5 °C, and so the high temperature target setting was 105 °C. The camera is clearly saturated and the NUC is not working any more, as shown in Figure 8k :

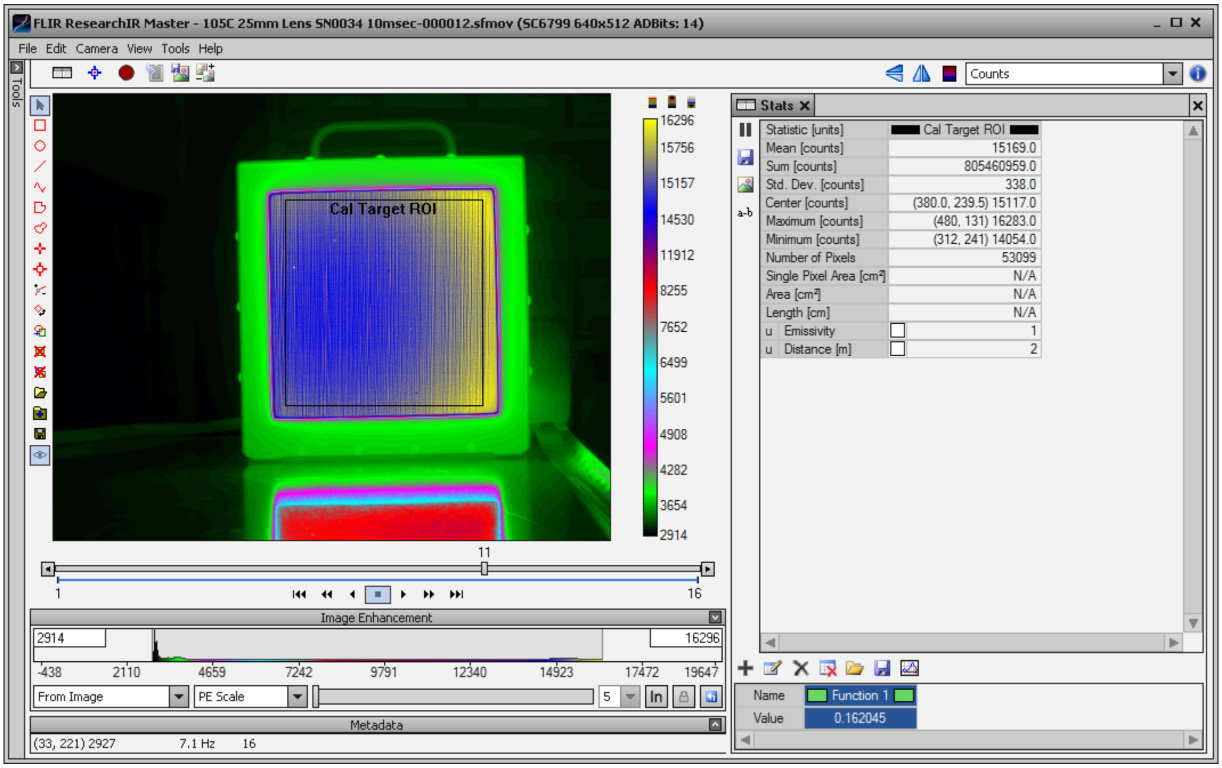


**Figure 8k. A 16-frame movie of a 105 °C blackbody. The NUC is no longer working correctly on the 105 °C blackbody emitter, since the pixels are saturated.**

14) Once you have all the movies recorded, you can then go back and compute the mean counts for each temperature, and then input those mean values and the target temperatures into the "measurement grid" of the RIR utility called CalibratIR. I prefer doing the calibrations this way, rather than using the calibration wizard within RIR. Using CalibratIR makes it easy to go back and recalculate the ROI stats, or you can move the ROI around on the emitter surface. You can also see what the images looked like during the calibration, to look for problems like bad pixels in the ROI that can skew the ROI statistics. If RIR crashes during a session, you can lose the work you did on the calibration. Figure 8l shows how you get to CalibratIR, which I always end up putting into the system tray of Windows so I can launch it quickly.

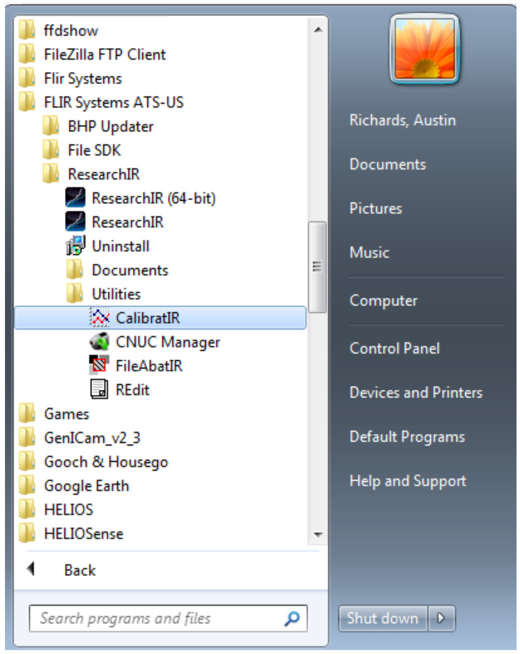


**Figure 8l. CalibratIR located in a subfolder of ResearchIR**

Open CalibratIR and choose the Expert mode if you are adventurous, or you can try the other options that will guide you with a Wizard:

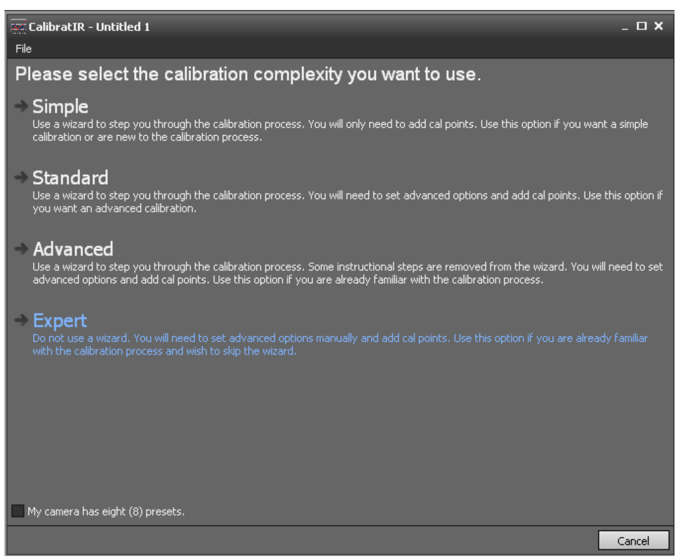


**Figure 8m. Selecting the mode of CalibratIR**

15) The first tab in Expert mode is the response tab. The default is a "Tophat" spectral response with the band limits set to 3 and 5µm. It is called that because the response curve looks something like a top hat. The default spectral limits for the tophat does not give the correct response for the camera I am calibrating in this app note. The SC6700 See-Spot camera has a different spectral response than a standard 3-5µm InSb camera. A

good top hat approximation for the thermal portion of the See-Spot camera's spectral response is 4.5 to 4.7µm. Later in the troubleshooting section, I will show you what happens when you have the WRONG spectral response distribution.

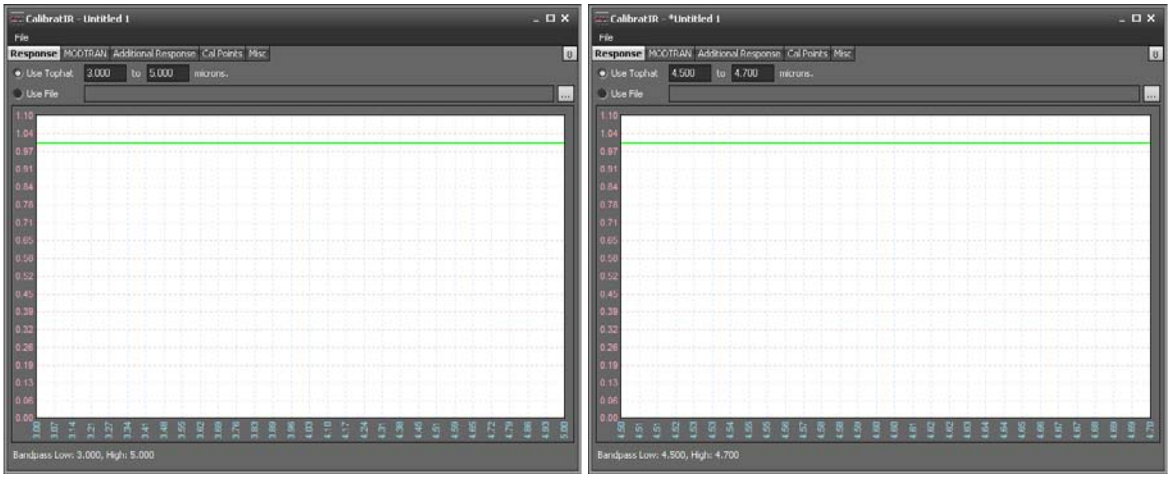


**Figure 8n. Default spectral response and See Spot midwave IR "Tophat" approximation**

16) The other better option is to use a spectral response file. The advantage of this is you will get very nice linear fits to the radiance versus counts data sets, better than if you used the top-hat approximation. FLIR Niceville can measure the spectral response of the camera and lens combination using a monochromator system for an extra cost. A more attractive solution is to obtain a spectral response file for a similar camera. I am using a spectral response file for a different See Spot camera and lens, but the differences are likely very small. The SR measurements done in Niceville will generate a relative spectral response curve file (in the appropriate power units), which can be loaded as the response. The files we provide have the file extension *.prn. It is a text file with two columns; one is wavelength in microns, the other is the relative spectral response normalized to 1 at the peak. Here is what the file I used looks like in a text editor. You can make your own *.prn file by creating a text file with wavelength in microns in the first column and the relative power spectral response in the second, normalized to 1 at the peak. You then change the *.txt file extension to *.prn. Figure 8o shows a screen shot of a *.prn file open in Windows Notepad:

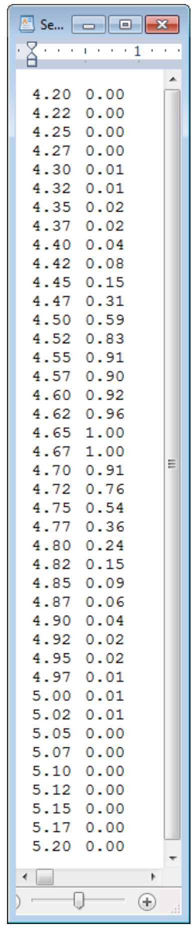

| | |
|---|---|
| 4.20 | 0.00 |
| 4.22 | 0.00 |
| 4.25 | 0.00 |
| 4.27 | 0.00 |
| 4.30 | 0.01 |
| 4.32 | 0.01 |
| 4.35 | 0.02 |
| 4.37 | 0.02 |
| 4.40 | 0.04 |
| 4.42 | 0.08 |
| 4.45 | 0.15 |
| 4.47 | 0.31 |
| 4.50 | 0.59 |
| 4.52 | 0.83 |
| 4.55 | 0.91 |
| 4.57 | 0.90 |
| 4.60 | 0.92 |
| 4.62 | 0.96 |
| 4.65 | 1.00 |
| 4.67 | 1.00 |
| 4.70 | 0.91 |
| 4.72 | 0.76 |
| 4.75 | 0.54 |
| 4.77 | 0.36 |
| 4.80 | 0.24 |
| 4.82 | 0.15 |
| 4.85 | 0.09 |
| 4.87 | 0.06 |
| 4.90 | 0.04 |
| 4.92 | 0.02 |
| 4.95 | 0.02 |
| 4.97 | 0.01 |
| 5.00 | 0.01 |
| 5.02 | 0.01 |
| 5.05 | 0.00 |
| 5.07 | 0.00 |
| 5.10 | 0.00 |
| 5.12 | 0.00 |
| 5.15 | 0.00 |
| 5.17 | 0.00 |
| 5.20 | 0.00 |

**Figure 8o**

Here is what the *.prn file looks like when I loaded it into CalibratIR:

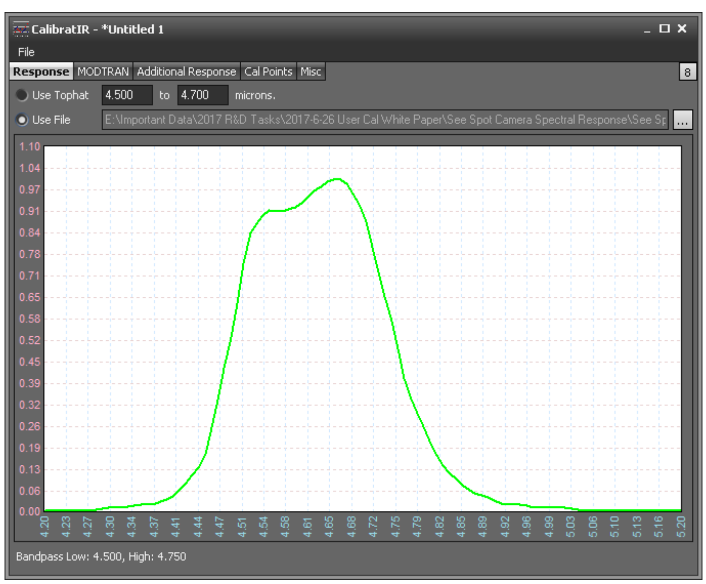


**Figure 8p. Spectral response curve for See Spot *.prn file**

17) There are other tabs in the expert mode view of CalibratIR. One of them is called MODTRAN and it is used for putting in the spectral transmission curve for the air in the path between the camera and the calibration target. The default settings, shown below, have the transmission set to 1 over the passband of the camera. It is just fine to use this when the distance is a meter or less. Some customers who want to calibrate a camera with a long focal length lens will have to place the calibration source at a significant distance from the camera to be within the minimum focus distance of the long lens. One can run MODTRAN software to try to model the transmission, but I have found MODTRAN hard to use, particularly the Ontar Corporation Windows version. We will leave this in the default state, which is 1 at all wavelengths, as shown in Figure 8q. Note: there is a sample MODTRAN file in RIR that one can load to see what these types of files look like. It is called MODTRAN Sample.txt and is located at: C:\Program Files\FLIR Systems\ATS-US\ResearchIR\Samples.

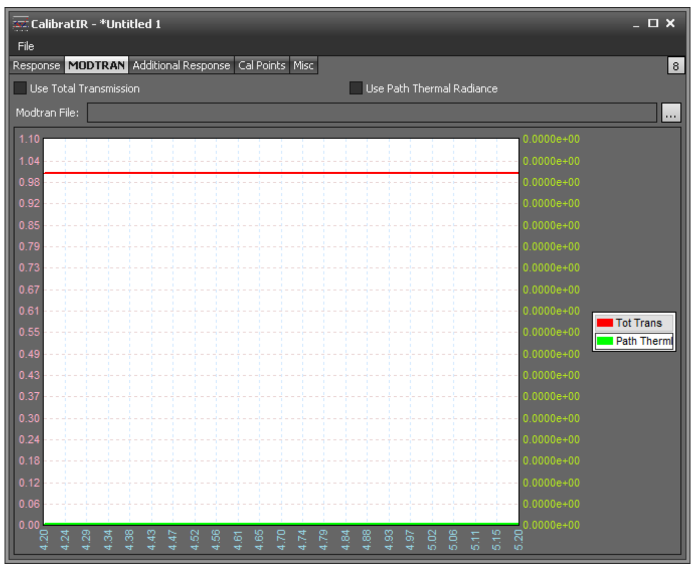


**Figure 8q. Default curve for MODTRAN tab in CalibratIR**

18) The next tab is Additional Response as shown in Figure 8r. It is used to correct the radiance for the presence of a window or a mirror in the optical path. If the spectral response *.prn file you load was generated using a camera without a lens in place during the measurement, you could load the spectral transmission curve for the lens here.

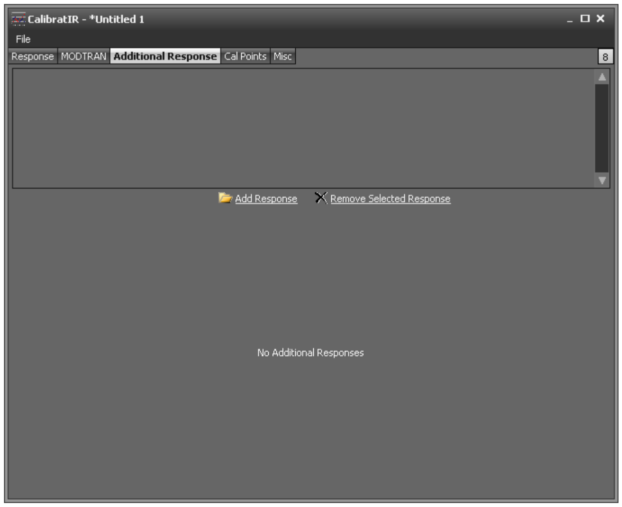


**Figure 8r. Additional response tab in CalibratIR**

19) Now we are ready to populate the measurement grid. The tab is called Cal Points.

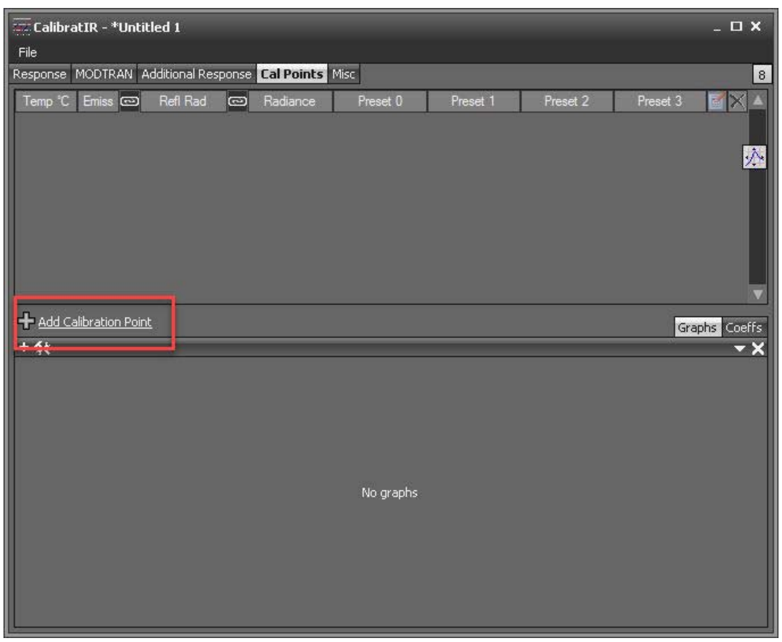


**Figure 8s. Cal Points tab where one adds calibration points to the measurement grid**

The first thing you do is click on the Add Calibration Point link in the middle on the left, which produces this view:

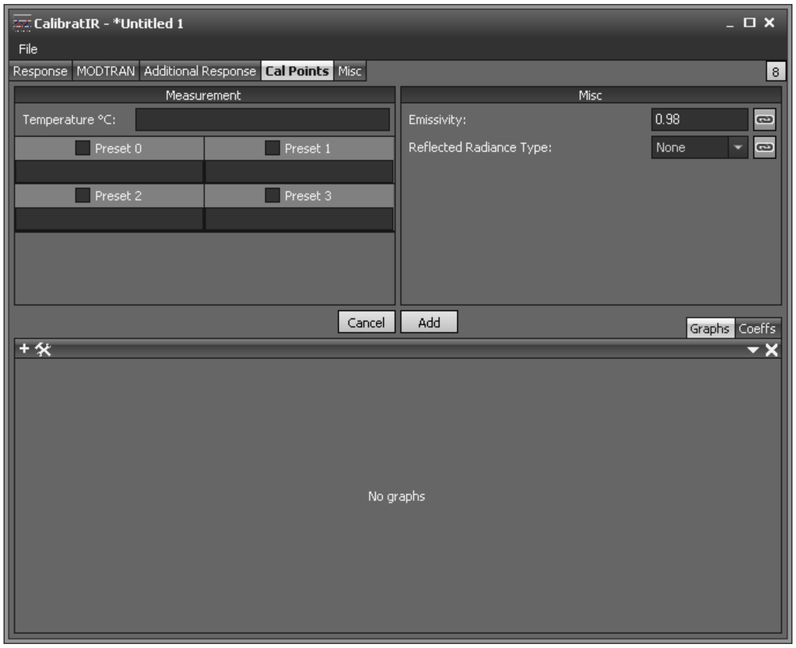


**Figure 8t. Ready to add a calibration point.**

First type in the temperature of the calibration target in the upper left of the window. Then type in the emissivity value in the upper right corner, noting that the default value is 0.98, which I changed to 0.96. Then check the Preset 0 checkbox, enter the mean digital counts in the ROI drawn on the calibration target and hit the Add button. I obtain this ROI mean value with the Frame Average filter on, and play the movie through a few times to get that ROI mean value to settle to a 16-frame average value, a process that suppresses temporal noise:

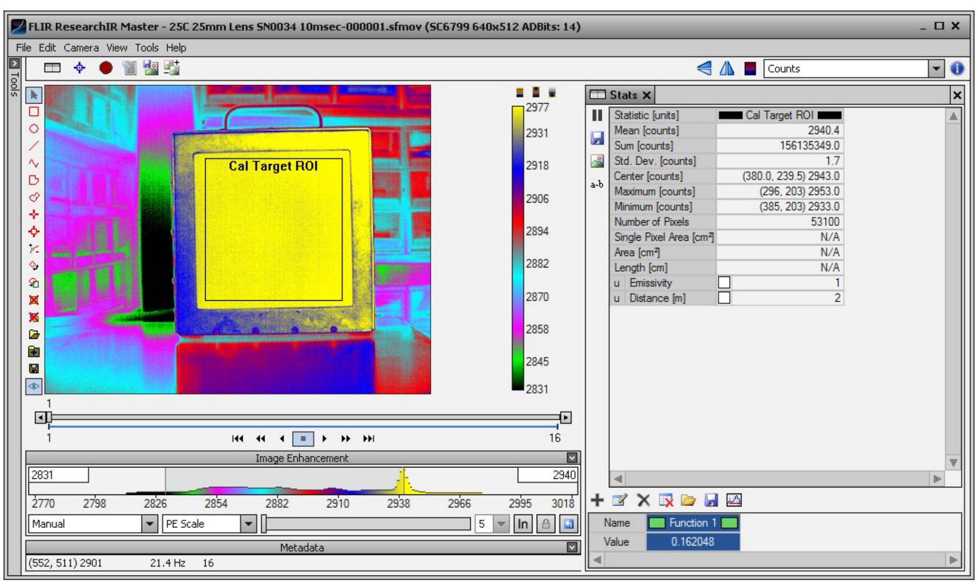


**Figure 8u. Playing a movie with frame average filter enabled to settle the ROI mean value and reduce temporal noise**

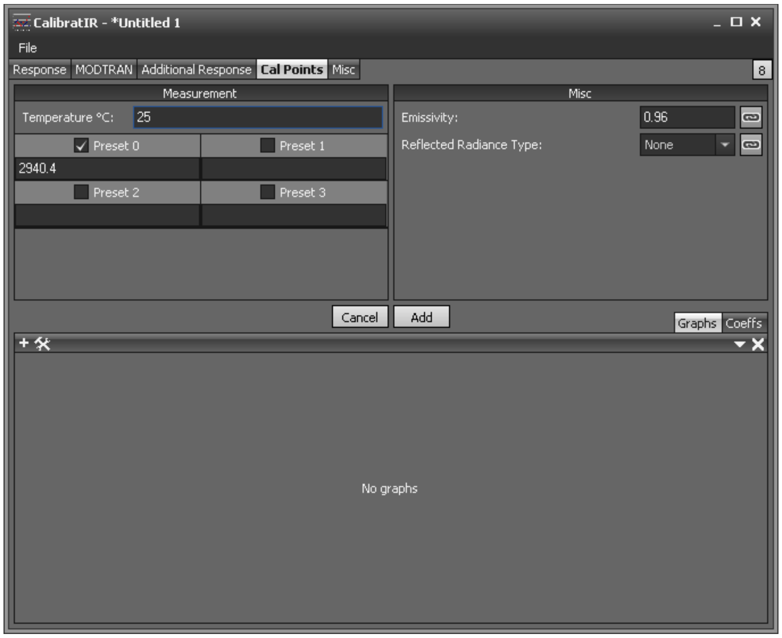


**Figure 8v. Entering the time-averaged ROI mean into CalibratIR**

After hitting Add, the software adds that point to the measurement grid, along with the calculated in-band radiance. The reflected radiance is set to None. Selecting the Reflected Radiance type pull down menu opens a dialog window where you can input the room temperature where you are doing the calibration, and the program will calculate the reflected radiance. In this case, the calculated value is 4% of the radiance of a 25 °C blackbody. This is because the calibration target has an emissivity of 0.96, which means it has a reflectivity of 0.04 and it reflects 4% of the radiation emitted by the surrounding room. You don’t have to use this feature if the calibration target is a cavity blackbody, because the reflectivity of a cavity blackbody is typically less than 1%, and the reflected radiation is negligible. If the calibration sources are much hotter than ambient temperature, then this step can also be skipped, because the reflection of radiation emitted by an object at ambient temperature will be orders of magnitude smaller than a calibration target with a temperature of >200 °C, for example.

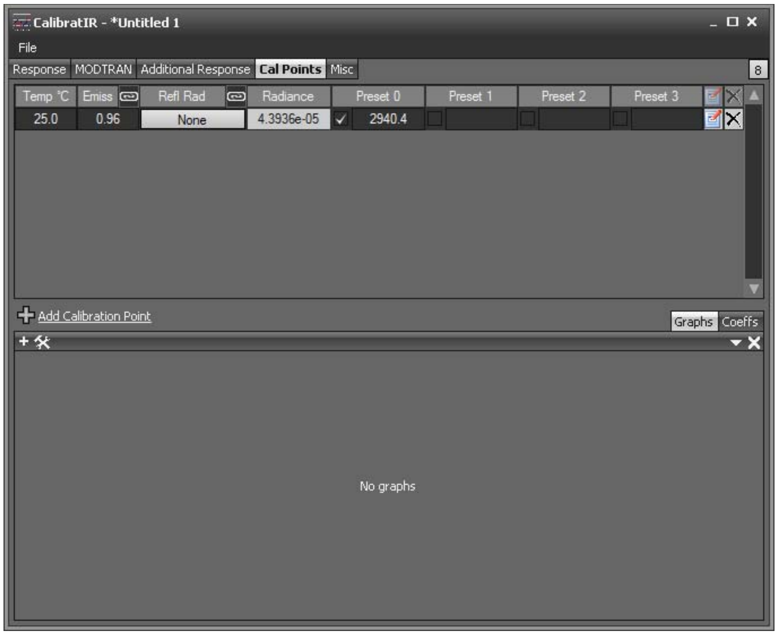


**Figure 8w. Measurement grid getting populated with a calibration point**

When I hit the “Refl Rad” button where it says None, I get this dialog:

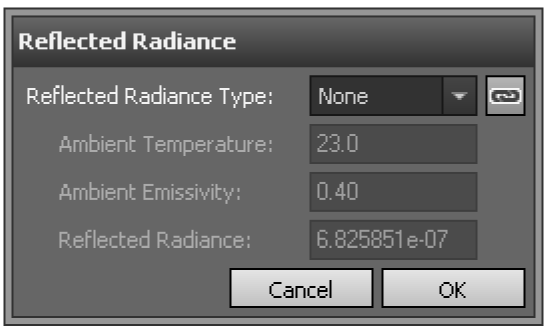


**Figure 8x.**

I now pull down to Calculated type radiance and input the room temperature. Because I completed this calibration in a closed room (essentially a big cavity blackbody), I use “1” as the ambient emissivity:

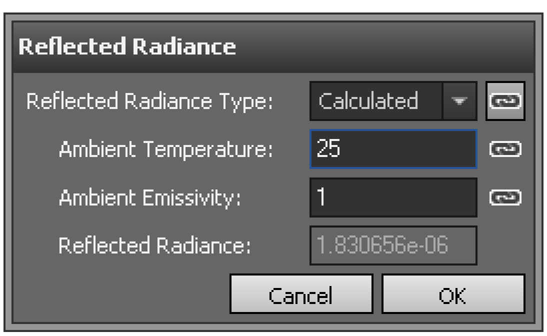


**Figure 8y.**

Note the reflected radiance is a small value, less than 5% of the in-band radiance for a 25 °C target. This reflected radiance correction is useful for cold calibration targets (less than ~20 °C), and it becomes very important for very cold (<0 °C) calibration targets! But you can ignore it for calibrations starting at or over 200 °C.

Now I repeat the process with all the other movies I captured, and I fill out the measurement grid. I have ResearchIR open in one window, and CalibratIR open in another. All of this calibration is taking place on Preset 0, which is why the other presets are blank in the measurement grid. Doing a user calibration on multiple presets for superframing can be done as well.

I did a lot of points for this calibration! Not all of them fit in the displayed measurement grid window size, and there is no way (yet) of making that part of the window bigger so there are two screen shots with a scroll down. Note the graph has a very nice linear shape to it:

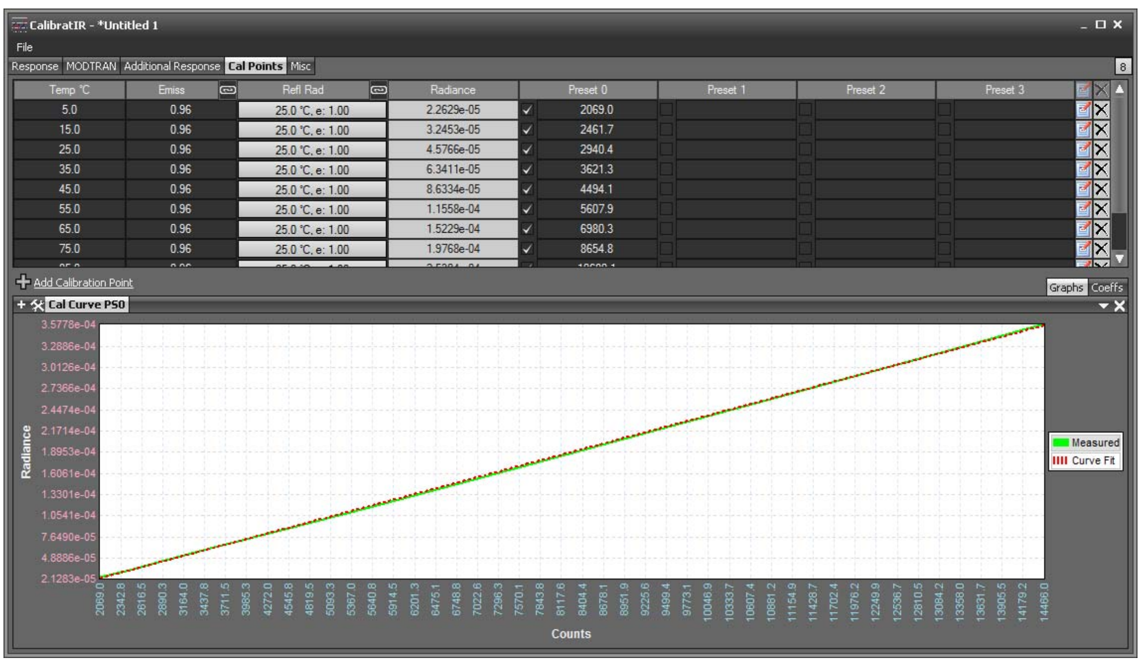


**Figure 8z. A complete measurement grid with lots of calibration points at 10 °C intervals.**

Neither CalibratIR nor the Calibration Wizard in RIR allow you to display the individual points in the graphs, so I have plotted the points with a fitted trendline in Excel to show you how well the data points lie on a line:

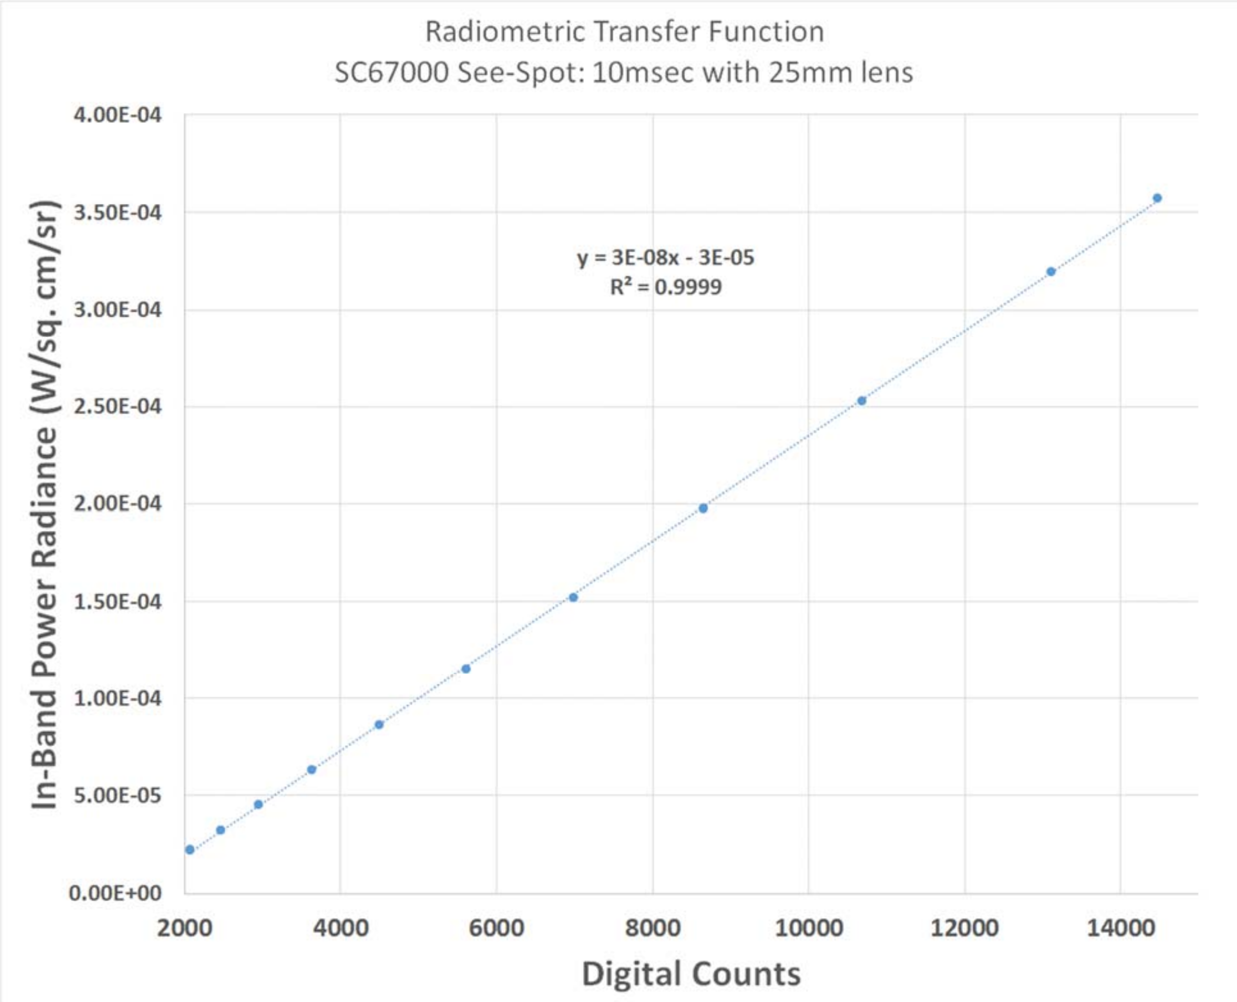


**Figure 9a. Excel plot of the data from the measurement grid**

I click on the "Coeffs" tab to see the goodness of fit, which is 0.9999, or very close to perfect. Note that in the CalibratIR screenshot below in Figure 9b, the counts to radiance polynomial fit is of order 1, a straight line. This is the default value used by CalibratIR and what one should always use. There is no good reason to ever use order 2 or higher, especially because the NUCs that are applied to the raw FPA data are linear transformations.

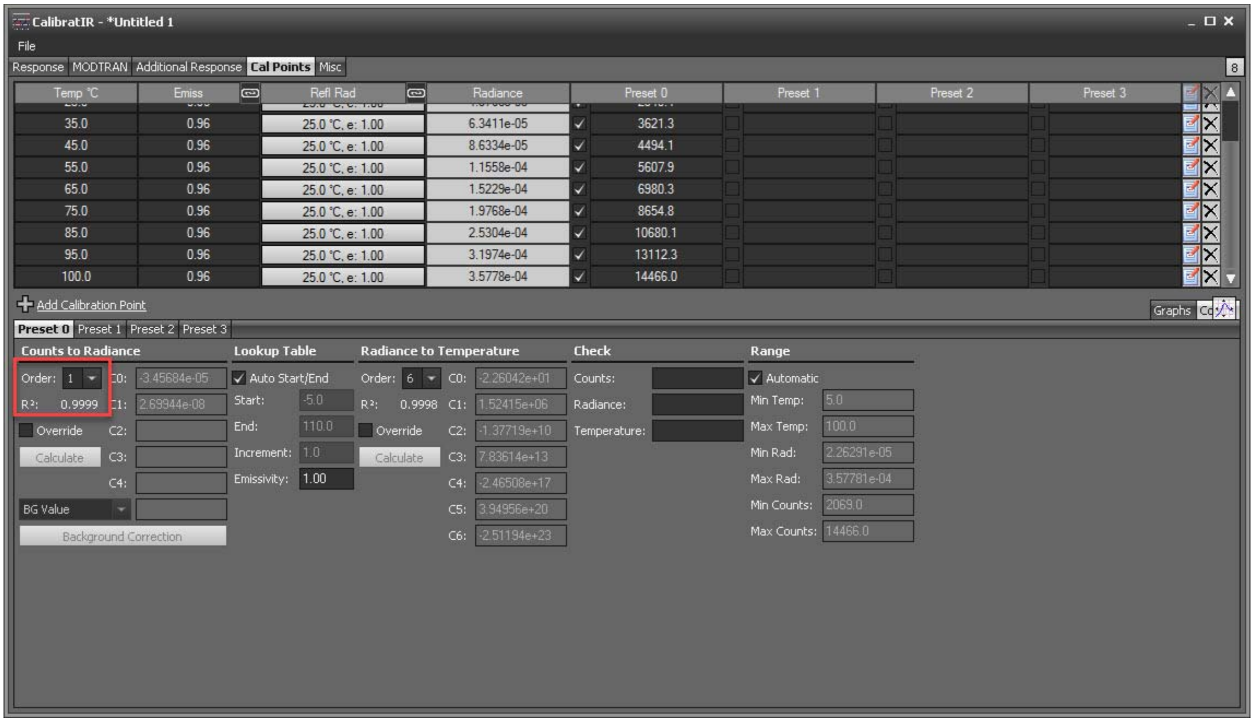


**Figure 9b. Counts to Radiance with polynomial order 1 selected**

The Lookup Table is automatically running a table of temperatures and radiance from -5 to 110 °C. This is not going to work well, because, as we saw earlier, when the target was set to 105 °C, the camera digitizer was saturated and the NUC broke down. It is a good idea to uncheck this Auto Start/End and set the limits to the range of the calibration target temperatures that I used, namely 5 to 100 °C:

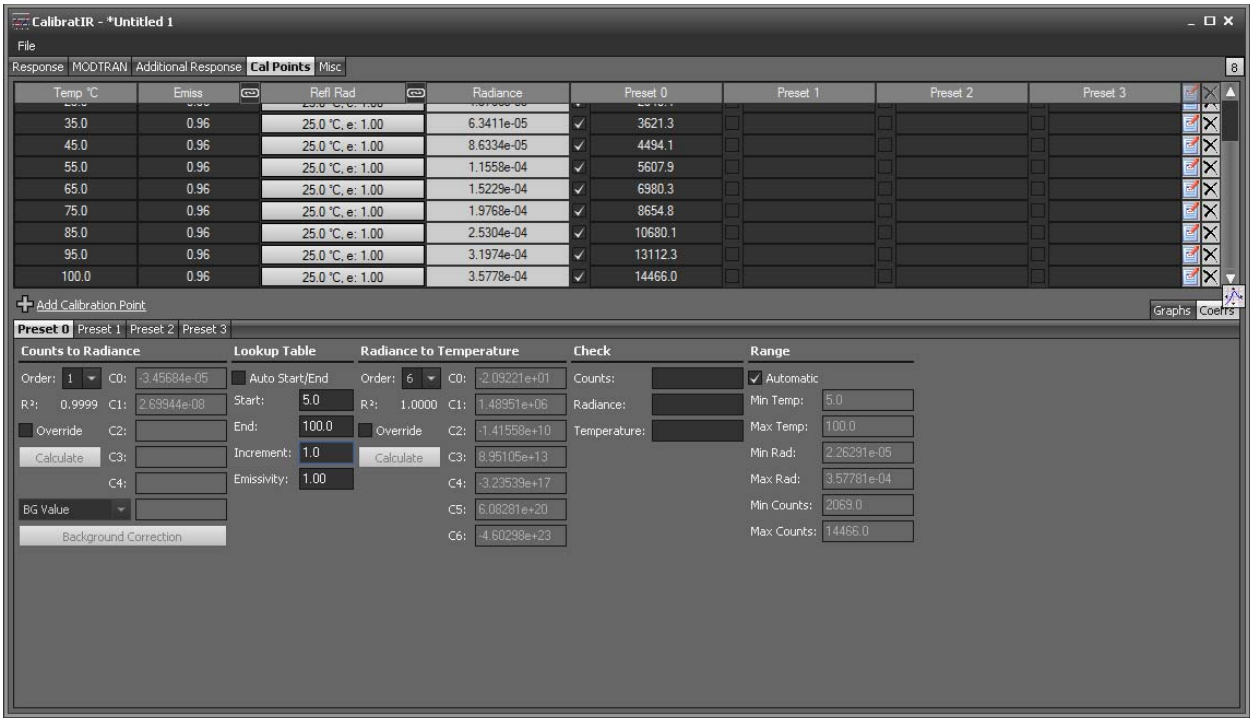


**Figure 9c. Uncheck 105 °C and end the calibration range at 100 °C.**

I can now save this calibration as a *.cal file. The file name is SC6700 See Spot 10msec 25mm Lens SN0034.cal. I then have to go back to File/Save As, and then I need to save as an include file. You need both file types to make this user calibration work fully.

**Note: If you only have the include file, then you can get radiance but not temperature measurements.**

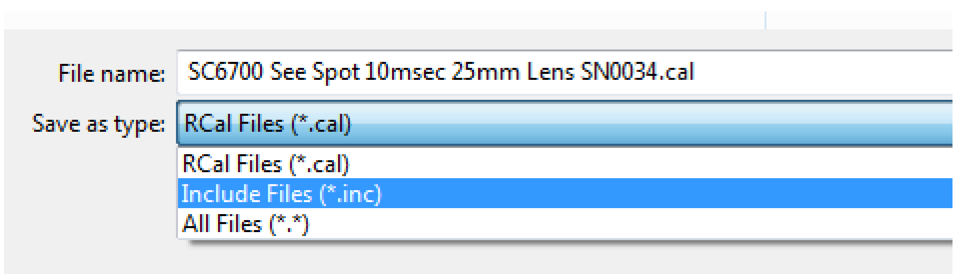


**Figure 9d. Saving a cal file, then saving an include file with the same name. Keep these files together!**

Now I have two files and I am ready to load them into ResearchIR:

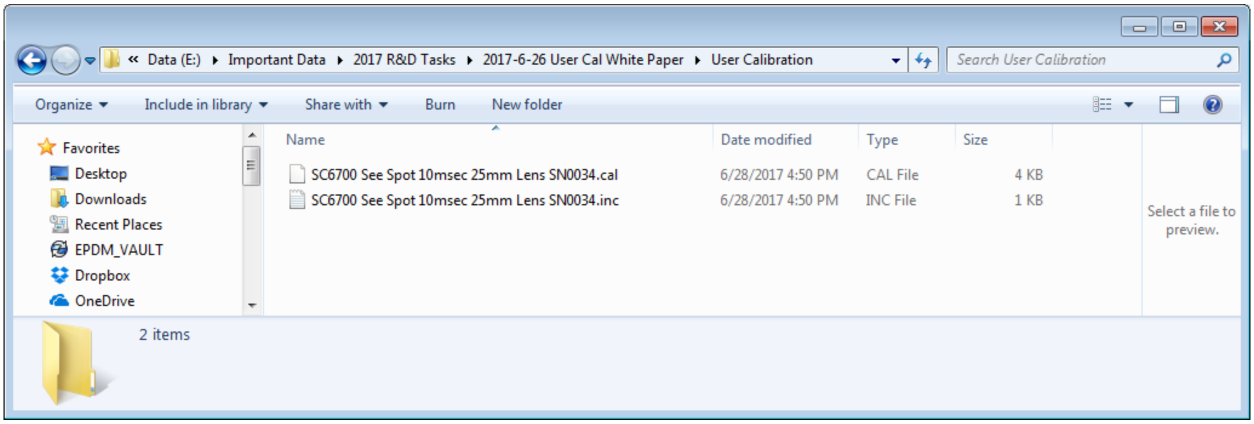


**Figure 9e. Windows folder showing the two files needed for calibration to work**

Now we load either one of the files, and either one will work because the software knows what to do – the other file is automatically loaded as well. **Just always keep the two files in the same directory!**

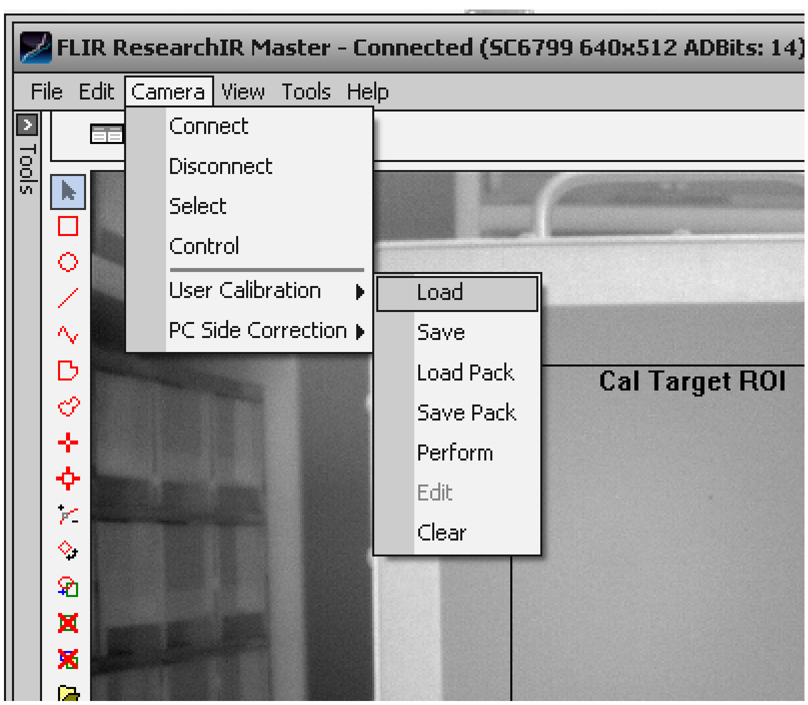


**Figure 9f. Loading calibration files in RIR**

There are two ways you will know that a user calibration is loaded into RIR. One: The Edit option is no longer greyed out. Two: You can now pull down the menu in the upper right corner to get Temperature (User) and Radiance (User). If you already have a factory calibration, the two factory options will be there as well: Temperature (Factory) and Radiance (Factory).

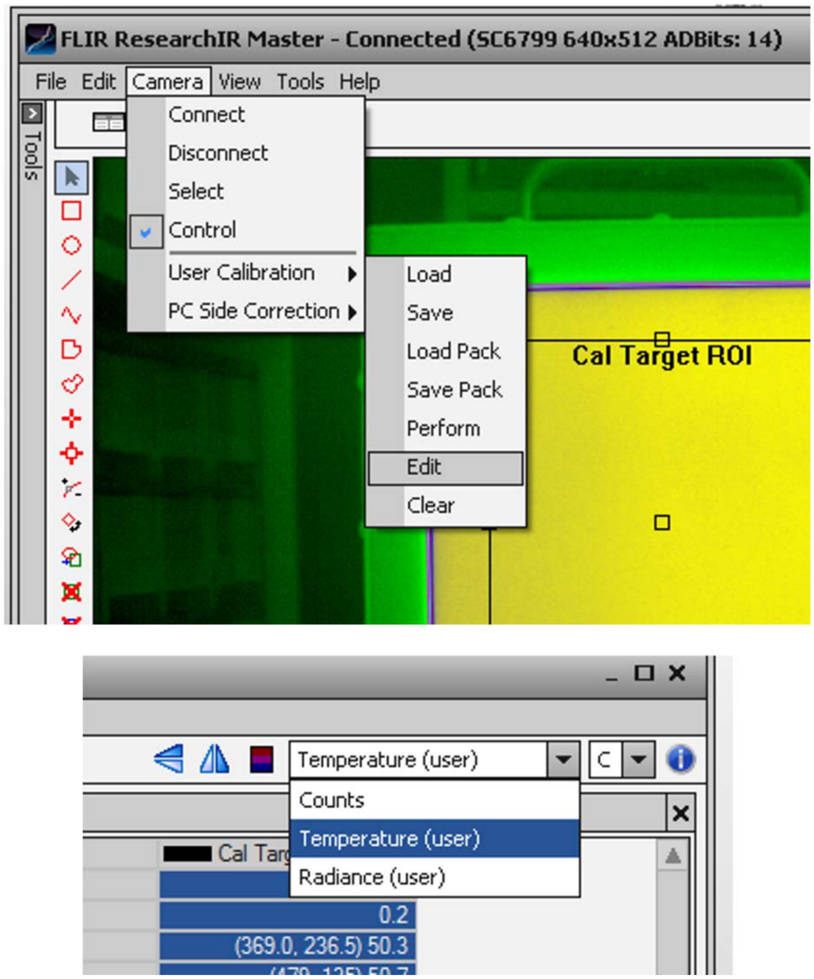


**Figure 9g. Two ways to tell that a calibration is loaded**

We can look at the calibration file that is loaded by going to Camera/User Calibration/Edit. Note that the Auto Start/Stop setting state is not preserved when using the ResearchIR Calibration Wizard when you are connected to a camera. When I saved the *.cal and the *.inc files, the checkbox was unchecked. But…it will revert back to Auto Start/End as being checked. It does not do this in CalibratIR, which is yet another reason why I like to use CalibratIR for generating calibrations and editing them, even when I have a live camera connected to RIR.

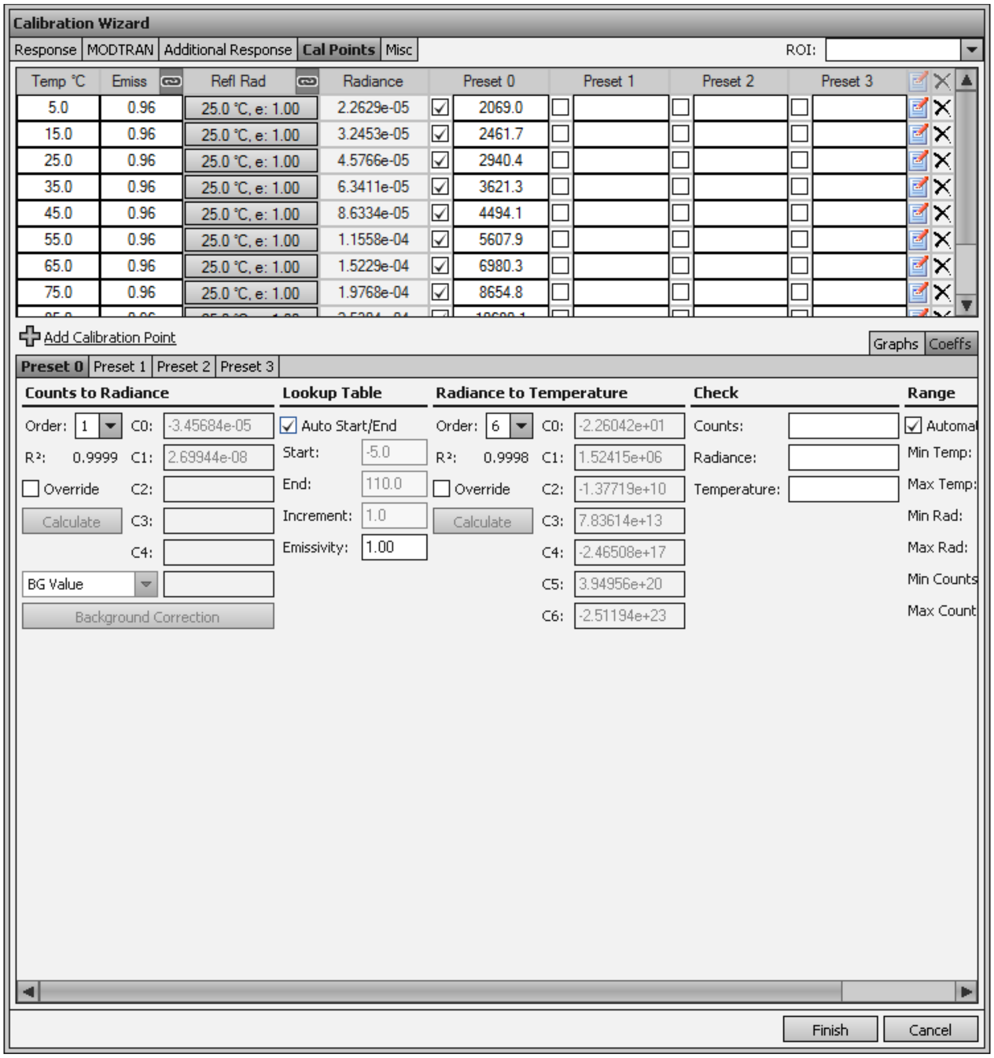


**Figure 9h. Examining the loaded calibration within RIR using the edit calibration function**

You should immediately test the calibration by looking at the calibration target at various temperatures, especially at the limits of the calibration to see if you get the correct temperature values. Here is a test image where I set the calibration target to 48.8 °C, and I get 50.2 °C. I am within 2 °C of the indicated temperature, which is matches the expected accuracy.

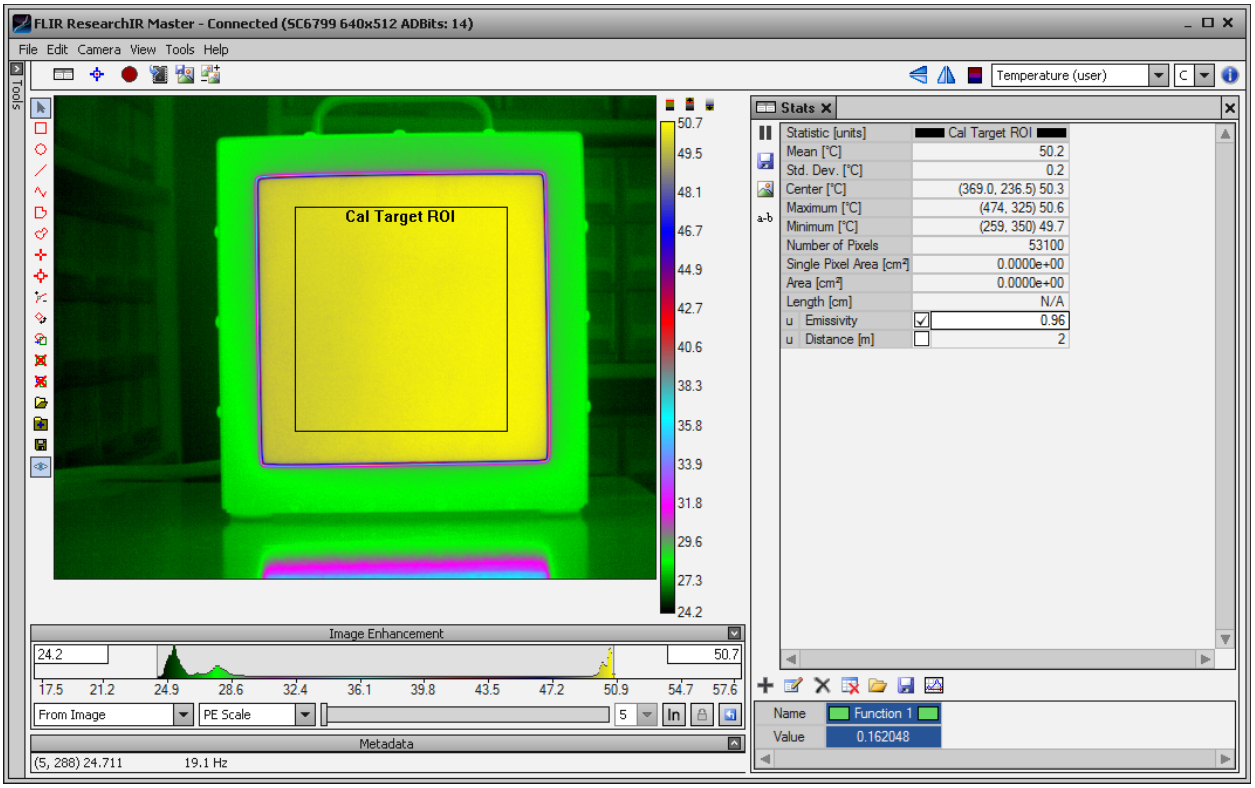


**Figure 10a. Testing a user calibration**

A 5 °C indicated temperature on the blackbody controller gets me a measured value of 6 °C as shown in Figure 10b. My user calibration appears to work well. However, I am cheating a little, because I am using the same blackbody to test the calibration as I did to MAKE the calibration. The emissivity value that I use to make the calibration is the same value that I use in the Stats window. But the true value of the emissivity in the camera's waveband might not be exactly 0.96, and I would not know. It is therefore good practice to use a *different* reference standard blackbody with a precisely known emissivity to test the calibration accuracy. A factory-calibrated cavity blackbody that get frequent re-calibrations is well suited to this task.

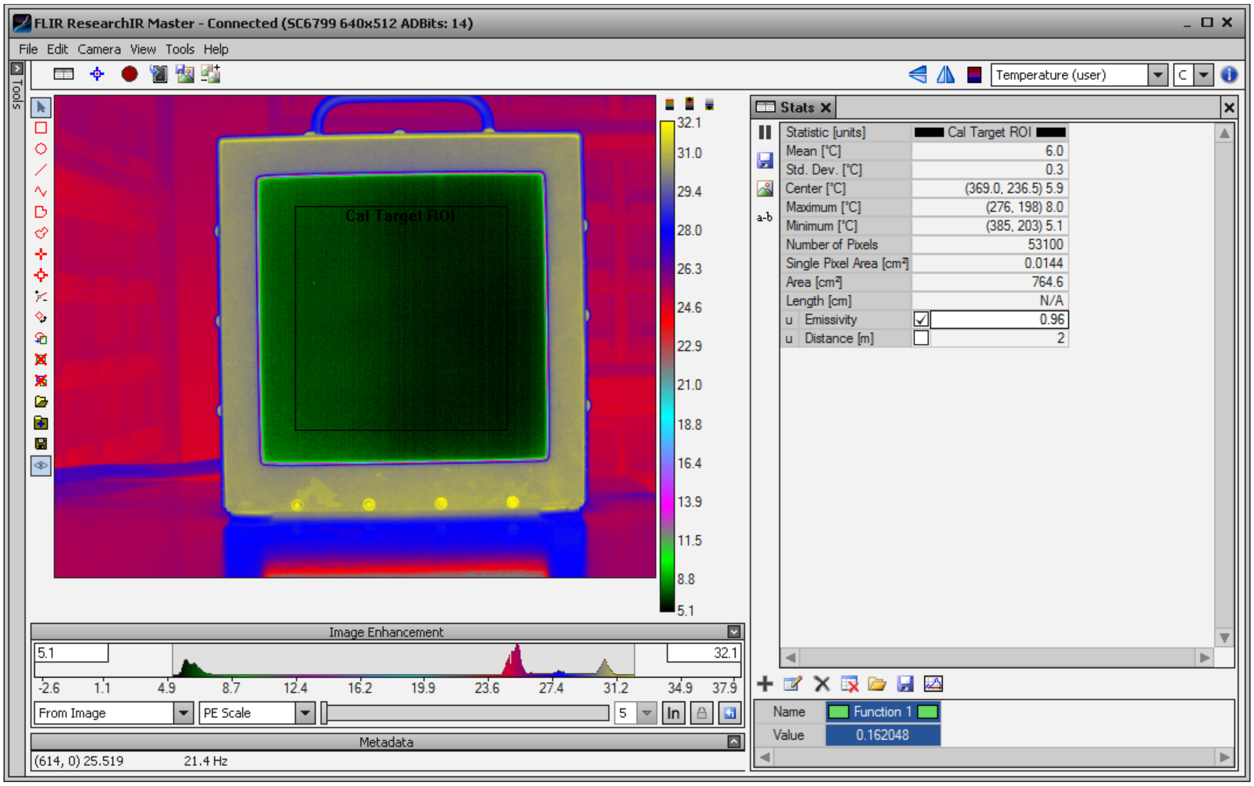


**Figure 10b. Testing a user calibration with a 5 °C blackbody. I get 6 °C as a measurement.**

## Troubleshooting a user calibration

1) Make sure you get points that sample over a pretty wide range of digital counts with fairly even spacing between them in counts units. The maximum ROI mean count value should be about 14,000. Above that and the NUC stops working because pixels in the middle of the array will be saturated, i.e. the digitizer will be at the 14-bit limit of 16,383 or something close to that. The lower end is determined by the temperature range you are working with. I like to get down to about 2000 counts if I can. There is an offset of 500 to 1000 counts in the cameras by design, to prevent one from bumping up against the lower rail of the digitizer with the wings of a distribution. A camera set up for a high temperature calibration will typically be running at a short integration time, and it may also have a neutral density filter in the optical path to further limit the sensitivity. When this low-sensitivity camera looks at a typical 22 °C scene in an air-conditioned laboratory, the image will appear very noisy. For a good calibration, you will want to set your lowest temperature calibration point to a temperature that gives you **at least** 100 digital counts above what you get when looking at an ambient temperature target. For a typical user calibration on a midwave camera that runs from -20 °C to 50 °C, you may not have a calibration source that goes below 5 °C or so. Just go as low as you can with the calibration sources. The calibrations can be extrapolated to measure down to lower temperatures (and higher temperatures) than the calibration points you took in the measurement grid by unchecking Auto Start/End and manually changing the limits. The limits and the state of the checkbox will be preserved through a Save of the *.cal and the *.inc, an exit of CalibratIR and a reopening of the calibration, unlike what happens when one uses the Calibration Wizard.

2) If you use the wrong spectral response, the curve of counts versus radiance will be curved. It should always be a very straight line! The camera FPA and digitizer electronics are designed to be very linear. One common error is using the default 3-5µm spectral range values even when the camera spectral response is something different. Here is a screen shot of the graph from CalibratIR with the wrong spectral response:

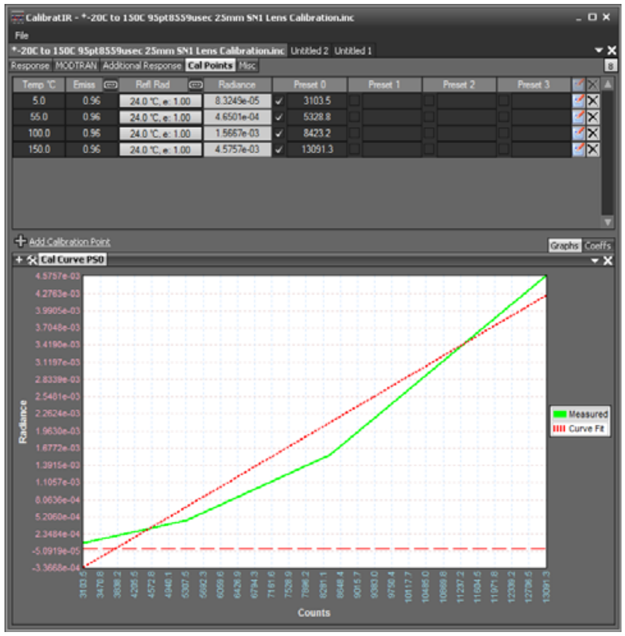


**Figure 11a. User calibration for SLS camera with 7.5-10.5µm sensitivity with the spectral response *incorrectly* set to Tophat and 3-5µm. There is severe curvature in the data because the radiance values are being calculated incorrectly.**

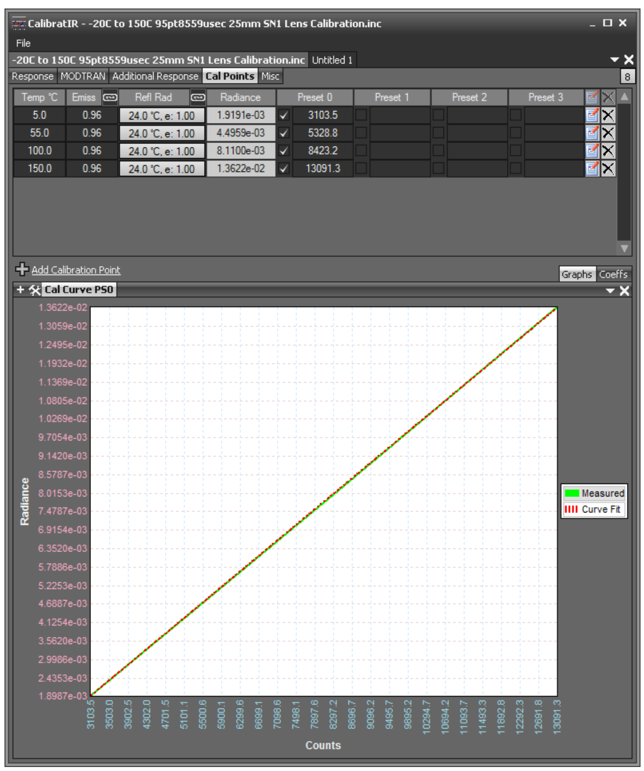


**Figure 11b. User calibration for an SLS camera with a bandpass of 7.5-10.5µm with the spectral response *correctly* set to "Tophat" with correct limits of 7.5 and 10. 5µm. Now there is a nice linear fit with no curvature.**

3) Figure 11c is an example of a calibration that is not useful. The camera is saturated, and all the count values in the measurement grid are at nearly the same saturated value. The curve fit to the data is terrible. The most likely problem here is that the integration time of the camera is far too long. I would want to see counts at the maximum temperature of 1929C around 14,000 but no higher, and the counts for the 853 °C point should be on the order of 2000 counts, or ~1000 counts above the noise floor. The ROIC in the FPA is most linear between about 2000 counts and 14,000 counts. It is best to stay in that range.

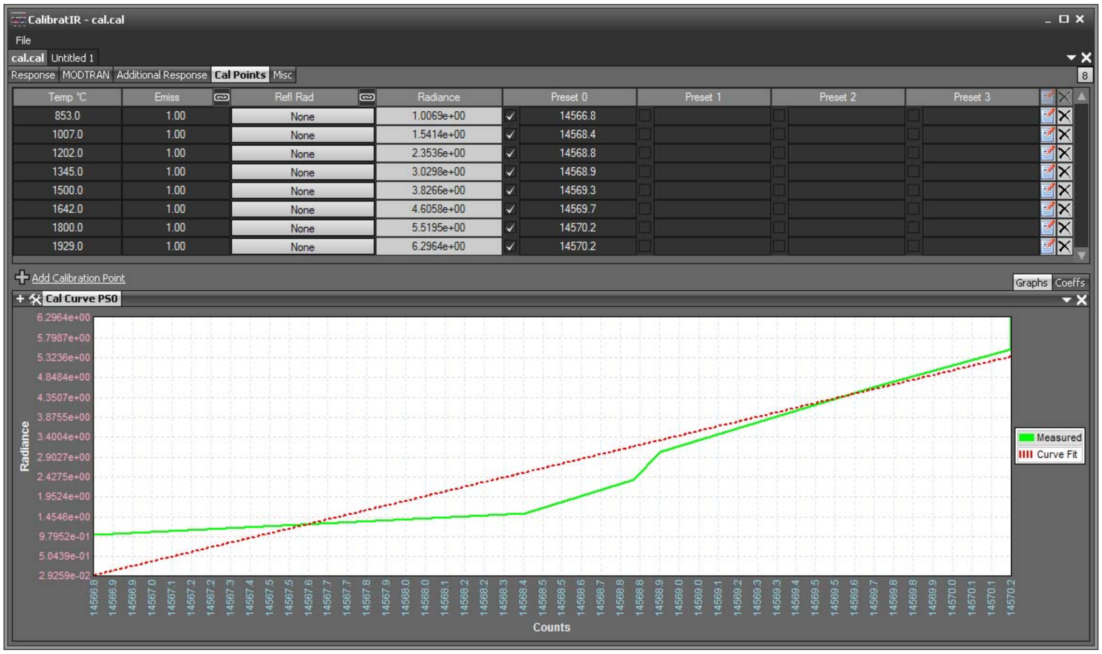


**Figure 11c. Very poor calibration!**

4) The *.cal file is all you need to get temperature and radiance values, but you should ALWAYS have both the *.cal and *.inc files in the same directory as each other, so RIR can find both! If you forget to keep the *.cal file in there, the temperature measurements will be incorrect, and you may just see a black screen, even though radiance still works. Radiance measurements only need the *.inc file, but temperature measurements need both files!

5) The *.inc file has some other useful things you can add to it, like the two IFOV values that you need to measure radiant intensity. This is done in the Misc tab of CalibratIR, shown in Figure 12, or with a spatial calibration when you are connected to a camera in RIR, or the two IFOV tags can be manually added to the include file using a text editor like Notepad or WordPad. The tags are called IHFOV and IVFOV, and they are numeric values with microradian angular units. You can determine their values by dividing the pixel pitch in microns by the lens focal length in millimeters and multiplying by 1000.

All thermal IR FLIR science cameras have square pixels, so the two tag values will be the same.

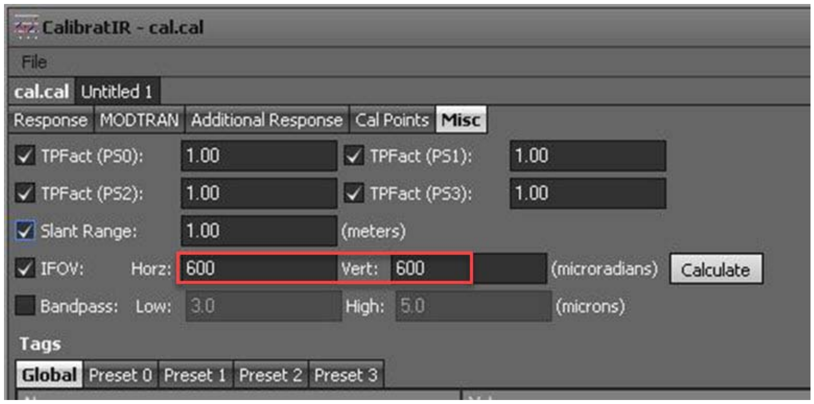


**Figure 12. Misc tab of CalibratIR showing where you can edit in the IFOV tags**

6) Go to Camera/User Calibration and Save either file type (*.inc or *.cal), and you save the other file automatically.

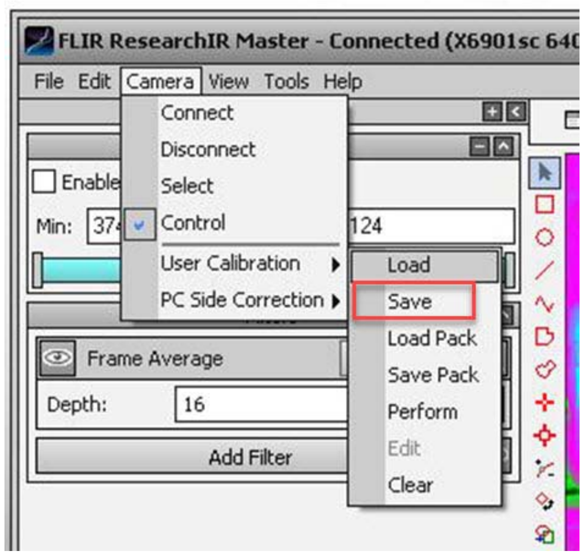


**Figure 13. Saving user calibrations**

7) If you perform a spatial calibration using the Edit/Spatial Calibration window, which you can only do when you have a camera connected to RIR, and then you save the user calibration, the IFOV tags will NOT show up in the new include file! You have to edit the calibration and add the IFOV tags manually here in the Misc page of the Calibration Wizard (which again you can only access when you are connected to a camera), then save the user calibration, and the tags will be added to the include file. To ensure they are in the include file, you should view the file in either Notepad or REdit. I always have my Windows set up to open a file with the *.inc extension in Notepad.

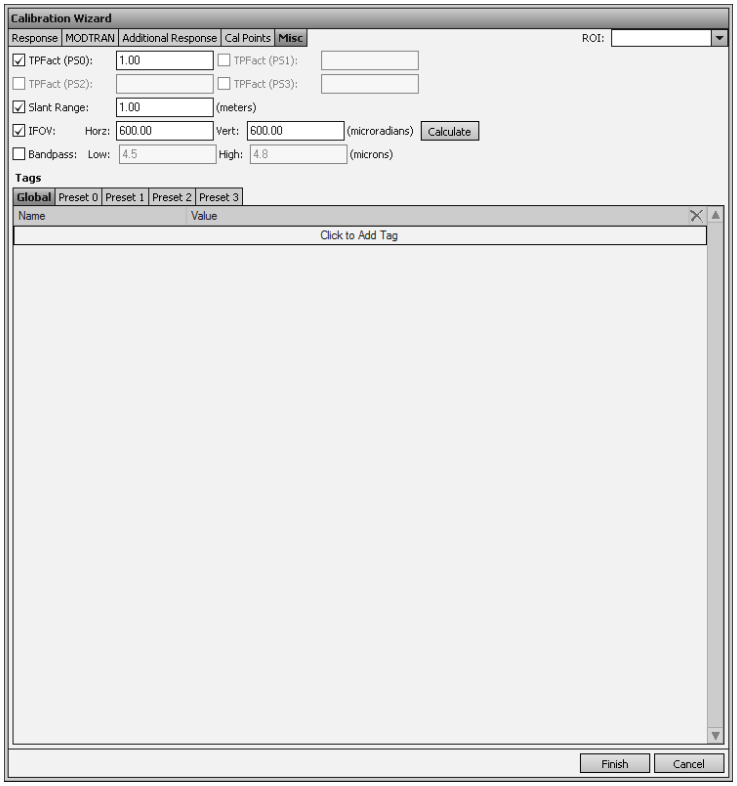


**Figure 14. Misc tab of the calibration wizard**

Figure 15 shows a screen shot of an include file with the IFOV tags marked:

```
HdSize auto
Group Misc Tags
BgFile_0
BgType_0 none
BgValu_0 0
CAFile -20C to 150C 95pt8559usec 25mm SN1 Lens Calibration.cal
Coeff0 -1.739221e-03
Coeff0_0 -1.739221e-03
Coeff1 1.172209e-06
Coeff1_0 1.172209e-06
DaUnit_0 W/(sr-cm^2)
EuRaw_0 Raw
IHFOV 1.000000e+03
IVFOV 1.000000e+03
PolyOrder_0 1
SBPLo 7.500000e+00
SBPUp 1.050000e+01
SltRng 1.000000e+00
StdUnt_0 17
TempCoeff0 -5.877592e+01
TempCoeff0_0 -5.877592e+01
TempCoeff1 4.500317e+04
TempCoeff1_0 4.500317e+04
TempCoeff2 -7.509197e+06
TempCoeff2_0 -7.509197e+06
TempCoeff3 9.549586e+08
TempCoeff3_0 9.549586e+08
TempCoeff4 -7.223645e+10
TempCoeff4_0 -7.223645e+10
TempCoeff5 2.912450e+12
TempCoeff5_0 2.912450e+12
TempCoeff6 -4.805150e+13
TempCoeff6_0 -4.805150e+13
TempPolyOrder_0 6
TPFact 1.0
TPFact_0 1.0
DATA
```

**Figure 15. Screen shot of an include file with the IFOV tags highlighted**

# Chapter 4 : Radiometric Measurement Accuracy

Introduction ........ 4—221
Terminology ........ 4—221
Definition of Accuracy ........ 4—222
Temporal and Position-Dependent Radiance Errors ........ 4—222
Radiometric Corrections to Radiance Values ........ 4—225
Object Parameters in ResearchIR ........ 4—226
Simulating Camera Performance with a Software Model ........ 4—229
Emissivity Error ........ 4—231
Cold Targets and Emissivity Error ........ 4—231
Hot Targets and Emissivity Error ........ 4—232
Pathological Emissivity Case ........ 4—232
Reflected Temperature Error ........ 4—232
Additional Sources of Error ........ 4—233
Atmospheric Transmission ........ 4—233
Radiance Measurements ........ 4—234
Source Radiance and Apparent Radiance ........ 4—234
Relative Radiance ........ 4—235
Accuracy at Performance Limits ........ 4—235
Very Short Integration Time: ........ 4—235
Radiometric Measurement of Small Targets ........ 4—239
Radiant Intensity of Small Targets ........ 4—242
Radiometry on Very Cold Targets: ........ 4—243
Atmospheric Transmission ........ 4—245
Rocket Fuel Tank Measurement ........ 4—245
Summary ........ 4—248

## Introduction

The accuracy of radiometric measurements is highly dependent on many effects, some of which are beyond the control of the engineer performing the measurements. In ideal laboratory conditions, the temperature measurement error from a calibrated infrared camera can be less than 1% of the measured temperature, or 1 °C, whichever is larger. Measurement accuracy of radiance and radiant intensity are harder to define, because the quantity being measured is the spectrally weighted in-band radiance, which is highly dependent on the spectral response of the camera/lens system. The spectral response is never a perfect "square band", with zero response outside band limits and flat response inside the band limits. FLIR Niceville does not certify accuracy for radiance measurements, only temperature measurements. The camera records digital counts which are converted into in-band radiance values according to a calibration. The temperature is then derived from the radiance measurement, after making corrections for target emissivity and other effects described in this chapter.

## Terminology

The terminology of radiometry is worth reviewing. A **target** is an object or point in a scene that is measured by a camera. The temperature that an infrared camera measured is the **apparent temperature**, that is, it is a temperature value that is derived from the infrared radiance of the surface. If the camera and radiometry software is given the correct values of parameters associated with the target and the environment of the target, then the apparent temperature value can closely match the **kinetic surface temperature**, which is the temperature given by a "perfect" temperature probe placed against the target surface[14]. I will refer to this as **KST** for short. Note that the target needs to have some kind of surface to it – it can be solid or liquid, but if it is a gas, then there is no good way to define the kinetic surface temperature of it. Similarly, a target that has some degree of transmission through it can be problematic for temperature measurement. The current FLIR science camera software does not have a means to accommodate radiance transmission through a target; the assumption is that targets are opaque and that they only emit or reflect infrared radiance.

For translucent targets like gas clouds or thin materials like plastic film, it is often useful to measure what I call the **apparent blackbody temperature**; this is the apparent temperature with the emissivity value set to 1. The camera measures a radiance value, and the temperature of a blackbody emitting that same value of radiance is the apparent blackbody temperature. It is a useful quantity because it gives the camera operator and anyone who looks at the data a physical sense of the target: saying that the apparent blackbody temperature is 900 °C is more illuminating as to whether the target is hot or cold than stating an in-band radiance value in units of Watts/cm$^2$/steradian.

[14] A perfect temperature probe is one that does not change the temperature value of the surface that it is measuring. In reality, any physical temperature probe either adds or subtracts thermal energy from the surface it is touching.

## Definition of Accuracy

Temperature measurement accuracy is defined as the deviation of a measurement from the true KST value of the temperature. Radiance and radiant intensity measurement accuracy would be defined in the same manner to be the deviation of the measurement from the "true" value.

Currently, Niceville science cameras have a factory-certified accuracy of +/-2 °C or +/-2%, whichever value is larger. An example of this level of measurement accuracy would be a target with an apparent temperature of 48.2 °C, but a kinetic surface temperature of 50.0 °C. Another example would be an object that measures at 257 °C but has a KST of 300 °C. In the latter case, the temperature error is 3C, which is less than 2% of the 300C actual temperature numeric value. Under some conditions, FLIR calibrations may have +/-1 °C or +/-1% accuracy, but this is hard to achieve, especially for very cold targets, and thus is not guaranteed. These levels of accuracy are only achieved when cameras are calibrated with laboratory blackbody sources that have an emissivity that is known and very close to 1. Cavity blackbodies typically will have emissivity values of 0.99 or greater, and they are by far the preferred tool for radiometric calibration. The blackbodies must be "in calibration", meaning that their stated accuracy condition is being met because they have been calibrated with a traceable temperature standard and that calibration is "fresh" enough to hold true. Over time, blackbodies will go out of calibration, and will need to be recalibrated. Blackbody manufacturers recommend annual recalibrations, which can be expensive and time consuming.

### Temporal and Position-Dependent Radiance Errors

The steps that lead to a temperature measurement using ResearchIR software start with a measurement of digital counts in a region of interest in the image. These digital count values are then converted to an in-band radiance measurement. Finally, the radiance is corrected using a number of radiometric parameters to be described in more detail below. Right off the bat, before the radiance is corrected, there are several sources of error in the radiance measurements. They are temporal noise error and position-dependent error. The temporal noise errors are caused by variations in a single pixel's signal over time which occur because of shot noise in the scene, and electronics noise in the camera. Figure 1a shows a plot of a 40 °C blackbody's radiance versus time over a 45-minute period. The peak-to-peak temporal error variation is 0.26% during that time. Figure 1b shows an image of the blackbody with the region of interest drawn in red.

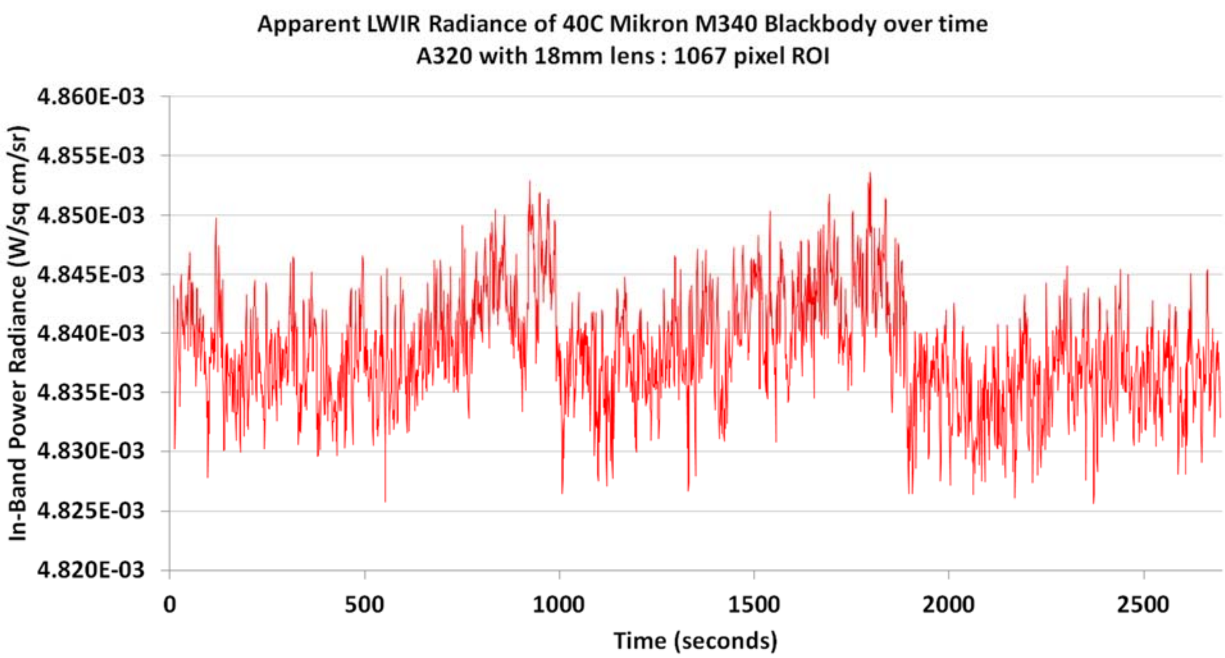


**Figure 1a. Radiance versus time plot for 40 °C Mikron M340 blackbody**

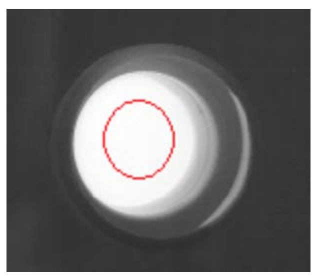

**Figure 1b. Region of interest on image of Mikron blackbody**

A single pixel in that region of interest will show a larger fluctuation in its count value. When a number of pixels are averaged, as in Figure 1a-b above, the measurement noise will be lower. A single pixel was tracked from the same dataset and the peak-to-peak temporal radiance error variation was 0.34%.

IR cameras are corrected with non-uniformity corrections (NUCs) that mitigate the spatial variations in the image. There is always a little bit of residual spatial non-uniformity left over, even with a very fresh NUC. It is absolutely critical to have the freshest possible NUC in place before radiometric image data is captured. An experiment was performed to quantify the remaining non-uniformity and its effect on radiometric measurements. The Mikron blackbody was imaged and the camera was moved around in pan and tilt while the mean radiance for the whole image was measured. Image segmentation was used to avoid having to use a moving ROI for each frame of the movie – only the pixels above a certain preset threshold are counted in the radiance statistics. The results are shown in Figure 2a and 2b below. The mean radiance values varied by 0.24%.

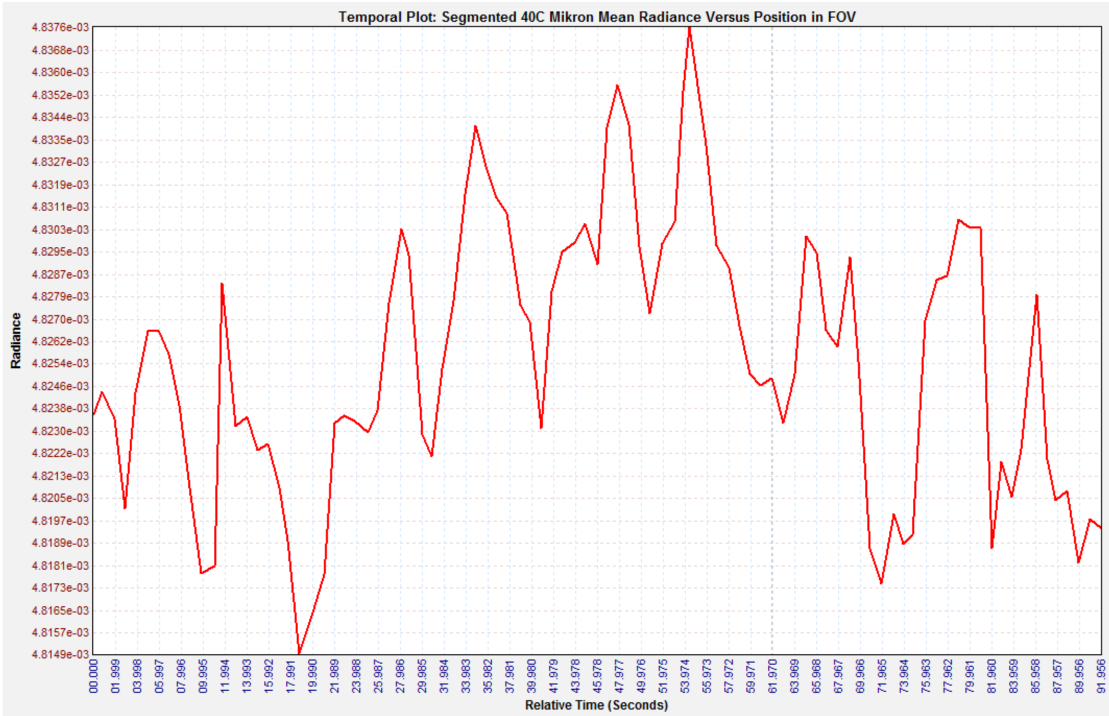


**Figure 2a. Mean radiance value as a function of position within the field of view**

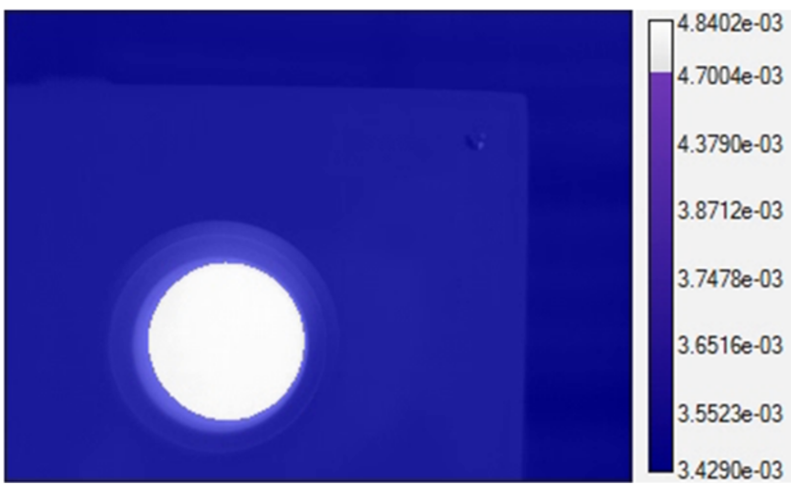


**Figure 2b. Mikron M340 blackbody imaged by longwave FLIR science camera with segmentation so that the measurement can be done with an ROI that encompasses the whole image**

The temporal and spatial error values for the ROI are 0.26% and 0.24%, respectively. These errors are assumed to be uncorrelated, so they can reasonably be added in quadrature using the RSS method. The error percentages are squared, added together, and the square root of the sum is taken. When this is done, the result is 0.35%. This is the percentage error in the radiance measurements of this blackbody. This error can be translated back into a temperature error by the use of a radiometric camera model which is discussed in more detail below. The result is a temperature error of +/-0.1 °C, a very small value. These two effects can generally be neglected. There are far more powerful sources of error that affect measurements of real-world targets in real-world situations.

### Radiometric Corrections to Radiance Values

Many use cases for radiometric cameras are not laboratory-based – they require that the camera be operated outdoors, and often at a significant distance from the target of interest. When the cameras are used outdoors on non-ideal targets, the measurement accuracy can be significantly worse over what is possible in the highly controlled environment of a laboratory. There are a number of reasons for this. The apparent temperature is the temperature of an object as determined by the infrared radiation emitted from it.[15] The goal of radiometric temperature measurement is to be able to measure apparent temperatures on a surface without physical contact, and for the measurement to be able to be "corrected back" to the KST. The correction involves knowing four things:

1. The emissivity
2. The reflected temperature
3. The atmospheric transmission
4. The air path temperature

If the emissivity of the real-world object was 1, then the reflected radiance off the object surface would be zero, and it would not have to be compensated. Put another way, the object would have the same apparent temperature if the reflected temperature was -20 °C (typical for an outdoor scene under a clear sky), or +20 °C (typical for an indoor scene where the reflected temperature was the temperature of the walls, ceiling and floor). Because real objects never have an emissivity of exactly 1, the emissivity typically has to be accurately known to better than 1% in order to have a chance at measuring a temperature value that is within 1% or 1 °C of the kinetic surface temperature.

The reflected temperature effect on a target measurement is very small when emissivity values are 0.9 or higher. If the reflected temperature changes by a few °C, typical target temperature measurements are hardly affected. Here is an example:

A target has a KST of 40 °C, and an in-band emissivity of 0.95. The in-band reflectivity is only 0.05 by conservation of energy (assuming that the object is not transmissive). If the reflected temperature is 23C (a ceiling for example), and there is only a 5% reflection of that temperature, then the reflected radiance value will be small compared to the emitted target radiance. For a 3-5μm midwave camera system, the emitted radiance is 0.95*L(40 °C) = 2.8e-4 Watt/cm$^2$/sr. The

---

[15] While it is certainly the case that other parts of the electromagnetic spectrum can be used to measure temperature, we restrict ourselves to infrared for the purposes of this discussion.

reflected radiance is 0.05*1.62e-4 Watt/cm$^2$/sr, or 8.1e6 Watt/cm$^2$/sr.  This is only 2.9% of the emitted radiance.  If the temperature of the ceiling heats up to 24 °C, the reflected radiance changes to 8.4e-6 Watt/cm$^2$/sr, which is 3.0% of the radiance.  The radiance the camera sees only changes by 0.1%, which has a negligible effect on the temperature measurement.  If the ceiling heats up to 25 °C, the apparent temperature changes to 40.05 °C, which is such a small change that it will be very difficult or impossible to even see it in the data, which is likely to have a temporal noise >0.1 °C.   In reality, one can and should compensate for the reflected temperature. This is done with Object Parameters.

### Object Parameters in ResearchIR

FLIR ResearchIR software has a set of inputs called Object Parameters.  The reflected temperature is one of those inputs.  If a laboratory measurement of a target temperature is performed with a goal of maximum accuracy, then the reflected temperature can and should be monitored to make sure that the Object Parameter entry for reflected temperature is always correct.  This reflected temperature can be controlled by placing a uniform temperature panel in a position where reflected rays from the panel reach the camera. The panel can have temperature sensors placed on it and the values monitored.

Running temperature measurement experiments in an air-conditioned laboratory space is highly recommended, as well as turning on the camera with the correct lens installed at least 2 hours before precision measurements are to be made. This allows the camera/lens system to reach thermal equilibrium.  For factory calibrations, the Tdrift correction works best when everything is stable in temperature.  Tdrift correction is a proprietary algorithm that runs continuously in Niceville science cameras.  The algorithm measures the temperature of points inside the camera housing and on the lens interface and applied a radiance offset based on those values.  The offset is computed using measurements of a blackbody by the camera system immersed in a thermal chamber.

The default values for the Object Parameters are shown below in Figure 3:

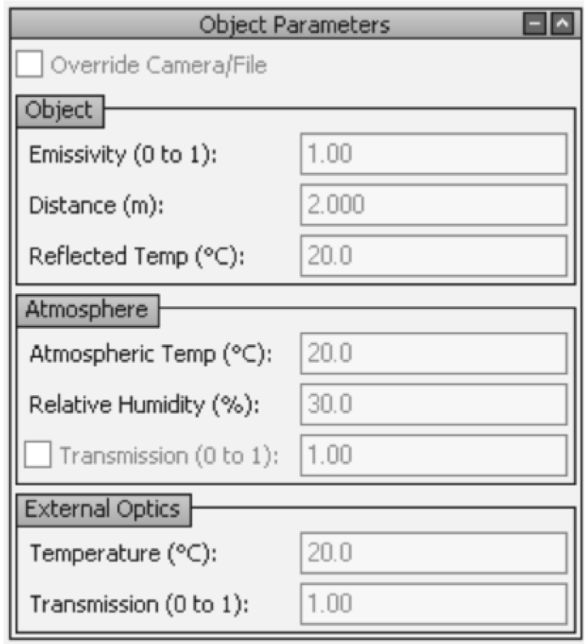


**Figure 3. Default values for ResearchIR Object Parameters. These inputs are applied whenever a factory calibration is active in the camera. The only parameter that applies to user calibrations is the distance.**

For a user calibration, the emissivity input has to be changed in the Stats window for each ROI. Changing it globally in the Object Parameter window will change the displayed value in the Stats window to match, but the value won't actually change! The user needs to use the checkbox to override the value in the Stats window, as shown in Figure 4.

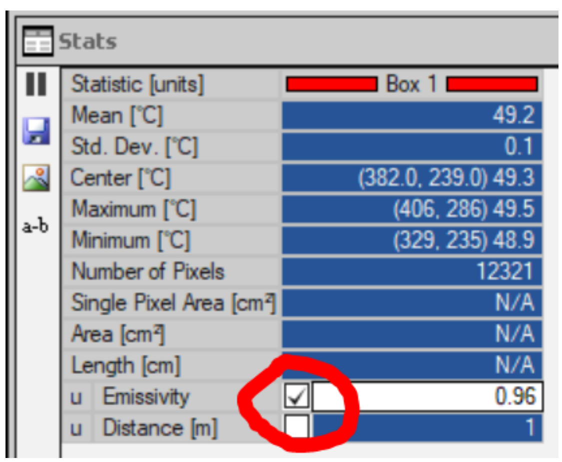


**Figure 4. When using a User Calibration, the emissivity for an ROI must be set locally in the Stats window by checking the circled checkbox and entering the appropriate value.**

Checking the Override Camera/File checkbox enables the camera system operator to change parameters to match current conditions. In the example screenshot below in Figure 5, the FLIR SC8303 camera has a factory calibration. The object parameters are based on the ambient air temperature, the distance from the camera to the hand and the known emissivity of human skin in the 3-5μm band. The mean temperature in the ROI is 26.7 °C. The software is computing an atmospheric transmission value of 0.97. This value is based on the distance, the atmospheric temperature, and the relative humidity. The inputs are compared to a library of model runs of MODTRAN, a radiative transport model designed by the Air Force Research Laboratory, and an in-band atmospheric transmission value is derived.

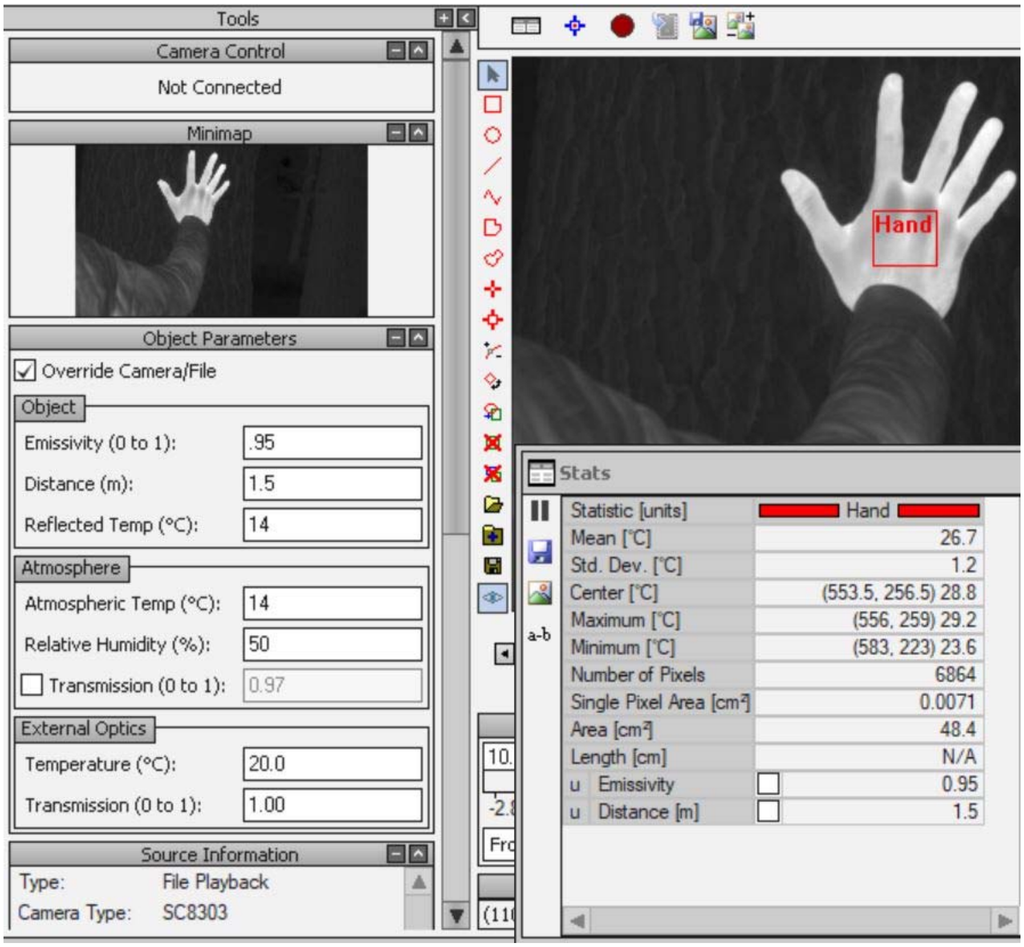


**Figure 5. Temperature of a hand against a tree, taken outdoors at night in cool weather. The Object Parameter values are optimized.**

Figure 6 is a screen shot of the same image and same ROI, but the Object Parameters are now the default values. The apparent temperature has dropped to 26.1 °C, mostly because the emissivity is set to 1, when it should be 0.95. The atmospheric transmission computes to 0.96, when it "should be" 0.97. This small change will make little difference. Note that I put should be in quotes because the atmospheric transmission is a calculated value, not a measured value. The calculation is derived from a library of spectral transmission curves generated by runs of MODTRAN, a radiative transport simulation originally developed by the Air Force Research Laboratory. The actual atmospheric transmission may vary from the simulation, which makes assumptions about the amount of $CO_2$ gas in the air path.

**Pro Tip: At a range of 1-2m, the air path effect on accuracy is small, but at longer ranges, the air path effect can become very pronounced.**

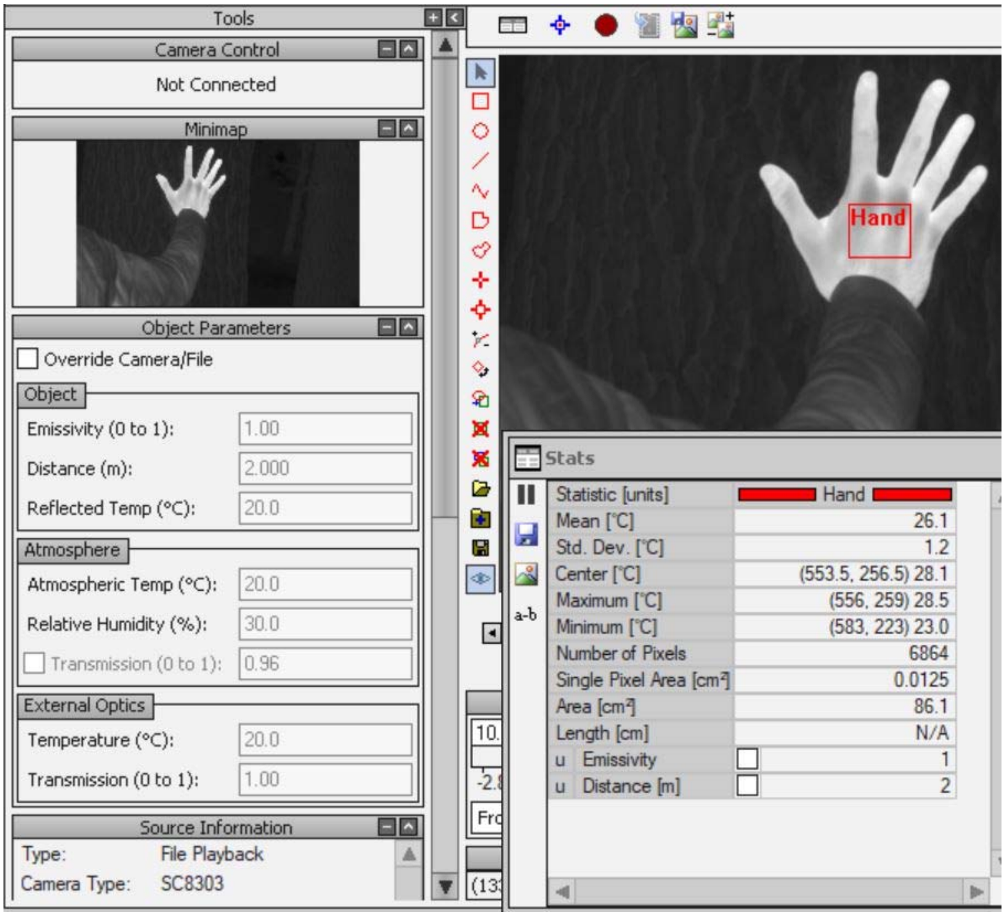


**Figure 6. Temperature of a hand against a tree, taken outdoors at night in cool weather. The Object Parameter values are their default values, which results in a lower temperature measurement value, mainly due to the emissivity being set higher than it should for human skin.**

## Simulating Camera Performance with a Software Model

In the next sections, I calculate errors induced on temperature measurements due to incorrect entry of object parameters. These calculations were done using a proprietary FLIR radiometric model constructed in Microsoft Excel. The model starts by using a Visual Basic function module to calculate the in-band radiance of a target using the Planck function, which is generated for a particular target temperature, then the area under the resulting curve is numerically integrated over the desired passband, in this case from 3 to 5 microns in wavelength. This in-band radiance, denoted as $L_T$, is then placed into a radiometric equation for the radiance the camera sees, both from that target and the atmosphere. The radiance the camera sees is denoted S. The equation is solved back for $L_T$, and the temperature of the target is derived from $L_T$ using a lookup table. The lookup table is a set of temperature inputs and the resulting in-band radiance outputs. The table is used backwards: an in-band radiance is computed, and the temperature that yields that radiance is determined. Here is an excerpt from the table. For a 31 °C input, the 3-5µm in-band radiance is 2.162e-4 Watts/cm$^2$/sr. If that same value of radiance was calculated by the model, then the lookup table would find that 31 °C temperature value, as shown in Table 1 below.

| Temp C | Temp K | Perfect Blackbody In-Band 3-5um Power Radiance |
|---|---|---|
| 25.00 | 298.15 | 1.742E-04 |
| 26.00 | 299.15 | 1.807E-04 |
| 27.00 | 300.15 | 1.874E-04 |
| 28.00 | 301.15 | 1.943E-04 |
| 29.00 | 302.15 | 2.014E-04 |
| 30.00 | 303.15 | 2.087E-04 |
| 31.00 | 304.15 | 2.162E-04 |
| 32.00 | 305.15 | 2.240E-04 |
| 33.00 | 306.15 | 2.320E-04 |
| 34.00 | 307.15 | 2.402E-04 |
| 35.00 | 308.15 | 2.486E-04 |
| 36.00 | 309.15 | 2.573E-04 |
| 37.00 | 310.15 | 2.663E-04 |

**Table 1. Lookup table**

The camera equation takes this form:

$$S = T_A(eL_T + R_TL_{BG}) + A_AL_A$$

Where the variable are defined as follows:
S is the total radiance the camera sees at its optics aperture
$T_A$ is the atmospheric transmission
e is the target emissivity
$L_T$ is the radiance of a blackbody at the target KST
$R_T$ is the target reflectivity
$L_{BG}$ is the radiance of a blackbody at the background temperature
$A_A$ is the atmospheric absorptivity/emissivity
$L_A$ is the radiance of a blackbody at the air temperature

The model first computes S based on the above inputs, which are the "true" values. The radiance values of each component of S are computed using numeric integration under the Planck curve. The camera equation was then solved for $L_T$ using algebra. This equation is arranged so that the parameters like emissivity can be changed to study the dependency of $L_T$ on those parameters being entered into the software in error. The S value is retained from the equation above with the correct parameters as inputs. The equation solving for $L_T$ contains S (the camera signal), and the new entered parameters. The $L_T$ value that comes out will differ from the correct $L_T$ value if any of the input parameters are changed. A lookup table is used to convert this new incorrect $L_T$

value back to a temperature, and the discrepancy between the correct target KST and the derived temperature is returned. This model allows us to study the effect on the target temperature measurement of changing the emissivity by 1%, for example.

$$L_T = [S - A_A L_A - R_T T_A L_{BG}]/[e T_A]$$

The italics indicate that the parameter is computed from a second set of entered parameters. $L_T$ = $L_T$ when the parameters that are entered all have the same value as they had previously when S was computed. If any of the parameters are changed from their "correct" values, one can see the effect on the output temperature measurement.

## Emissivity Error

If the object under test has an emissivity that is not accurately known, then there is no way to accurately measure its temperature. The camera measures a radiance value. This value needs to be corrected for the object's emissivity and the reflected scene radiance. Emissivity is the parameter that creates the largest contribution to error in radiometric temperature measurement.

Example: A metal plate is painted with a black paint that has an emissivity of 95% and a temperature of 50 °C. A 3-5µm radiometric camera is pointed at the plate and measures a radiance value of 3.957e-4 Watt/cm$^2$/sr. The ceiling in the room is 23 °C, so there is a 5% reflection of that temperature. The operator of the camera is inexperienced, and mistakenly enters an emissivity of 1 into the measurement software. The software now "thinks" that the camera is measuring the radiance of a perfect blackbody, so it will not consider the reflected radiance, because if the emissivity is set to 1, then the reflectivity is 0 and there will be no reflected radiance. But in reality, the camera is getting 95% of the radiation that would have been emitted from a perfect blackbody at a temperature of 50 °C, plus the camera is also getting the reflected ceiling radiance (5% of the ceiling radiance). Under these conditions, the software will report 49 °C, a lower temperature than the "true" temperature of the object. If the operator puts in 0.9 for the emissivity, then the reported object temperature will be 51 °C. Thus, a 5% error in emissivity results in a 1 °C temperature measurement error. At lower object temperatures, the error caused by a 5% emissivity error increases. For a 0 °C object with e = 0.95 and a 23 °C reflected scene temperature, putting in e = 0.9 gives a -2C measurement, and putting in e = 1 gives a 1.65 °C measurement.

### Cold Targets and Emissivity Error

Errors in emissivity have a much stronger effect on temperature measurement accuracy for very low temperature targets. Using the Niceville standard calibrations, the coldest target that an InSb camera can measure is -20 °C. At that target temperature, the radiance of just about everything else in the system is much higher than the target radiance. Consider the same scenario as above, where the reflected scene temperature is 23 °C, the true in-band emissivity is 0.95 and the target is -20 °C.  For an entered emissivity of 0.9, the measured temperature is -26.95 °C.  An emissivity of 1 gives -15.15 °C.  A small error in emissivity leads to a large error in temperature measurement.

### Hot Targets and Emissivity Error

For a target KST of 1000 °C, an emissivity of 0.95 and a reflected temperature of 23 °C, the effect of changing the emissivity by 0.1 changes the temperature by ~30 °C.  For example, if the operator enters 0.9 for the emissivity, the temperature measurement is 1022 °C.  If e=1 is entered, the measurement is 980 °C.  This error is directly related to the corrected radiance, not the reflected radiance, which for a 23 °C reflected scene temperature is a tiny fraction of the target radiance.

### Pathological Emissivity Case

In the case of the target being the same temperature as the reflected scene, the measurement will be unaffected by the setting of the emissivity.  It makes sense because in that case, the radiance that the camera sees will always be equivalent to a blackbody at 23 °C, and the actual surface emissivity will modify how that radiance is apportioned between the emitted and the reflected components.

## Reflected Temperature Error

Let's say that the camera operator knows that the in-band emissivity of the surface is 0.95.  That means that the surface reflects 5% of the in-band radiation that impinges on it.  That radiance comes from the ceiling of the room, or the sky, or a hot machine above the object, or any number of things.  If the radiance of the reflected objects is not accurately known, then the correction for it in the software will not be accurate and this will cause a temperature measurement error.  The magnitude of the error depends on the situation. If the emissivity is really close to 1, then it is hard to affect what the camera sees with reflections.  If the target is very hot, with a radiance that is 1000 times the reflected scene radiance (for example), then the measurement will again be very insensitive to changes in the reflected scene temperature.

Take the previous example of the 50 °C plate with an emissivity of 95% and a 23 °C reflected scene temperature.  Suppose the operator correctly enters 0.95 as the emissivity, but instead of putting in the correct ceiling temperature of 23 °C, he mistypes, and accidently puts in 33 °C.  What happens to the apparent temperature? It will be 49.7 °C, which is a very small error.  Under these conditions, the apparent temperature is very insensitive to the reflected temperature value.  If the plate was at 0 °C, then the error in the reflected radiance becomes a much larger percentage of the total radiance.  In that case, putting in 33 °C for the ceiling temperature results in an apparent temperature of -1.4 °C, which is almost 5 times larger of an error than in the 50 °C case.  Here is a rule of thumb and its converse:

***When the target of interest is cold relative to the temperature of the reflected scene, the measurement will be much more sensitive to errors in the reflected-scene temperature input, particularly when the target emissivity is close to unity.***

***When the target is hot relative to the temperature of the reflected scene, the measurement will be insensitive to errors in the reflected-scene temperature, particularly when the target emissivity is close to unity.***

## Additional Sources of Error

As if emissivity and reflected temperature-induced errors weren't bad enough, there are additional effects that can also introduce significant errors. They include:

1. The atmospheric transmission between the camera and the target, especially at ranges beyond tens of meters
2. The air temperature, because hot air emits more radiation than cold air
3. The parasitic radiation of the camera optics and camera interior
4. The size of the target in pixels, especially below 20x20 pixels

### Atmospheric Transmission

The radiance from the target is reduced by the atmospheric transmission. For a 50-meter air path, 20 °C air and 30% relative humidity, the in-band atmospheric transmission as calculated by ResearchIR is 86%. The radiance from the target (which includes both the emitted and reflected components) will thus be reduced by 14%. By conservation of energy, and assuming thermal equilibrium between the air and the electromagnetic radiation field, the air path also emits like a blackbody with a 14% emissivity. If the air is hot, then this air path radiance (as it is called) can introduce a significant error, especially if the target is much colder than the air. The camera detects radiance from the target (both emitted and reflected), as well as radiance from the air path itself. The camera cannot tell the difference between those different components of the total radiance, so the operator has to tell the software the correct object parameter values in order for the software to correctly derive $L_T$, and from that derive a temperature measurement. Luckily, the air path radiance is usually very small compared to the target radiance, especially when the target is >100 °C and has an emissivity greater than 0.9. But if the target is cold and the air path is warm, and the atmospheric transmission is low, the measurement error can get severe! In any event, the radiance that reaches the camera optics aperture from the target is directly proportional to the atmospheric transmission, so entering an incorrect atmospheric transmission directly affects the corrected radiance value. The corrected radiance value is the measured radiance, minus the air path radiance, all divided by the atmospheric transmission. The temperature is then derived from the corrected radiance value. If the correction to the radiance value is not done properly, the apparent temperature will be in error!

For example, suppose the target is a 50 °C blackbody with an emissivity of 1. The air path temperature is 23 °C. The atmospheric transmission is 0.95. If the transmission is set to 1 in the software, then the apparent temperature measurement is 49 °C. If the transmission is set to 0.9, the apparent temperature measurement is 51 °C. A 5% error in emissivity translates to a 1 °C temperature error. It is relatively easy to measure in-band emissivity to within +/-1%, so for this scenario, the error due to emissivity can be minimized.

Notice anything here? I am using the exact same parameter values as I did when talking about emissivity. The effect is identical – both atmospheric transmission and emissivity act the same – the target radiance is linearly scaled by them. The air path temperature and the reflected scene temperature both add a radiance offset, the former when the transmission of the air is not 1, and the latter when the emissivity of the target is not 1.

## Radiance Measurements

Some FLIR science camera operators want to measure radiance of targets, as well as the radiant intensity, and they want to know what level of measurement accuracy can be achieved. The typical application for this sort of work is to measure the radiant intensity signature of a target like an aircraft as a function of distance. This might be done to see how the radiance drops off to a level where it would be undetectable to a particular sensor.

The measurement accuracy of radiance is also affected by the object parameters, but in a different way. For example, reflected radiance is part of the total radiance measured by the camera, so there is no attempt to correct for it. The radiance signature is a combination of emitted and reflected radiance. Since we are not trying to measure temperature, there is no need to correct for the reflected radiance component to try to back out the emitted radiance.

## Source Radiance and Apparent Radiance

In the case of atmospheric transmission, we note that the measurement that is affected is the measurement of **source radiance**. Source radiance is the “intrinsic” radiance of a target with no atmospheric effects. It is the measurement you would get in a vacuum, or at zero distance so there is negligible air path transmission loss. **Apparent radiance** is the radiance one measures at the aperture of the camera optics: it is what it is, and it is affected by air path transmission and air path radiance.

If the goal is to measure the apparent radiance on a given day with a given set of conditions, then knowing the actual air path transmission is not critical. An example of this might be a measurement of a hovering helicopter from 500m away. The camera operator wants to know how much radiance will reach a sensor aperture at that range under different sets of atmospheric conditions. Thus, the experiment might be conducted in different seasons and times of day, with the atmospheric conditions recorded. If the goal was to measure the helicopter source radiance, then it would be important to actually measure the air path transmission and air path radiance, so that the apparent radiance could be corrected back into a source radiance value.

## Relative Radiance

In many situations, the goal of the radiance measurement is not to get at absolute radiance, but to see the effect on radiance of some change in the target. For example, a hovering helicopter might be imaged from 200m away and during the measurement, do a slow 360-degree spin in place. The camera system will then be able to measure the radiant intensity of the helicopter as a function of view angle relative to the aircraft fuselage long axis, for example. The relative radiance is what matters in this case, and since the air path is the same or nearly the same during the maneuver, its value divides out.

## Accuracy at Performance Limits

This section is about measurement at performance limits: what happens to radiometric accuracy in the edge and corner cases? These cases could include one or more of the following features:

1. Very short integration times, where the performance of the readout IC in the camera focal plane array starts to degrade.
2. Very small targets, where optical effects start to degrade the sampling of the target.
3. Very cold targets, where the IR radiation emitted from the target is itself small compared to reflected radiation off the target surface, and/or the optics self-radiation in the system
4. Low atmospheric transmission, where the target radiance is reduced through a long air path and air path radiance can become significant.

## Very Short Integration Time:

FLIR Niceville runs factory calibrations of cameras in a range of integration times that vary from ~10ms down to a minimum integration time of about 10µs. The reason for this lower limit is that the FLIR ROICs can exhibit performance anomalies below this integration time, most noticeably the reduction of the counts span of linear performance. For instance, an SC8303 camera with a FLIR ISC0802 ROIC was operated at 0.48 microseconds and radiometrically calibrated with a cavity blackbody. A plot of in-band power radiance versus digital counts that resulted from this calibration is shown in Figure 7 below:

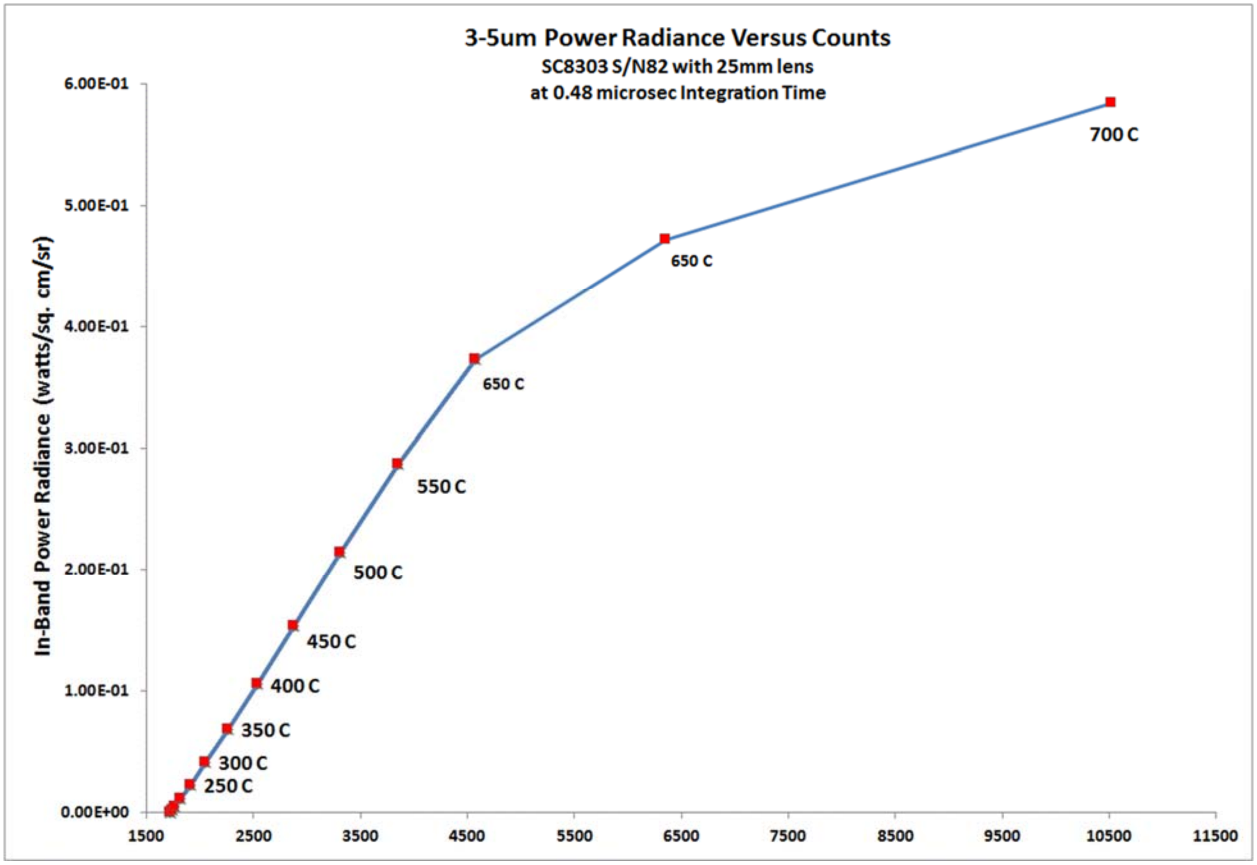


**Figure 7. Radiometric transfer function for a FLIR SC8303 operated at 0.5 microseconds integration time. The linear region of the sensor under these conditions extends up to about 550C.**

At this very short integration time, the maximum digital counts out of the camera in the **linear range** were 3850 counts, which corresponded to a cavity blackbody temperature of 550 °C. The camera counts continued to rise for higher temperatures, but the radiometric transfer function was no longer linear. When one is no longer in the linear region of the sensor, the non-uniformity correction no longer corrects the image, since the NUCs used by FLIR science cameras are linear transformations that only work well in the linear range of the sensor. The minimum target temperature that still produces an acceptable image is ~200 °C, as shown in Figure 8. The image quality was very noisy for target temperatures below 200 °C, as shown below in Figure 9.

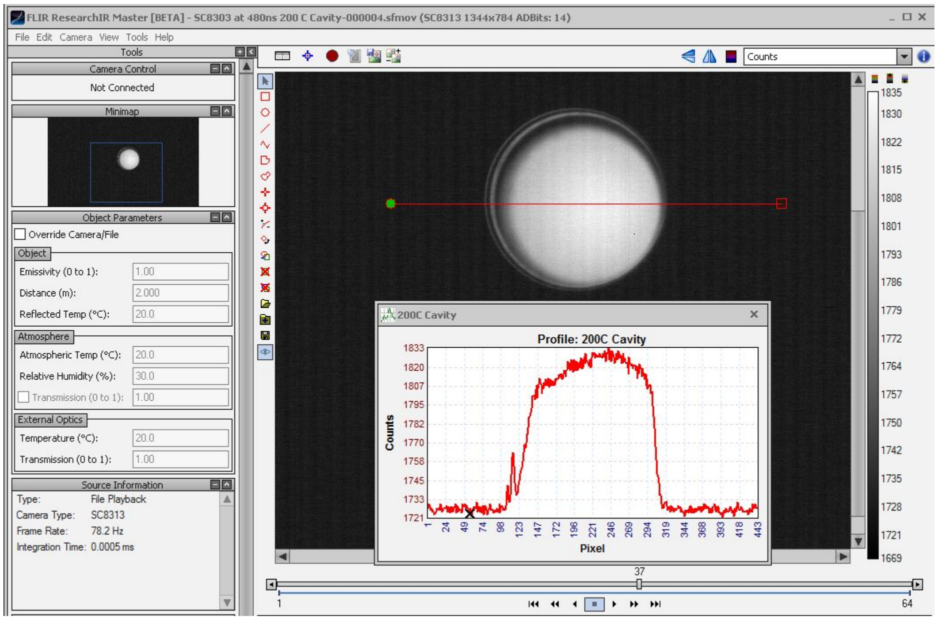


**Figure 8. Image of 200 °C cavity and line profile. There is a signal swing of about 100 counts for this target temperature. Given the noise level of the image, this is about as cold as one would want to try to measure a temperature.**

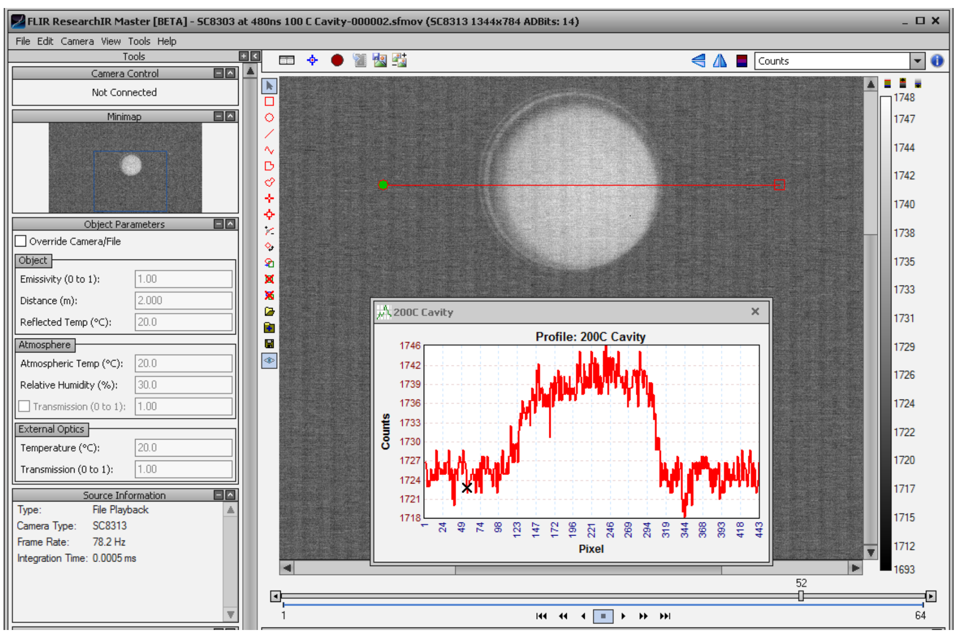


**Figure 9. Image of 100 °C cavity and line profile. The image is much noisier, and measurement of a target at 100 °C would be ill-advised due to the uncertainty caused by the noise.**

At 200 °C, the signal swing above ambient is only about 100 counts. At 100 °C, it is 20 counts and the images would be unsuitable for radiometric measurements. This particular camera configuration would only be useful for measurements between 200 °C and 550 °C. A neutral density filter could enable higher temperature measurements. If an ND1 filter were added, then the minimum measurement temperature would become ~410 °C and the high end would be ~1350 °C. These new temperatures were estimated by calculating the in-band radiance values for 200 °C and 550 °C (the limits without an ND filter), multiplying these radiance values by 10 and estimating the temperatures corresponding to these new 10X radiance values using a lookup table shown in Figure 10.

| SBIR Cavity Temperature (Celsius) | Temperature (Kelvin) | 3-5 micron In-Band Power RADIANCE (Watts/sq.cm/sr) |
|---|---|---|
| 23 | 296.15 | 1.62E-04 |
| 50 | 323.15 | 4.08E-04 |
| 100 | 373.15 | 1.62E-03 |
| 150 | 423.15 | 4.76E-03 |
| 200 | 473.15 | 1.13E-02 |
| 250 | 523.15 | 2.29E-02 |
| 300 | 573.15 | 4.13E-02 |
| 350 | 623.15 | 6.84E-02 |
| 410 | 683.15 | 1.14E-01 |
| 450 | 723.15 | 1.54E-01 |
| 500 | 773.15 | 2.14E-01 |
| 550 | 823.15 | 2.87E-01 |
| 600 | 873.15 | 3.73E-01 |
| 650 | 923.15 | 4.73E-01 |
| 700 | 973.15 | 5.85E-01 |
| 750 | 1023.15 | 7.10E-01 |
| 800 | 1073.15 | 8.48E-01 |
| 850 | 1123.15 | 9.98E-01 |
| 900 | 1173.15 | 1.16E+00 |
| 950 | 1223.15 | 1.33E+00 |
| 1000 | 1273.15 | 1.52E+00 |
| 1050 | 1323.15 | 1.71E+00 |
| 1100 | 1373.15 | 1.91E+00 |
| 1150 | 1423.15 | 2.12E+00 |
| 1200 | 1473.15 | 2.35E+00 |
| 1250 | 1523.15 | 2.57E+00 |
| 1300 | 1573.15 | 2.81E+00 |
| 1320 | 1593.15 | 2.91E+00 |

**Figure 10. The orange highlighted rows show the minimum temperature measurement values with and without an ND1 filter. Note the in-band radiance values differ by a factor of ten, which corresponds to the radiance reduction of 10X from an HD1 filter insertion. The red highlighted values are the high end of the temperature ranges with and without an ND1 filter.**

## Radiometric Measurement of Small Targets

It has been observed that when a target gets below a minimum spot size, measurement accuracy starts to degrade. The most common example of this is a hot target that is small and round on a cold background. When the diameter of the target gets below about 10 pixels, there is a noticeable decrease in the radiance of the target, which leads to an underestimation of its temperature. The effect is caused by the modulation transfer function (MTF) of camera systems, which blurs the edges of targets. It is a combination of diffraction and aberrations in the lens, as well as detector footprint MTF due to the non-zero size of the detectors themselves. Figures 11a-d show images of a cavity blackbody with an aperture wheel in front of the cavity exit port. The wheel contains a set of apertures that allow the user to present a target to the camera that has uniform radiance across an emitting circle, and then to vary the diameter of the circle. As the emitting circle gets smaller, the MTF effects become more noticeable. Two things happen: the target no longer looks like a uniformly emitting circle, and the maximum radiance of the target decreases.

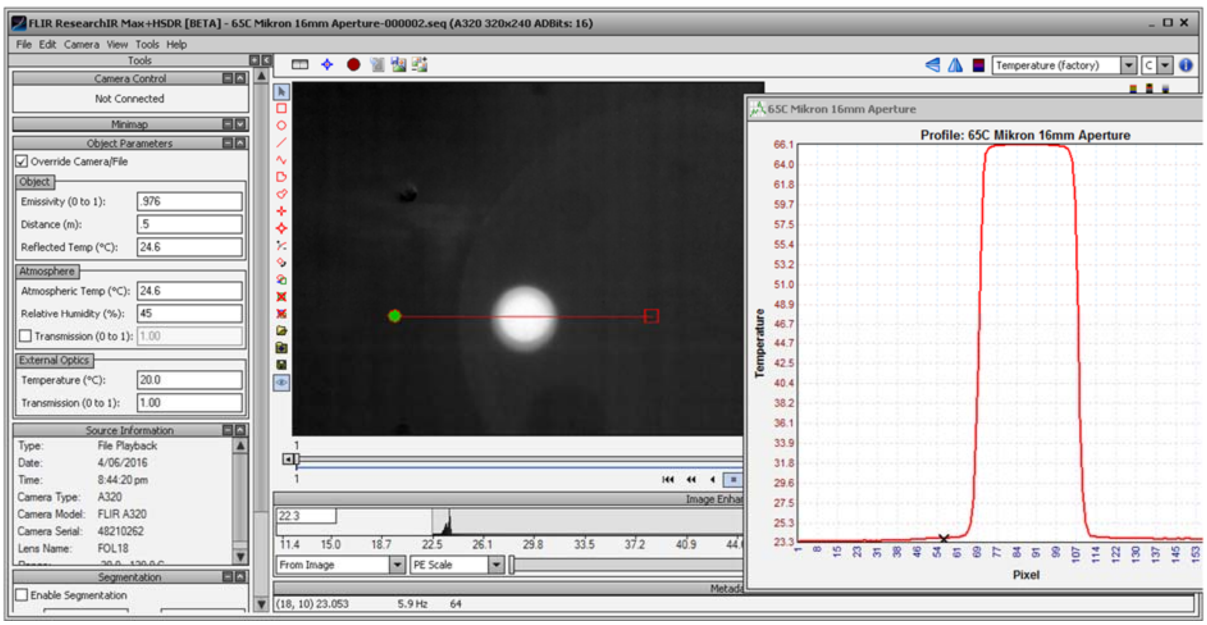


**Figure 11a. A 65C circular target that is 38 pixels in diameter. The temperature profile is flat across the top and has a peak value of 66.1 °C. There are sloping sides to the line profile peak, but the slopes do not interfere with the flat peak except near the edges of the circle.**

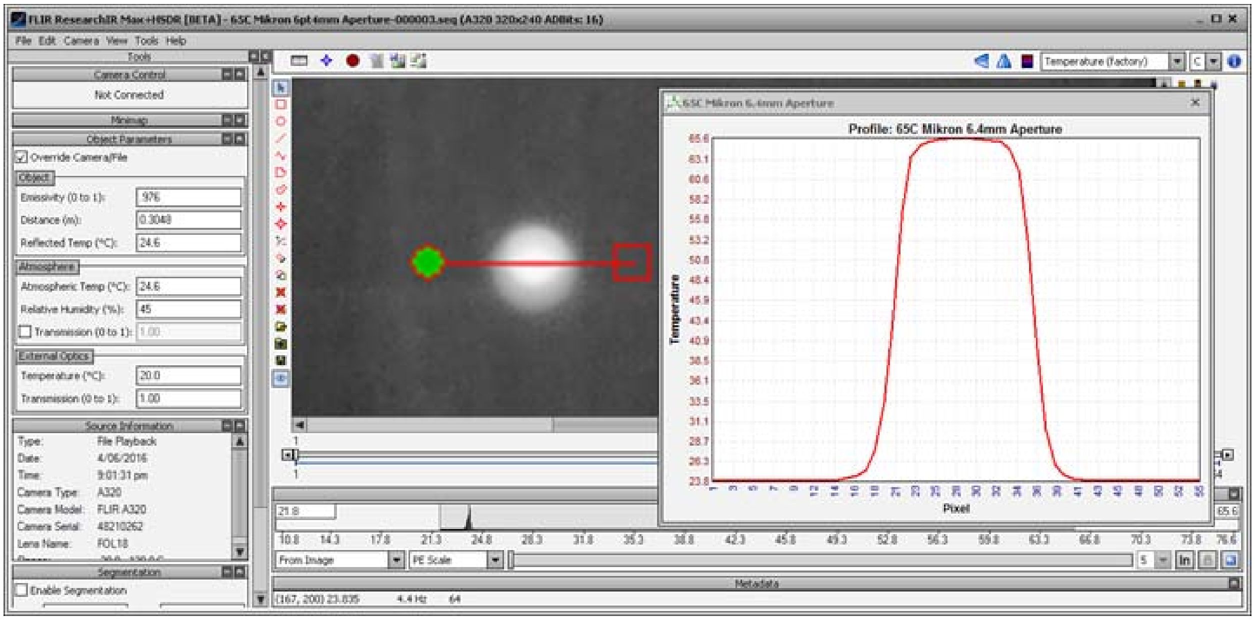

**Figure 11b. The 65 °C circular target is 15 pixels in diameter. The temperature profile is no longer flat across the top except for a few pixels, and the peak value has dropped to 65.6 °C. The sloping sides of the line profile are more noticeable since the horizontal scale is zoomed in compared to Figure 11a.**

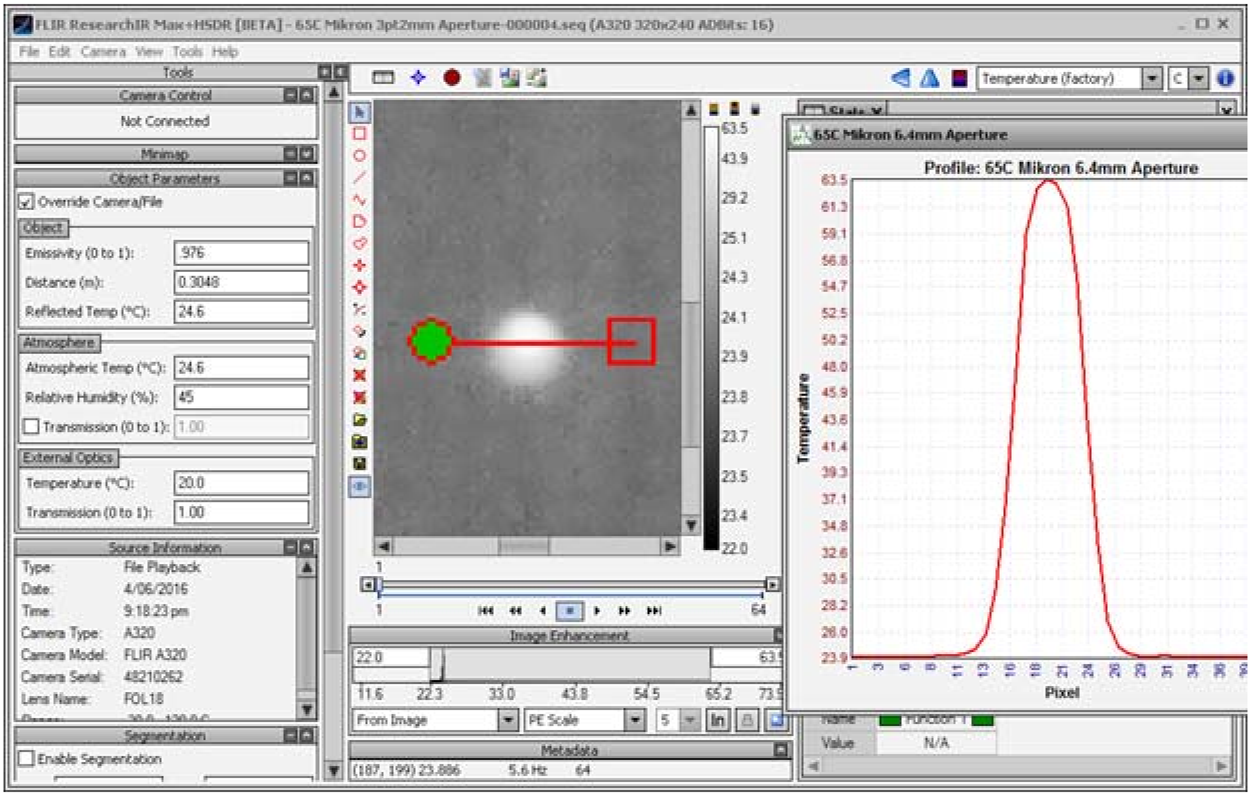

**Figure 11c. The 65 °C circular target is 7.5 pixels in diameter. The temperature profile is no longer flat across the top at all, and the peak value has dropped to 63.5 °C.**

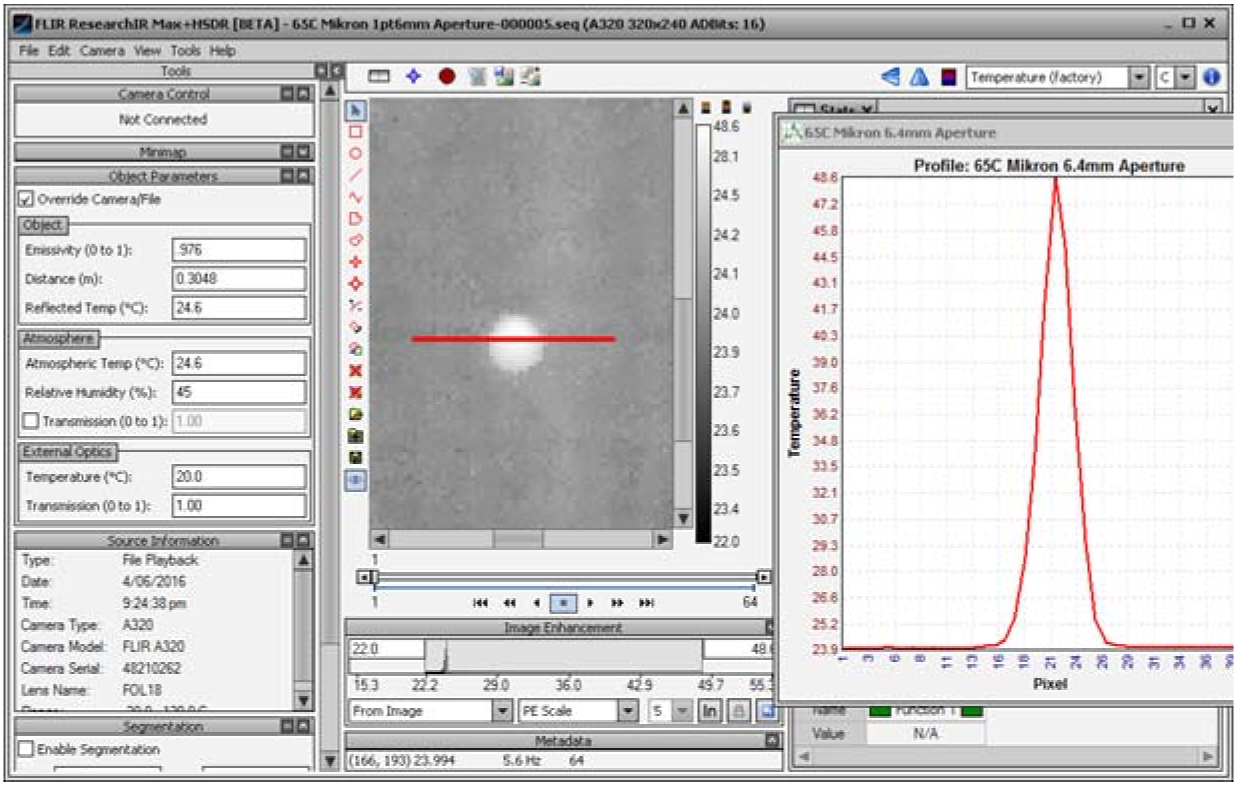


**Figure 11d. The 65 °C circular target is ~3 pixels in diameter. The temperature profile is spiky, and the peak value has dropped 48.6 °C, a very significant error.**

Conventional wisdom in the infrared camera industry is that a target can be as small as 3x3 pixels to get an accurate temperature measurement on it. This notion would appear to be quite wrong. What is true is that a 3x3 spotmeter can be used to obtain accurate temperature measurements on targets, but if the target is too small, the measurement will be incorrect. What is "too small"? It would appear to be any high-contrast target that is smaller than about ~15 pixels across in the smallest dimension. By high contrast, I mean that the target is much hotter or colder than its surroundings. When that is the case, it is as if the cold or hot pixels that surround the target get "mixed in" to the pixels subtended by the target itself.

This experiment was repeated with additional uncooled longwave cameras looking at the same target, with the results shown in Figure 12. The cameras all have different IFOVs. The data all line up well and show that below about 15 pixels, there is a drop-off of apparent temperature of a 100 °C blackbody target. Below 5 pixels, there is a very sudden drop off.

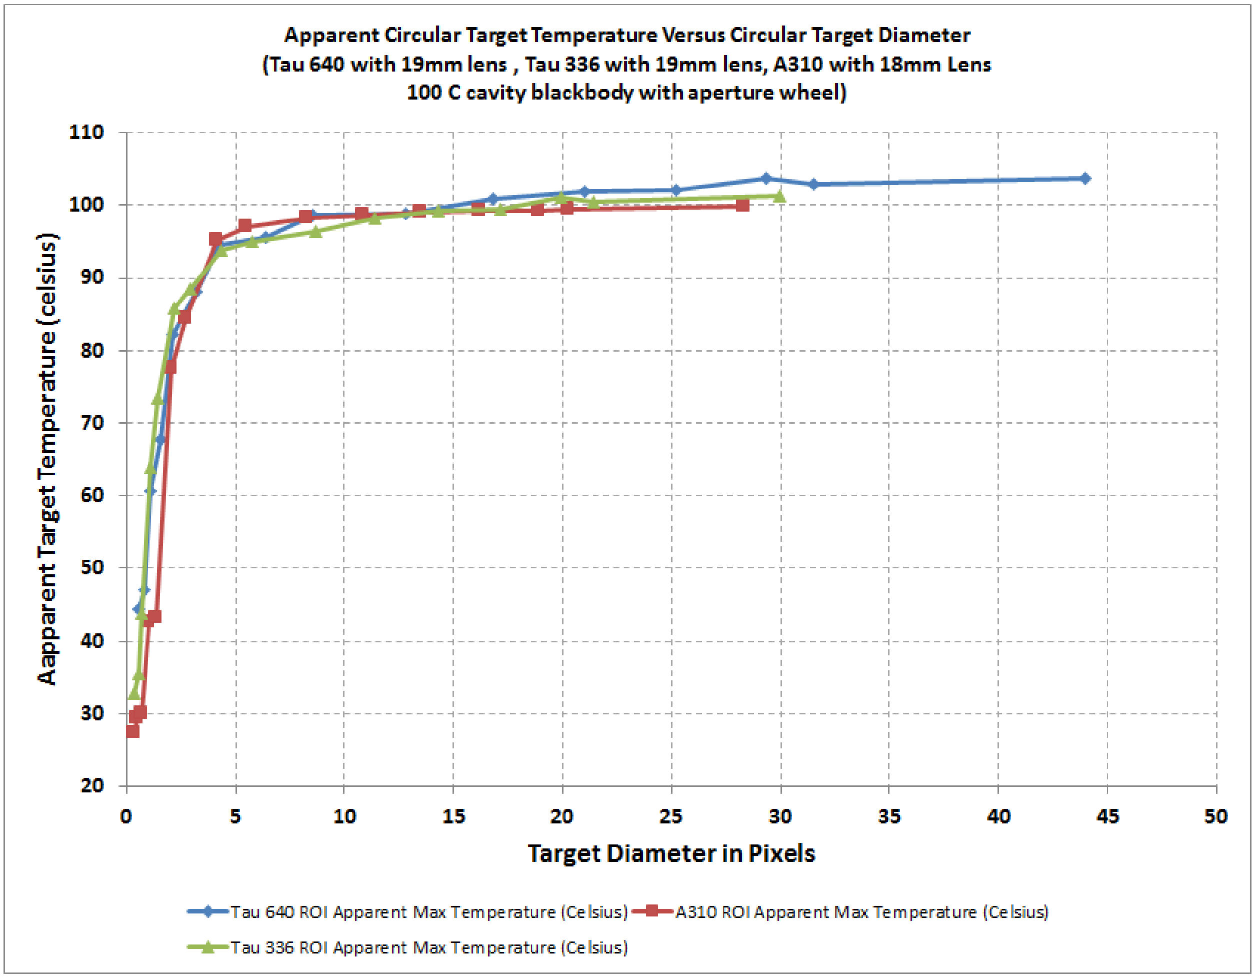


**Figure 12. Apparent target temperature versus target diameter in pixel units.**

The spot size effect is something that needs to be taken into consideration when trying to obtain accurate radiometric measurements of small targets. Targets need to be much bigger than what is commonly believed in order to achieve good temperature accuracy.

## Radiant Intensity of Small Targets

If the goal of the measurement is to determine the apparent radiant intensity contrast of the target, then the size of the target in the image is not as much of an issue, since the purpose is to image the target at range, not to correct back to the source radiance. In this situation, a longer lens is an advantage, because the longer the lens focal length, the better the sampling of the target image. Apparent radiant intensity contrast is the radiant intensity difference between a target and its background at the camera lens aperture, with no correction for atmospheric transmission.

One application for this type of measurement is to determine a camera system's ability to detect and track aircraft at long range. A typical jet aircraft at tens of km range or longer will most likely present as a point source, i.e. it will subtend an angle that is several times smaller than the

IFOV of the camera. In this case, the aircraft will look like a circular blob of light that covers a small number of pixels. Midwave IR cameras will never image a point source onto a single pixel because of the point spread function of the optics, even with a low f/number coldshield. At closer ranges, the aircraft geometry can be resolved. Figure 13 shows an out of focus image of what is likely a medium-sized commercial jetliner flying in the pattern around Ronald Reagan Airport in Washington DC. The image shows the two engines on the plane, which appears to be flying away from the camera.

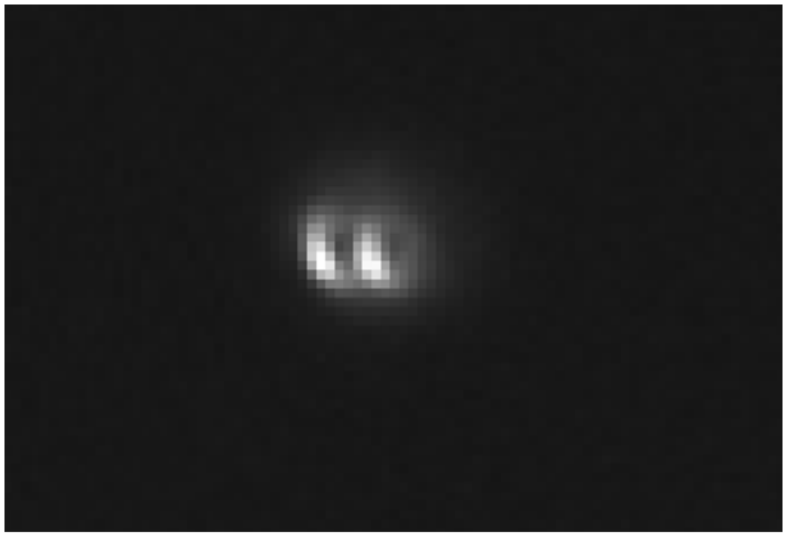

**Figure 13. An out of focus 3-5µm image of a twin-engine jet aircraft at an estimated 1200m range**

## Radiometry on Very Cold Targets:

For the purposes of infrared imaging radiometry, a very cold target is one that is -20 °C or colder. The in-band radiance in the 3-5µm band is very low for this target temperature, and it is very difficult to make an accurate temperature measurement for colder targets with a standard midwave camera. There is just too much parasitic radiation from the optics that competes with the weak signal from the target. The optics will typically be at 20-40 °C, depending on the ambient environment. A blackbody at 40 °C has ~11 times more in-band radiation than a target at -20 °C. If we treat the optics as a graybody having an emissivity of 0.1, then the radiation from the optics will equal the radiation from the cold target. For a -40 °C target, the ratio is even worse – there is 33 times more radiation from a 40 °C blackbody. Even slight changes in the lens temperature will have a significant effect on measurements of the cold target. Many of the FLIR Niceville cameras have $T_{Drift}$ correction that will counteract the effects of the parasitic radiation, but there are limits on its effectiveness.

A much better operational waveband for very cold temperature target measurements is the LWIR. The Niceville SLS cameras are designed to operate in the 7.5-11µm waveband, where there is a substantial increase in radiance compared to in the 3-5µm band. For a -20 °C target, the SLS camera will see **84** times more radiance than a 3-5µm camera. At -40 °C, the ratio is

**144**! So right off the bat, there is *much* more signal to work with. Secondly, the parasitic radiance of the ambient temperature optics is not substantially more that the target radiance. To use the previous example, a blackbody at 40 °C has only 3.3 times more radiance than a -20 °C target in the SLS band, and for a -40 °C target the ratio is still only 5.7. This is due to blackbody physics and the extent to which the in-band radiance changes with temperature in the different bands. We did a user calibration on an SLS camera using a CI System SR-800R blackbody which can be set to a temperature of -50 °C. The blackbody surface, shown in Figure 14a, was purged with dry nitrogen gas, and a tube-shaped plastic bag was installed between the camera lens and the blackbody to capture the purged gas and exclude moist air, as shown in Figure 14b.

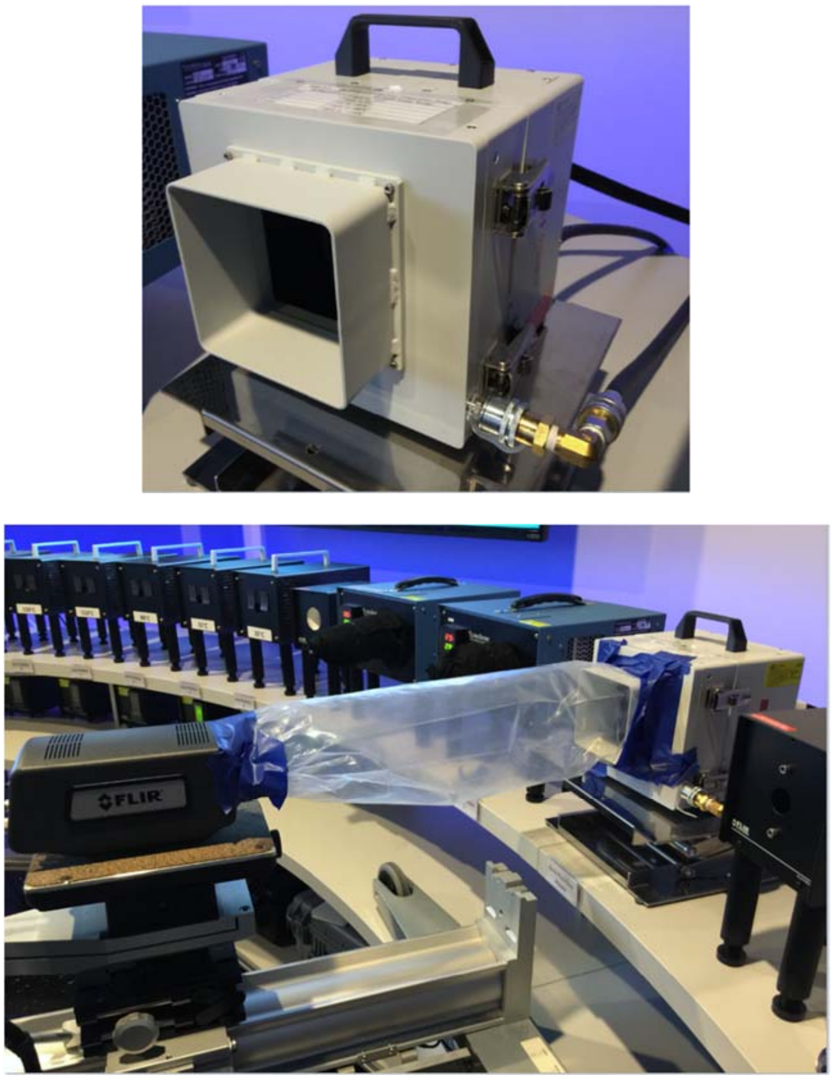


**Figures 14a-b. CI Systems SR-800R blackbody (a) and camera and blackbody setup with purged bag tube to exclude moist air (b).**

The blackbody got down to an indicated temperature of -48.3 °C, probably limited by the relative warmth of the dry nitrogen purge gas.  An improved setup would involve pre-cooling the purge gas with a coiled-up copper gas line immersed in a closed-circuit chiller bath.  The SR-800R emitter emissivity is 0.98 in the LWIR band, so that value of emissivity was used for the ROI in Figure 15 below. The measured value is within 2 °C of the indicated value, which is within the measurement specifications of standard calibrations.

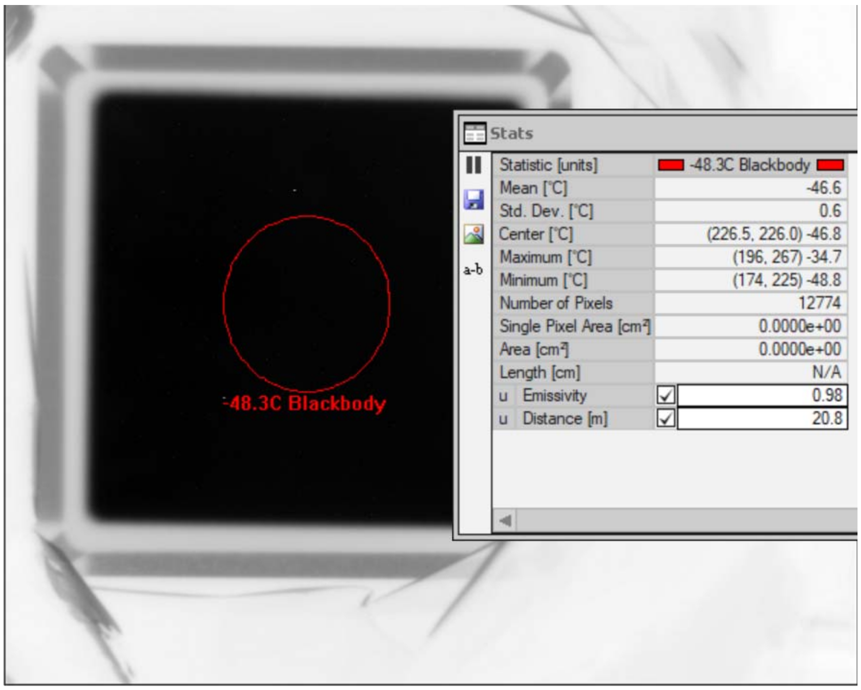


**Figure 15.  Apparent temperature image of -48.3 °C area blackbody taken with a FLIR A6751sc SLS camera with a user calibration**

## Atmospheric Transmission

The optical transmission of the air path between the target and the camera measuring it can vary from 100% transmission, which is realized at very close up distances, and essentially 0% transmission, which occurs when the air path is sufficiently long, and the air is thick.  Thick air is encountered at low altitudes, and horizontal air paths are the worst case for preserving radiometric accuracy.

### Rocket Fuel Tank Measurement

A particularly difficult measurement to make is the temperature of a cold object imaged at considerable distance through a warm, humid air path.  An example of this is an MWIR measurement of a liquid hydrogen fuel tank outer surface on the Space Shuttle.  I don't have the actual data on this scenario, but I can make some guesses for model inputs and see just how fraught with peril this sort of measurement is.

**Note: In this example, I am assuming that the $T_{Drift}$ correction is being made to the digital data perfectly. In other words, the parasitic radiation inside the camera lens and lens mounts is compensated. That is not necessarily the case unless the camera and lens system are very stable in temperature.**

Figure 16 shows the original model inputs, which are the "true" values. I will then change the inputs and see what happens to the temperature measurement. This is all a simulation, but the results should apply to an actual camera measurement.

1. The KST of the tank outer surface is -20 °C
2. The 3-5µm in-band emissivity is 0.8
3. The reflected background temperature is 35 °C (Central Florida in the summer)
4. The atmospheric temperature is 35 °C as well
5. The atmospheric transmission at the distance between the camera and the tank is 0.67

The atmospheric transmission was calculated using MODTRAN software for a 200m air path with air temperature of 35 °C and 80% relative humidity. The air is very "soupy" under these conditions. At 1000m range, the air path transmission would drop to ~13%.

If all the Object Parameter inputs to the software are correct, then the temperature value that comes out would be -20 °C. Note the External Optics transmission is 1, because there is no external optics being used.

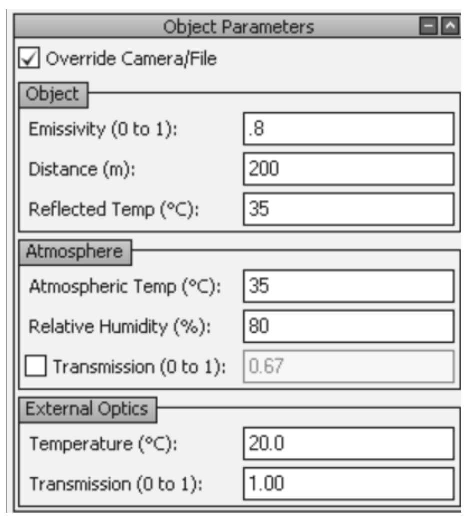


**Figure 16. Object Parameters for rocket fuel tank exterior temperature measurement**

Now let's change some of the parameters one at a time and look at the effect on the rocket tank temperature measurement.

If the emissivity is actually 0.8, but the camera operator thinks it is 0.85, then the measured temperature is -11.4 °C. That is a large error, since the actual temperature is -20 °C.

If the reflected temperature is actually 35 °C, but the camera operator enters 30 °C, then the temperature is -13.2 °C.

The worst source of error is the air temperature, since there is a significant amount of radiance from the air path itself relative to the radiance of the cold target. If the air temperature is 35 °C, but the operator enters 30 °C, the measurement is -5.75 °C, a whopping error to be sure!

Finally, if the air path is really 0.67, but the entered air path is 0.72, the temperature is -10.2 °C.

What happens if we put in a number of incorrect parameters all at the same time?

Table 2 shows the entered values in red, and the actual values:

| | Entered Values | Actual Values |
|---|---|---|
| Entered Target Emissivity: | 0.85 | 0.8 |
| Entered Background (Reflected-Scene) Temperature (Celsius): | 30 | 35 |
| Entered Atmospheric Temperature (Celsius): | 30 | 35 |
| Entered Atmospheric Transmission: | 0.670 | 0.670 |

**Table 2.**

The result is a temperature measurement of 1.5 °C, which is 21.5 °C warmer than the true temperature. This is the most difficult type of measurement in radiometry: cold object with low emissivity reflecting warm scene, imaged at long distance through warm air.

At longer ranges, the situation is even more hopeless. The estimated air path transmission at 1km is 0.13. Assuming that is correct, if the emissivity, reflected temperature and air temperature are just slightly incorrect, the resulting error is severe! In Table 3, I changed the inputs by 1% for emissivity and 1 °C for the Object Parameter temperatures, and the result was a measurement of 10.6 °C, which is 30.6 °C too high!

| | Entered Values | Actual Values |
|---|---|---|
| Entered Target Emissivity: | 0.81 | 0.8 |
| Entered Background (Reflected-Scene) Temperature (Celsius): | 34 | 35 |
| Entered Atmospheric Temperature (Celsius): | 34 | 35 |
| Entered Atmospheric Transmission: | 0.130 | 0.130 |

**Table 3.**

The reason is simple – the radiance that reaches the camera is a tiny fraction of the atmospheric radiance, which itself depends on the air temperature. If the emissivity is slightly off, then the reflected background radiance changes, and since it is twice the emitted radiance, that has a large effect on the outcome. The bar graph in Figure 17 shows the relative apportioning of the radiance components that go into S, the camera signal. In this example, there is no window between the camera lens and the target. Adding a window always makes accuracy worse, as it absorbs target radiance, emits its own radiance and reflects the radiance of the inside of the enclosure.

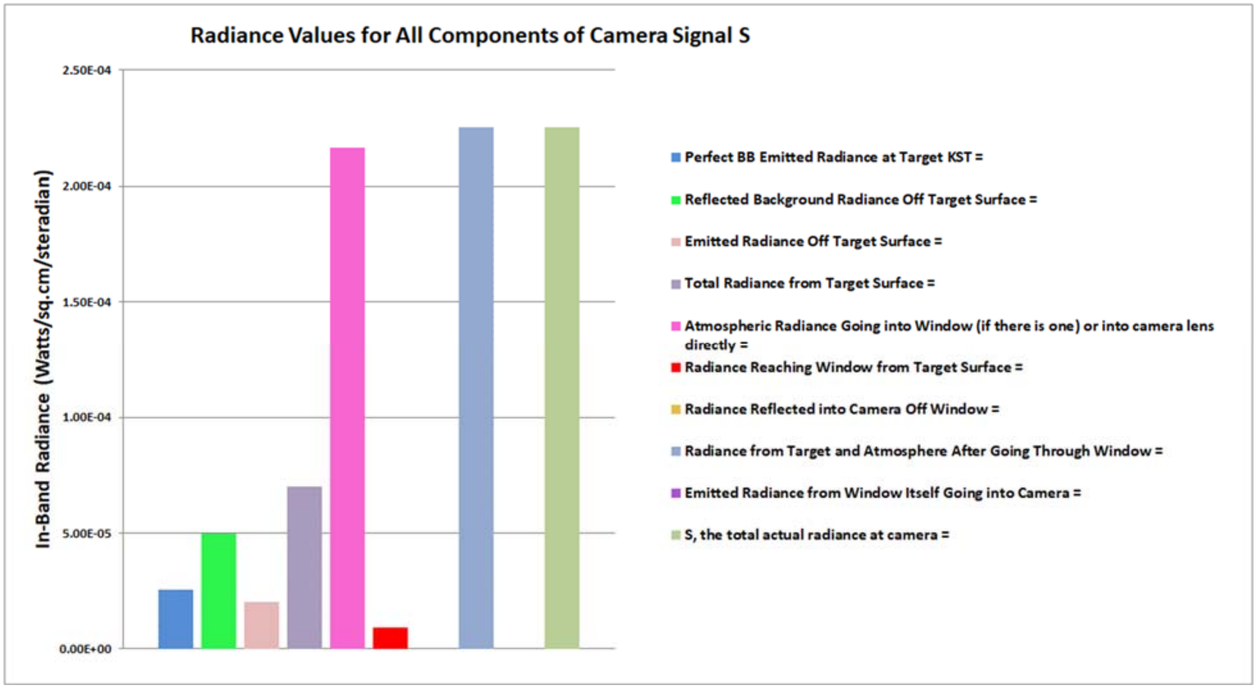


**Figure 17. Radiance values for rocket fuel tank measurement scenario**

## Summary

Radiometric accuracy is difficult to predict for real-world situations. FLIR Niceville measures accuracy by measuring the apparent temperature of periodically calibrated cavity blackbodies and comparing the results to temperature measurements of the same cavities made by a calibrated pyrometer. When a radiometrically-calibrated camera is pointed at a target, it will always measure an apparent temperature, but correcting that apparent temperature back to an accurate kinetic surface temperature is complicated by the effects of emissivity, reflection, atmospheric transmission and the extent to which the $T_{Drift}$ correction works properly, which in turns depends on some level of thermal equilibrium being established in the camera interior and optics in order that the camera system's prediction of the radiance offset determined by the internal temperature sensors is accurate. The Niceville cameras are calibrated in radiance and temperature, but we do not specify the accuracy of radiance measurements, in part because it is difficult to verify the absolute accuracy of a camera measurement of spectrally weighted radiance.

# Chapter 5 : Optical Filters in Radiometric Cameras

**Image of man with gas-power backpack leaf blower taken with midwave IR camera equipped with a spike filter that enhances hot carbon dioxide emissions. This is a frame from the film Racing Extinction, which used a FLIR SC8303 camera with a warm filter to image carbon dioxide gas emissions.**

# Contents

Chapter 5 : Optical Filters in Radiometric Cameras ........ 5—249
Introduction ........ 5—251
Bandpass Filters ........ 5—252
FLIR Warm Filters: ........ 5—252
Warm-Filtered Camera ........ 5—257
Case Study - Warm filter versus Cold Filter ........ 5—263
Warm-Filtered Methane Camera ........ 5—264
Cold-Filtered Methane Camera ........ 5—265
Hotter Scene Temperatures ........ 5—266
Non-Uniformity with No NUC applied: ........ 5—273
External Offset Update Time Drift ........ 5—274
Summary ........ 5—277

## Introduction

It is often desirable to restrict the spectral range of an infrared camera to a particular band to observe gas, plastic or other materials. These materials may be highly transparent to most wavelengths of infrared radiation yet will absorb strongly in certain narrow ranges of wavelength. The most common embodiment of this concept are the optical gas imaging cameras, built by FLIR Systems and others. These cameras use a narrowband cold filter to make the camera see in a very particular spectral range. The camera will not detect radiation outside the narrow band, making these cameras fairly insensitive to scene radiance. They are really suitable only for specific imaging applications, though they can be calibrated to measure scene temperatures. Figure 1 shows an image of a large propane leak taken with a hydrocarbon gas imaging camera. The camera's cold filter transmits in the 3.2-3.4-micron band only, where hydrocarbon gases show significant absorption due to a C-H bond stretch mode resonance.


**Figure 1. Propane leak imaged with FLIR GF320 hydrocarbon gas camera**

It is also possible to use bandpass filters that are not cold, i.e. not built into the refrigerated compartment where the sensor is located. A "warm filter" as it is called, can be installed in a filter wheel between the lens and the sensor, behind the lens in a static filter holder, or in front of the lens.[16] The warm filter concept is useful for applications where the object of interest has a fairly high radiance, i.e. the hotter the better. Neutral density filters are used to limit the radiation from the scene to something that the camera sensor can manage. For example, a typical midwave science camera can image blackbodies up to 350C or so. For higher scene temperatures, a filter is needed to restrict the levels of the incoming radiation. The most commonly used filter is a neutral density type, which has uniform transmission within its band of operation. In the case of an IR camera imaging up to 350C without a filter, the addition of an ND1 filter (which transmits $10^{-ND}$ or 0.1) would allow the camera to see a scene with a radiance that is 10 times higher than the radiance of a 350C blackbody. For the 3-5-micron midwave band, this 10X high radiance is reached for a temperature of 750C. It is never advisable to use a neutral density filter with a long

[16] The latter is the rarest configuration, because the filter should be bigger than the clear aperture of the front lens element, which can be quite large in diameter (>2 inches). Midwave filters this big are custom and expensive.

integration time to look at a cool scene.  The filter must be removed. Otherwise it is like wearing your sunglasses at night.[17]

## Bandpass Filters

These bandpass filters are constructed with multiple interference layers on an IR-transparent substrate. Figure 2 shows a typical midwave IR bandpass filter transmission curve for a filter suitable to image carbon dioxide gas, which absorbs strongly at 4.3 microns wavelength:

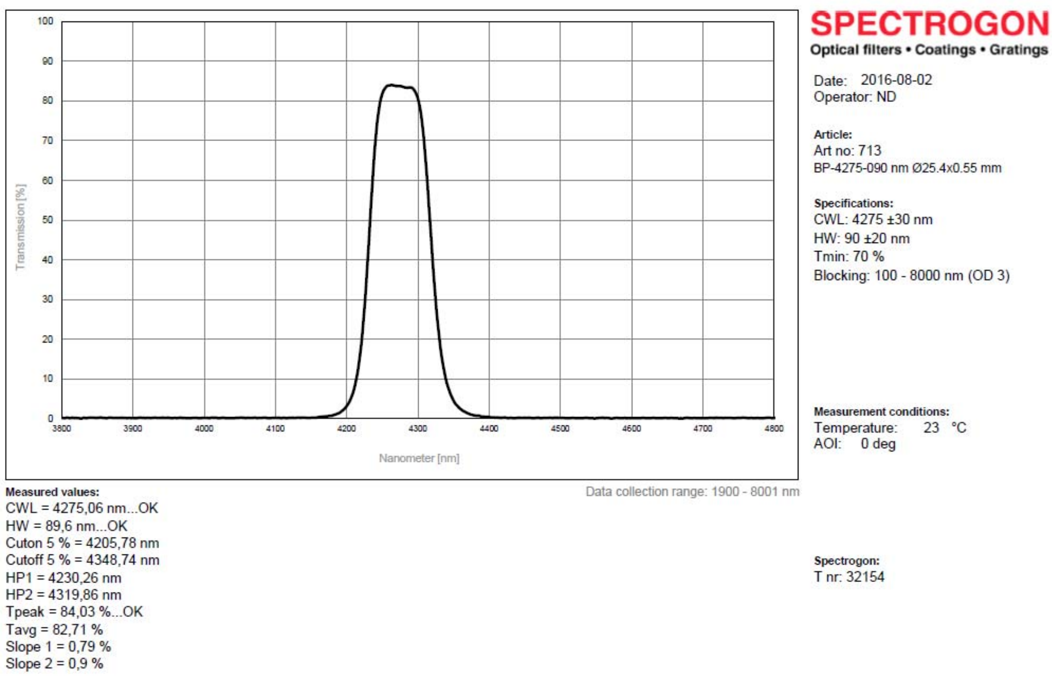


**Figure 2.  Typical interference-type filter for carbon dioxide gas imaging**

The CO2 filter has ~80% transmission in the bandpass and blocks radiation down to 0.1% transmission or less outside the main peak.  This type of filter can be built into the coldshield of an infrared camera, typically one with an InSb detector.  These filters have a center wavelength that shifts to smaller values as the filter temperature decreases.  The peak of the filter shown in Figure 2 is right on a strong absorption line of carbon dioxide gas when the filter is at 23 degrees C.  Taking it to 77K where cold filters are operated will shift this wavelength down about 50nm, a shift which should always be considered by the camera designer when cold filters are specified.

## FLIR Warm Filters:

FLIR has a number of warm filters that are sold for use with the Niceville science camera line. They are mounted in filter holders that screw into the back of lenses, or in different holders that fit into cameras equipped with filter wheels.  The filter wheels can hold 4 filters, and the wheel has a magnetic or optical sensor that can read codes off the filter holders.  This ensures that the

[17] This is not advisable unless you are a pop singer from the 1980's named Corey Hart.

camera "knows" what filter in in what position. This is particularly important for cameras that have factory calibrations that require a particular filter to be in the optical path, such as a neutral density filter for high temperature calibrations.

Here is a table of our current offerings for the Niceville science cameras:

| **Description** | **Center Wavelength and FWHM (nm)** | **Application** |
|---|---|---|
| ND1 | Broadband | 45-600C (MW), 250-1500C (LW) |
| ND2 | Broadband | 250-1500C (MW), 500-3000C (LW), 400-1200C (SW) |
| ND3 | Broadband | 500-3000C (MW), 500-3000C (LW), 1100-2200C (SW) |
| MWIR | 4000, 2000 | MWIR imaging (used on a broadband camera) |
| Thru Glass | 2360, 80 | Imaging through glass |
| Glass Surface | 5000, 145 | Glass surface becomes opaque for surface temp measure) |
| Flame | 3900, 200 | Sees through clean gas flames (no soot) |
| Plastic | 3440, 140 | Makes thin film plastic opaque |
| CO2 | 4300, 200 | Sees CO2 gas |
| Nitrous Oxide | 4500, 170 | Sees Nitrous Oxide gas |
| COS | 4270, 70 | Flame Imaging (makes clean flames brighter) |
| Thru Glass HT | 2360, 80 with ND2 | Through glass high temperature 600-2500C |
| Glass Surface HT | 5000, 145 with ND1 | Measure glass surface temperature (200-2500C) |
| Flame HT | 3900, 200 with ND1 | Measure high temps through clean flames (250-1500C) |

**Table 1. Currently available warm filters for Niceville Science cameras**

The thru glass filters are highlighted in blue, because they can only be used with a broadband InSb camera. The standard 3-5-micron camera's cold filter is incompatible with a 2.36-micron filter – no scene radiation will reach the sensor. Other spectral filters are available for special applications, and many of these filters can be built into cameras as the cold filter. These custom cold-filtered cameras must be approved by Niceville engineering, carry substantial extra cost and lead time, and are not guaranteed to meet factory specifications or customer expectations. In some cases, they can be calibrated with a custom calibration, which incurs extra cost lead time and has no guarantee of accuracy.

The filter holders for the X series of cameras have magnets that are installed at the factory in numbered positions as shown in Figure 3a. These give the filter holder a unique ID which can be read by the filter wheel to determine what filter is installed in a given wheel position.

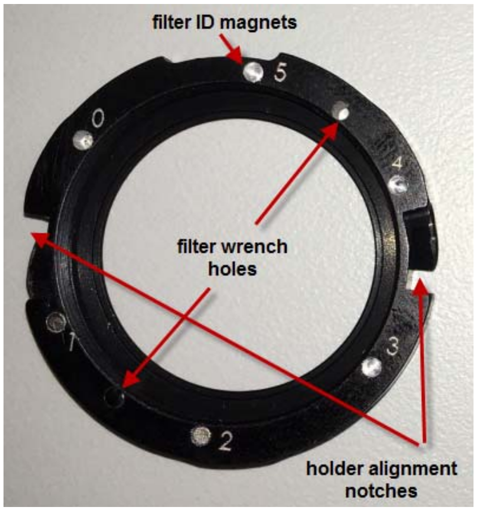


**Figure 3a. X series camera filter holder and filter wheel inside camera**

Other science cameras do not have a filter wheel. The warm filters are installed in a filter holder that screws into the back of lenses. The filters are held a few mm from the warm window, shown in Figure 3b. Filters are glued into filter holders by FLIR at the factory. Customers can buy empty filter holders from FLIR and mount their own 1-inch diameter filters in holders, as shown in Figure 3c.

**Figure 3b. SC-Series camera without filter wheel and lens with filter holder threads. The reddish circle in the camera is the warm window on the camera's dewar.**

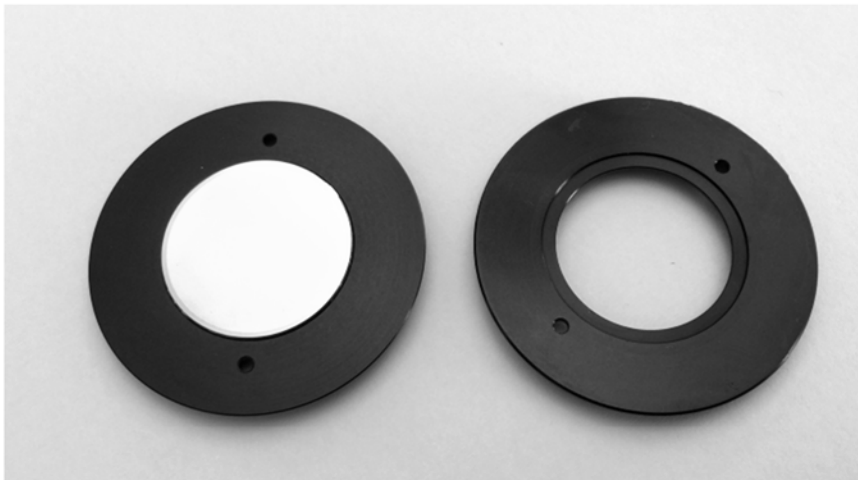

**Figure 3c. Behind-the-lens filter holders for three-tab bayonet lenses. The left one has a filter, the right one is empty. Users can use Kapton tape to hold filters in place, or silicon RTV. It is also possible to use 4 drops of fingernail polish around the filter perimeter to hold it. This can later be dissolved with fingernail polish remover.**

Warm filters are especially useful for imaging hot objects in different narrow bands. Hot objects have plenty of radiance, so the shortcomings of operating a thermal camera with a warm filter are greatly mitigated. Figures 4a-c show some images of a propane flame that show how its radiance changes drastically in different wavebands.

**Figure 4a. Visible light image of a clean propane flame**

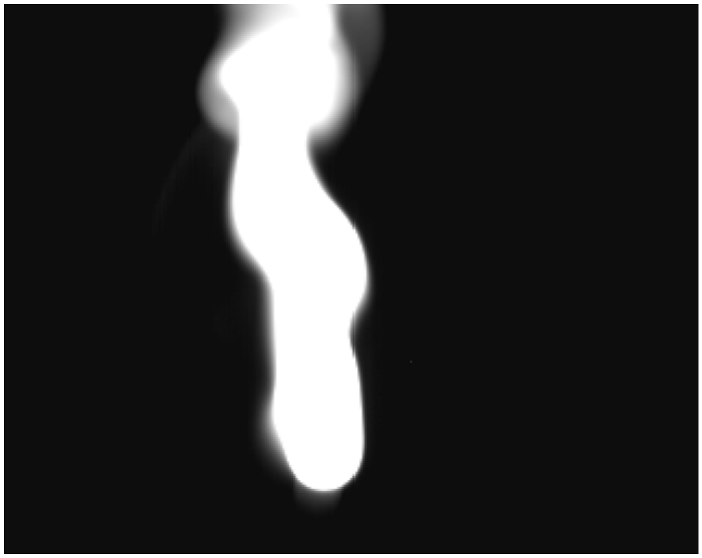

**Figure 4b. Propane flame imaged with a 4.2-4.4-micron warm filter (CO2 filter). CO2 absorbs and emits at this wavelength very strongly. Clean burning flames like this become very opaque.**

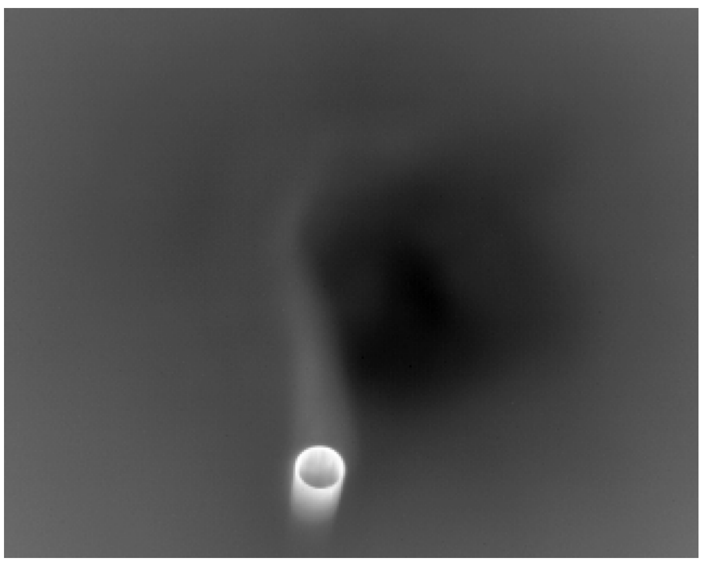

**Figure 4c. Propane flame imaged with a 3.8-4.0-micron warm filter (flame filter). CO2 does not emit or absorb in this band, nor does water vapor, the two constituents in clean-burning gas flames. You can measure the metal surface temperature of the inside of the torch nozzle in this band.**

The propane flame has very hot CO2 gas in it that emits very strongly at 4.3 microns. Right next door, at 3.9 microns, there is very little emission from the CO2 or water vapor in the flame. Figure 4d shows the spectral radiance curve for a Bunsen burner flame, which has a similar emission spectrum to the propane in the midwave band. You can see that there is a lot of emission at 4.3 microns and very little emission around 3.9 microns, just as Figures 4b and 4c show.

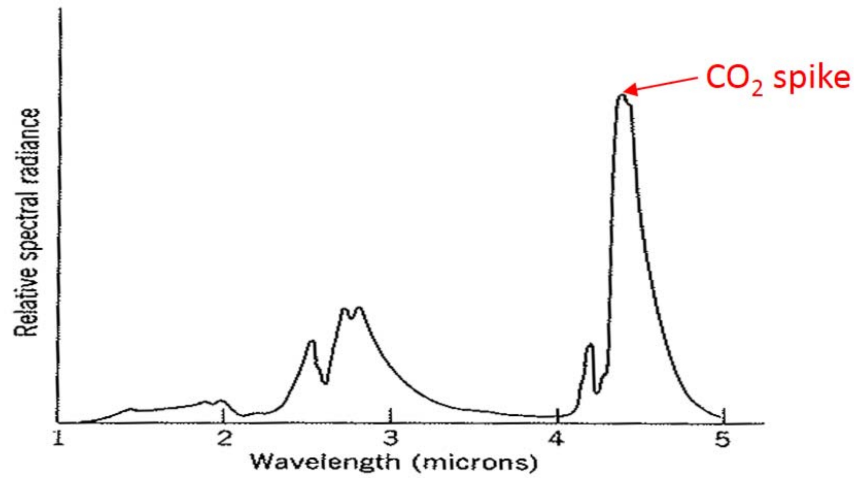


Infrared emission from a Bunsen flame (adapted from Plyler[36])

**Figure 4d. Bunsen flame emission spectrum**

## Warm-Filtered Camera

If used with an InSb camera, the filter shown back in Figure 2 will make the camera's photon spectral response to the scene have very nearly the same shape as the filter transmission curve, because the InSb sensor itself has a pretty flat photon spectral response between 3 and 5 microns. But there is a caveat. The filter is warm, which means it is placed outside of the camera's dewar (the refrigerated compartment where the IR sensor is housed). There will be some self-radiation from the filter itself as well as self-radiation from the inside of the camera housing reflected off the back side of the filter, both of which will contribute IR background to the camera's pixels, which creates a level offset and additional noise. Figure 5 shows a schematic diagram of this effect.

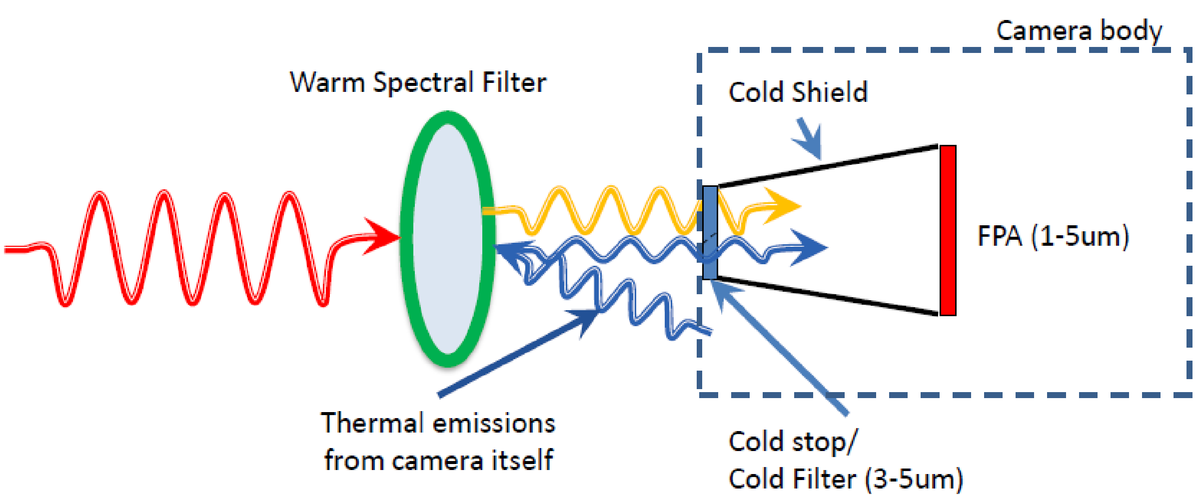

**Figure 5. The red light rays are all the radiation from the scene. The yellow rays are the "good stuff": the part of the red rays that are in the desired passband of the warm spectral filter. The undesirable blue rays are self-radiation from the inside of the camera body reflected off the back of the filter into the dewar, and self- radiation that comes from the filter itself. If the warm spectral filter is not within the passband of the cold filter, then no radiation from the scene will reach the FPA. Some users buy cameras without cold filters so they can image outside of the 3-5-micron band. These are the so-called broadband cameras.**

This self-radiation background is not correlated with the scene content, it is not very spatially uniform, and it has the unfortunate effect of filling the pixels' integration capacitors (wells) with unwanted background electrons (raising the so-called pedestal). These effects add to the temporal and spatial noise. One can NUC out the spatial non-uniformity using an external NUC source, but the results are only temporary, and the non-uniformity will drift back over time as the optics changes temperature. This effect is quite pronounced with narrow filters, and it can require doing an offset update about every ten minutes or so when imaging an ambient scene. At short integration times, the drift is much smaller, and the non-uniformity caused by the warm filter is negligible. The catch is that you can only get good measurements of hot targets – ambient targets will be imaged with very poor signal to noise ratios.

Radiometric calibration with a warm filter and a long integration time is very unstable and hard to control. The warm filter's temperature can become a big problem if the filter is very narrowband. Figure 6 shows a time plot of the mean counts for a 30C blackbody imaged by a camera with a 3.95-micron bandpass filter with a FWHM of 0.04 microns. The integration time is 50 msec, which is very long for an InSb camera with an f/4 coldshield. The large downward excursion of counts was caused by my blowing air out of my mouth across the warm filter! This configuration cannot be calibrated for radiometry with this much sensitivity to ambient temperature of the optics unless great care is taken to stabilize the temperature of the optics, which is not very practical or easy to achieve.

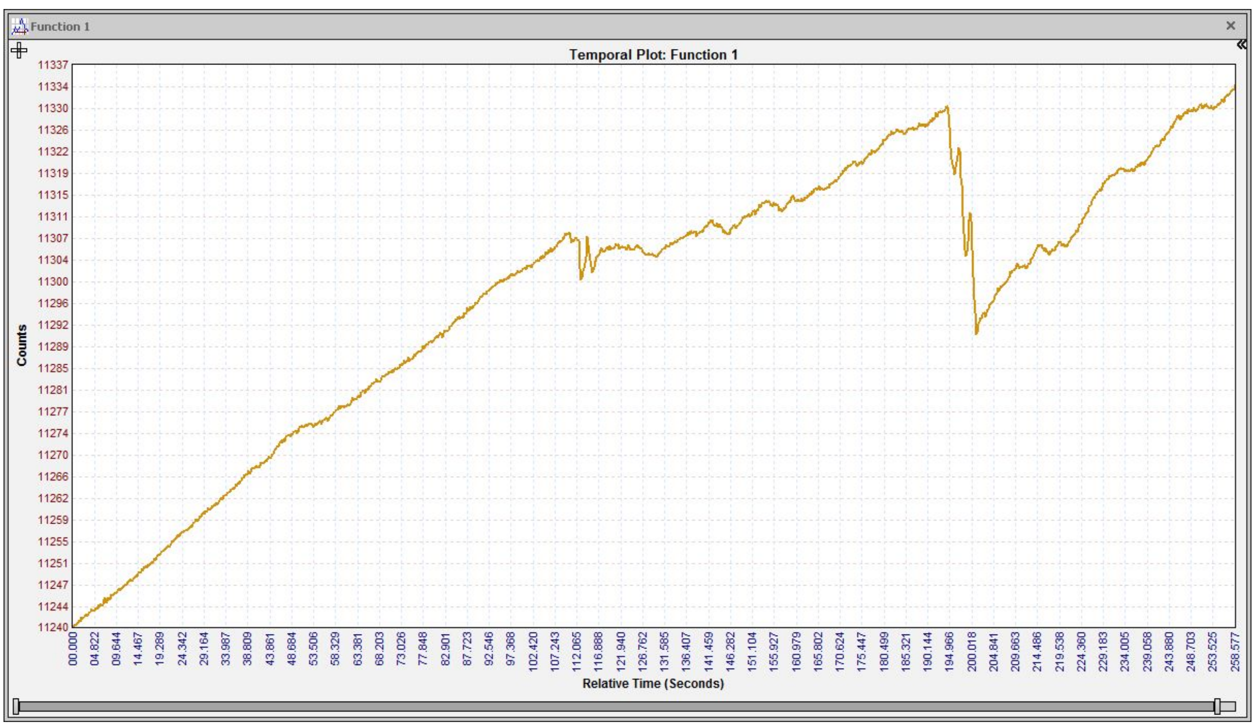

**Figure 6. Mean counts on a 30C blackbody versus time for a narrowband warm-filtered SC8300 camera with f/4 coldshield and a 50msec integration time. The time axis is 285 seconds long. The excursion at 200 seconds is about 40 counts.**

The problem of self-radiation raising the pedestal becomes particularly acute if the camera is operated at long integration times, which is necessary to do if the scene is not particularly hot (~ambient terrestrial temperatures), and the passband of the warm filter is narrow (FWHM < 0.5 microns).

One useful way to think about the self-radiation effect with warm filters is to consider an automobile driven at night with the dome light on. This situation makes it much harder for the driver to see the dark scenery outside, even with the headlights on. The ratio of outside illumination (the scene) to inside illumination (the self-radiation) is poor, as shown in Figure 7 below. Dark parts of the scene will be very hard for the driver to see. Night driving is analogous to not having much IR radiation from the scene relative to self-radiation inside the optics and camera body. In the midwave band, this is the case if the infrared scene is cool (<50C), or the camera is being operated with a narrowband warm filter. Turning up the integration time on the camera helps to some extent by improving the amount of signal, but then the camera will just see more and more of the self-radiation which will fill up the well. You end up running the system in a state such that your scene's mean count value is over 10,000 counts, which for a 14-bit camera means you do not have much scene dynamic range to work with before the camera saturates.

**Figure 7. The equivalent of an IR camera imaging an ambient scene through a narrowband warm filter with a long integration time. The couple in the car cannot see the outside world very well with all the interior light reflecting off the inside of the windows and into their eyes. Good thing they appear to be parked!**

If one is driving during the day, then it hardly matters if the dome light is on or off. The eye adjusts to the bright conditions outside, and the amount of dome light reflected off the windshield into the driver's eyes is negligible. This is analogous to an IR camera imaging a hot scene. The integration time of the camera must then be set to a smaller value (to avoid saturation) and thus the camera becomes very insensitive to any self-radiation.

If you are imaging a rather cool scene (<50C), which requires you to run at a long integration time (multiple msec), then a much more desirable place to install a narrow spectral filter is **inside** the dewar of the camera. Then there is no self-radiation from the filter itself, nor is there any reflection of radiation off the filter and onto the focal plane array, as shown in Figure 8. In this case, the "dome light" is always off and the driver does not see any reflected light off the inside of the "windshield".

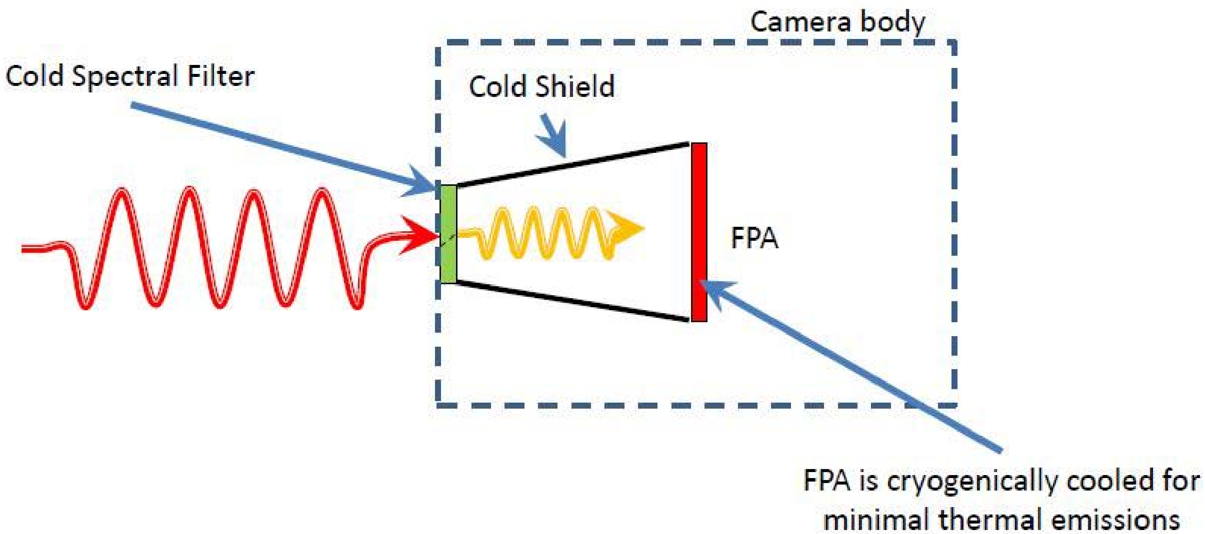


**Figure 8. In this configuration, the cold spectral filter has negligible self-radiation, and its back side is reflecting the interior of the coldshield and the focal plane array, both of which are at cryogenic temperatures and therefore have negligible IR radiance. The red rays represent the unfiltered radiation from the scene, and the yellow rays are the "good stuff" or the desired band of radiation that makes it through the cold filter.**

There is a catch, however. Modern IR cameras with cryocoolers have vacuum-sealed dewars, and the cold filter must be chosen at the time the camera is manufactured. Once that decision is made, the cold filter is physically "baked into" the camera, and only a very invasive procedure can change the cold filter for a different one. Most midwave infrared cameras are made with a 3-5-micron bandpass cold filter, and these cameras can then be equipped with an additional warm filter to restrict the spectral response of the camera to a sub-band of the 3-5-micron band.

If the scene being imaged is hot enough, the camera can be operated at a short integration time (<1msec for an f/2.5 camera), and the self-radiation due to the warm filter will not have a serious effect on the imagery. For ambient scene temperatures in the 40C and lower temperature range, the self-radiation can a much bigger problem, as we will see in the next section.

For an InSb camera, the typical spectral response of the InSb itself is from 1.0 microns to 5.7 microns. On most of the Niceville science cameras, the silicon warm window on the dewar is coated to transmit above 1.5 microns.[18] FLIR InSb science cameras can be made without any cold filter – a variant called "broadband InSb". The camera user can then add a warm spectral filter to the optical assembly outside of the dewar and restrict the camera to imaging in a sub-band of the broadband InSb spectral response – this is often done to access the SWIR part of the spectrum (1-3 microns). The user can also use a 3-5-micron

[18] The A6700 camera series uses a sapphire warm window that transmits down into the visible, enabling 1 micron and up camera spectral response

warm filter, which will mimic the spectral response of a standard 3-5 micron cold-filtered camera. On paper, this option is quite attractive for users that want to work in both the midwave and SWIR bands with a single camera. In practice, this approach works well for higher scene temperatures, but generally fails to give good image quality for cold scene temperatures. This is because there is a limited amount of background from a cold scene, requiring long integration times which end up causing the camera to integrate unwanted self-radiation up to 5.7 microns.

The typical camera is operated in a 25C environment and the interior temperature of the camera body is generally around 36C. A 36C scene has a Planck curve that is steeply rising in the 1.5-5.7-micron band, as shown in Figure 9. A camera with a warm filter that is reflecting significant amount of camera internal self-radiation onto the sensor will have its image quality strongly affected, unless the scene is sufficiently bright, i.e. hot.

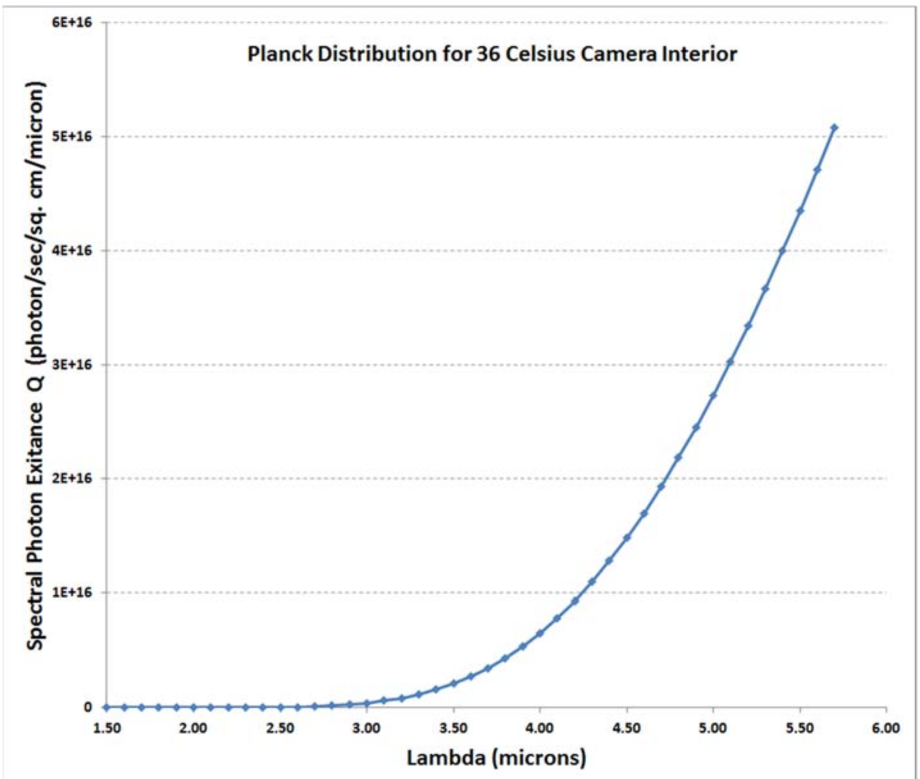


**Figure 9. Planck distribution for 36C camera interior**

A narrow warm filter will reflect camera internal self-radiation onto the sensor. The shorter the center wavelength of the filter, the worse the ratio of scene radiation to camera internal self-radiation for a typical ambient temperature scene.

The camera internal self-radiation (CISR) is also not uniformly spread across the camera's sensor. The filter is close to the warm window and much of what is reflected off the back of the warm filter is the sensor itself, which is at cryogenic temperature. Still, there is a "leak" around the edges caused by the gap between the warm filter and the warm window on the dewar. There is always a little space between a warm filter attached to the back of a lens, and the warm window in a science camera. That space can be small in the case of a camera with a filter wheel, but it still is a few millimeters. That is how the reflected self-radiation ends up getting into the dewar and onto the focal plane, and how the appearance of the self-radiation is that of a cold spot in the center of the image and warmer pixels around it, as shown in Figure 10:

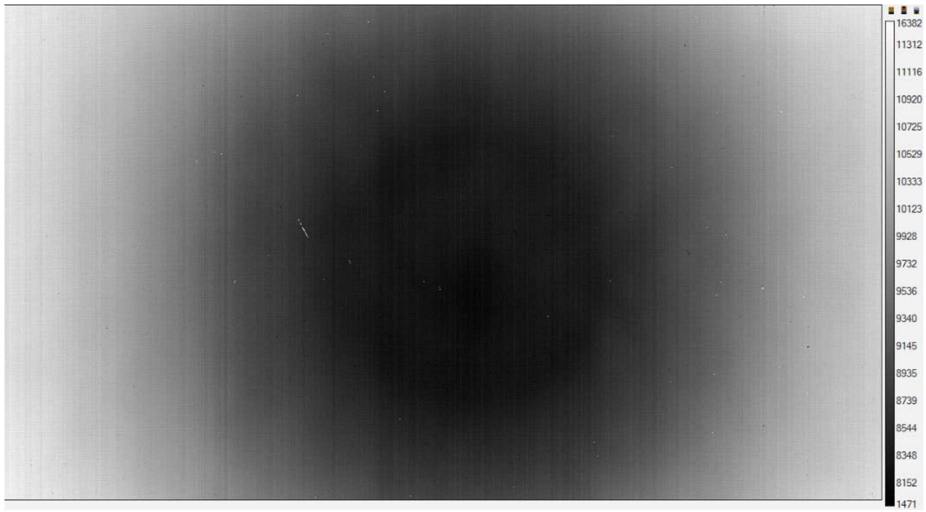


**Figure 10.  Narcissus reflection of the cold stop by a warm filter mounted behind the lens**

In some cases, it can be advantageous to install a warm filter right on top of the warm window of a camera dewar to minimize the non-uniformity of the self-radiation, since there is no significant gap for CISR to leak in.  Figure 11 shows a warm filter installation done using Kapton tape.  This approach has to be done carefully to make sure that the offset update shutter does not interfere with the filter. As long as the filter is around 1mm thick, there is sufficient clearance for the shutter on the FLIR Niceville InSb cameras.

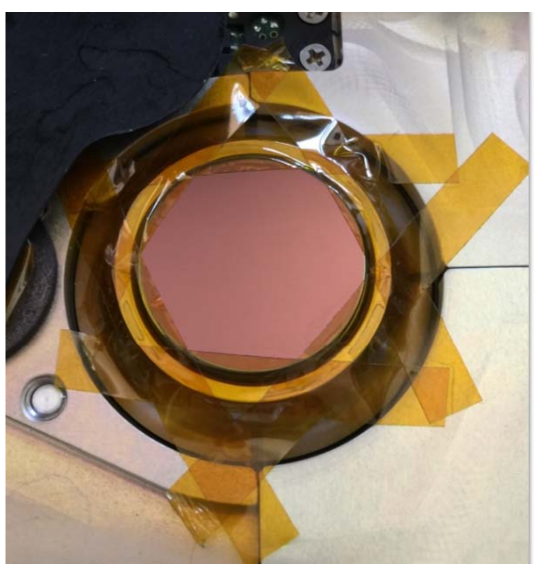

**Figure 11. A narrowband warm filter taped over the warm window on an SC8303**

## Case Study - Warm filter versus Cold Filter

A study was made comparing the behavior of two InSb science cameras which were identical in many respects, except that one camera had a narrowband cold filter, and the other camera had an identical narrowband warm filter mounted behind the lens, as well as a 3-5-micron bandpass cold filter.

The two cameras (shown in Figure 12) were both from the A8300HD series, with 1344x784 InSb sensors with 14-micron pixels, f/4 coldshields and identical lenses. The narrowband cold filter and warm filter both had the same 3.2-3.4 micron bandpass, which is the band of the hydrocarbon gas imaging filter.

**Figure 12. Two narrowband InSb cameras looking at an area blackbody**

The two cameras were initially set to 20msec integration times and the digital counts observed on a region of interest of a blackbody at four different temperatures were measured. The blackbody was an ISDC 6-inch area blackbody with an emissivity of 0.96.

There are three main differences between the behavior of the two cameras at this rather long integration time.

1) The warm filtered camera had a significantly higher count value when looking at the same temperature blackbodies. It is integrating more CISR that the cold filtered camera.

2) The uncorrected image of the warm filtered camera showed a strong narcissus effect – reflections of the cold stop from the back of the warm filter. This could be temporarily eliminated with an external offset update. The cold-filtered camera does not show a narcissus, since there is no warm filter that is highly reflective out of band.

3) The warm-filtered camera's optics and internal temperature drifted, and the narcissus non-uniformity tended to return after about 15 minutes of operation. The cold-filtered camera's non-uniformity was much more stable over time.

## Warm-Filtered Methane Camera

Table 2 shows the image mean count values observed with the two cameras on the ISDC blackbody.

| Radiance Source | Warm-Filtered Camera | Cold-Filtered Camera | Integration time (msec) |
|---|---|---|---|
| | | | |
| 20C Area BB | **9849** (higher than 25C point) | 2168 | 20 |
| 25C Area BB | 9766 | 2208 | 20 |
| 30C Area BB | 9875 | 2260 | 20 |
| 40C Area BB | 10264 | 2397 | 20 |

**Table 2. Camera count values for floor illumination with various blackbody temperatures.**

At 20 msec, the warm filtered camera has about 7600 counts of excess radiation on the sensor compared to the cold-filtered camera. This is about half of the available counts in the camera's digitizer. There is a lot of self-radiation in the system, and it greatly limits the temperature range of the warm-filtered camera, since it is already at 10,264 counts for a 40C scene. This camera cannot be radiometrically calibrated, since the 20C data point for the warm-filtered camera has a higher count value than the 25C camera point, indicated that there is a big drift in the CISR. The camera was heating up over a period of ~10 minutes when these data points were captured, and the 20C point was taken after the 25C and 30C points. The camera interior temperature increased about 0.5C between the 25C point and the 20C point collections. This illustrates just how susceptible the system is to drifts induced by CISR, and why long integration times combined with narrowband warm filters add up to a very "drifty" radiometric cameras system that cannot be reliably calibrated. The radiometric transfer function looks terrible, as shown in Figure 13, and there is no way to calibrate the camera using these data.

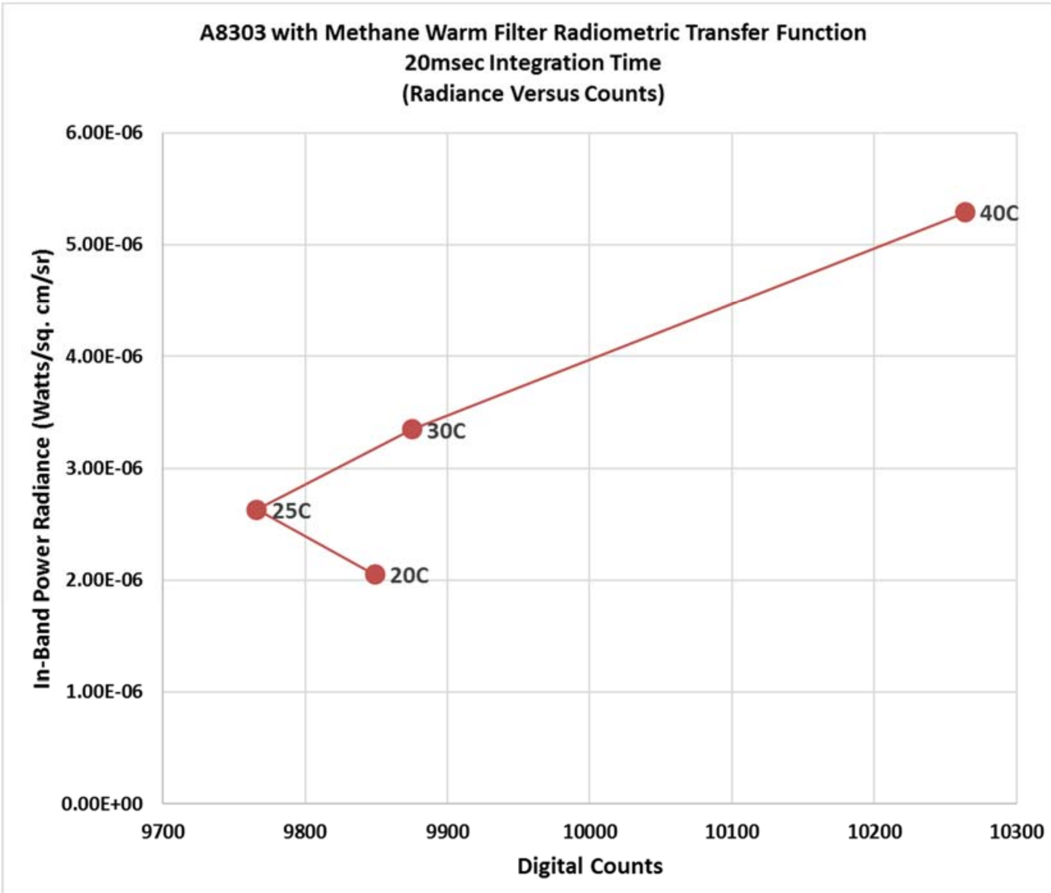


**Figure 13. Radiometric transfer function for warm-filtered methane camera calibration at 20msec. There are large systematic errors in the count values due to the CISR changing during the calibration process from the camera self-heating.**

## Cold-Filtered Methane Camera

In contrast, the cold-filtered methane camera can be precisely radiometric calibrated to measure a 20C scene (or colder) up to an estimated 140C scene temperature, according to a radiance versus counts model done in FLIR CalibratIR software and extrapolated to 14,000 counts, the typical limit for linear behavior of the science-grade cameras. Figure 14a and 14b show the calibration measurement grid in the utility CalibratIR, and a plot of radiance versus counts for the cold-filtered camera data.

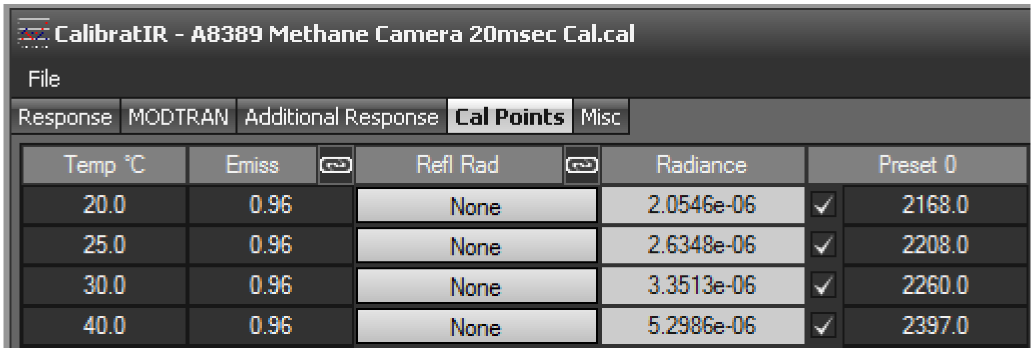

CalibratIR - A8389 Methane Camera 20msec Cal.cal

File

Response | MODTRAN | Additional Response | Cal Points | Misc

| Temp °C | Emiss | Refl Rad | Radiance | | Preset 0 |
|---|---|---|---|---|---|
| 20.0 | 0.96 | None | 2.0546e-06 | ✓ | 2168.0 |
| 25.0 | 0.96 | None | 2.6348e-06 | ✓ | 2208.0 |
| 30.0 | 0.96 | None | 3.3513e-06 | ✓ | 2260.0 |
| 40.0 | 0.96 | None | 5.2986e-06 | ✓ | 2397.0 |

**Figure 14a. Measurement grid in CalibratIR for cold-filtered methane camera calibration at 20msec**

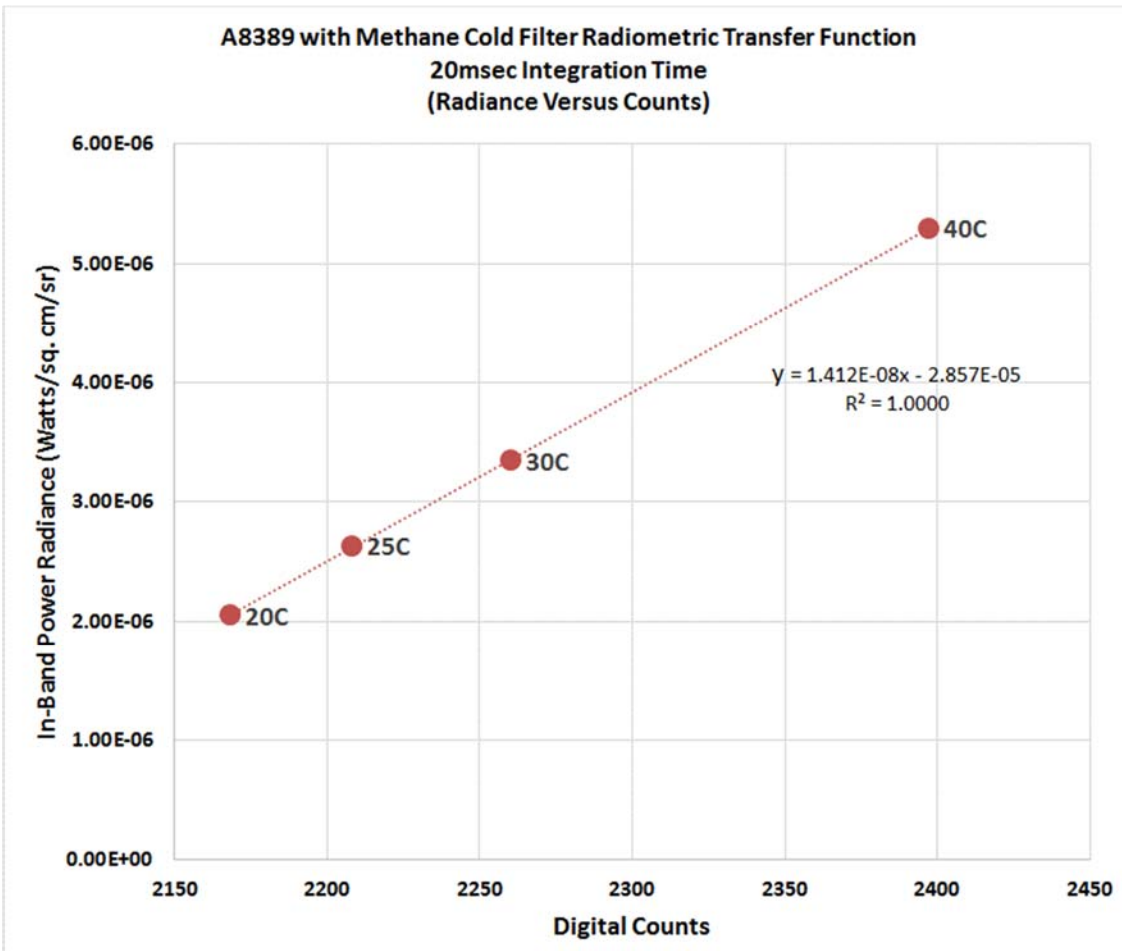


**Figure 14b. Radiometric transfer function for cold-filtered methane camera calibration at 20msec**

## Hotter Scene Temperatures

I turned down the integration time to 3msec and imaged a 200C cavity blackbody. The warm-filtered camera is still seeing an excess of counts, no doubt due to the greater CISR caused by the warm filter.

At still shorter integration times and higher scene temperatures, the count differences become less, as shown in Table 3.  It is not clear why the 400C cavity produces ~400 counts less in the warm-filtered camera, unless it is the difference in optics transmission, since there is an extra filter.  But then why are the 300C measurement count values so close?  I don't know, and further study is warranted.

| Radiance Source | Warm Filtered Camera | Cold Filtered Camera | Integration time (msec) |
|---|---|---|---|
| 200C Cavity BB | 8818 | 7724 | 3 |
| 300C Cavity BB | 6548 | 6575 | 0.5 |
| 400C Cavity BB | 7375 | 7763 | 0.2 |

**Table 3.  Cavity blackbody measurements with the two cameras**

At still shorter integration times, I would expect that the two cameras would behave very similarly in terms of their susceptibility to optics temperature drift. This suggests an axiom:

***As the scene temperature range is increased, a warm-filtered camera behaves more like a cold-filtered camera.***

Here is some quantitative data to back up this assertion.  Consider a camera with a warm filter installed behind the lens, about 3mm from the warm window. A schematic is shown in Figure 15. The camera has a 3-5-micron cold filter in it.  The warm filter has a bandpass of 3.2-3.4 microns.  The interior of the camera and the optics are at a temperature of 33C, which is a typical value measured for a warmed-up camera operating in a 23C lab.  The radiance of a 33C blackbody in the 3.0-3.2-micron range is calculated and summed with the radiance of a 33C blackbody between 3.4 and 5 microns. The assumption is that the warm filter is reflecting the inside of the camera body outside of the filter's passband of 3.2-3.4 microns, so there are two bands of reflected radiation because the 3-5-micron radiance is notched out between 3.2 and 3.4 microns by the filter being transmissive in that range. This is a fair assumption for these types of interference filters:  They transmit in-band and reflect out of band, without significant absorption or emission.  Neutral density filters can have much higher absorption values.

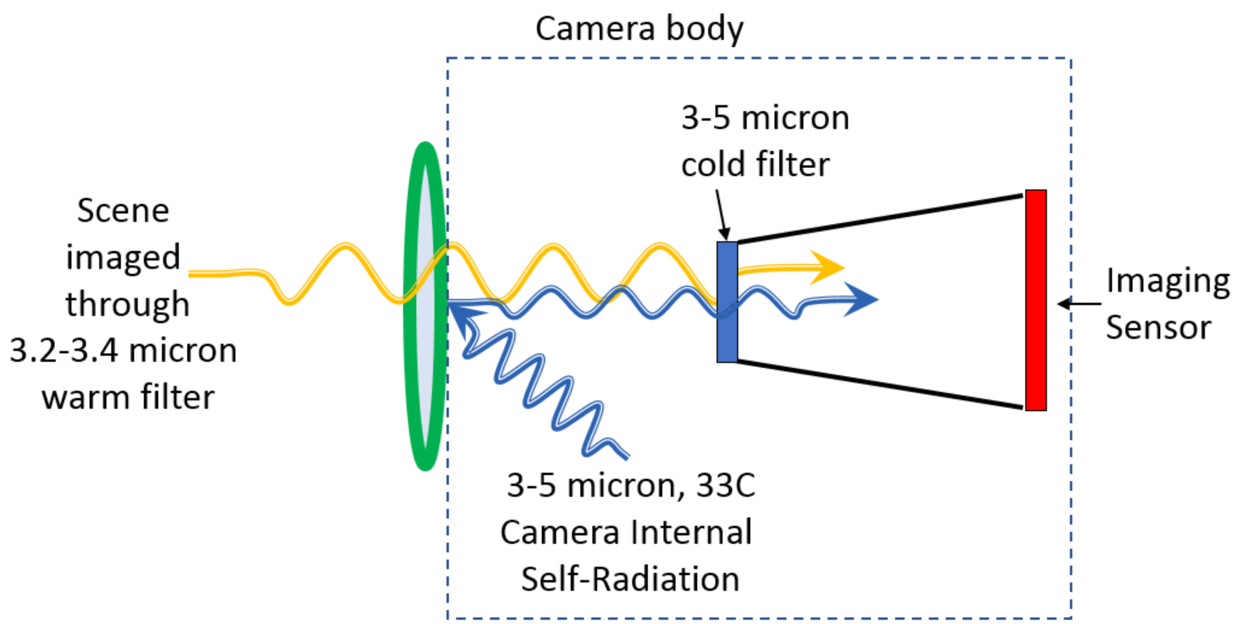


**Figure 15. Schematic diagram of warm-filtered methane camera**

This summed radiance is actually a very conservative estimate of the amount of CISR reaching the sensor for this warm filter and a 33C internal camera temperature. In reality, the filter is within a few mm of the warm window and will reflect back the 77K cold shield temperature, particularly in the center of the image. The total CISR photon radiance is 5.03e15 photons/sec/sq. cm/sr (3-3.2 and 3.4-5 micron, 33C blackbody. The warm filter in-band scene radiance is calculated for various scene temperatures between 20 and 400C and plotted in Figure 16a. The scene is really at a disadvantage relative to the CISR and the camera will be very noisy and "drifty" for scene temperatures around ambient. Consider a scene at the same temperature as the inside of the camera and optics. For a 33C scene temperature, the scene radiance is only 1.3% of the 33C internal temperature CISR. The scene hardly contributes to the radiation striking the sensor compared to the CISR. In this situation, the camera is much more sensitive to the internal camera temperature than to the scene itself.

**The critical takeaway is that the maximum useful integration time one can set for a camera with a narrowband warm filter is limited by the camera internal self-radiation (CISR), NOT the scene radiance.**

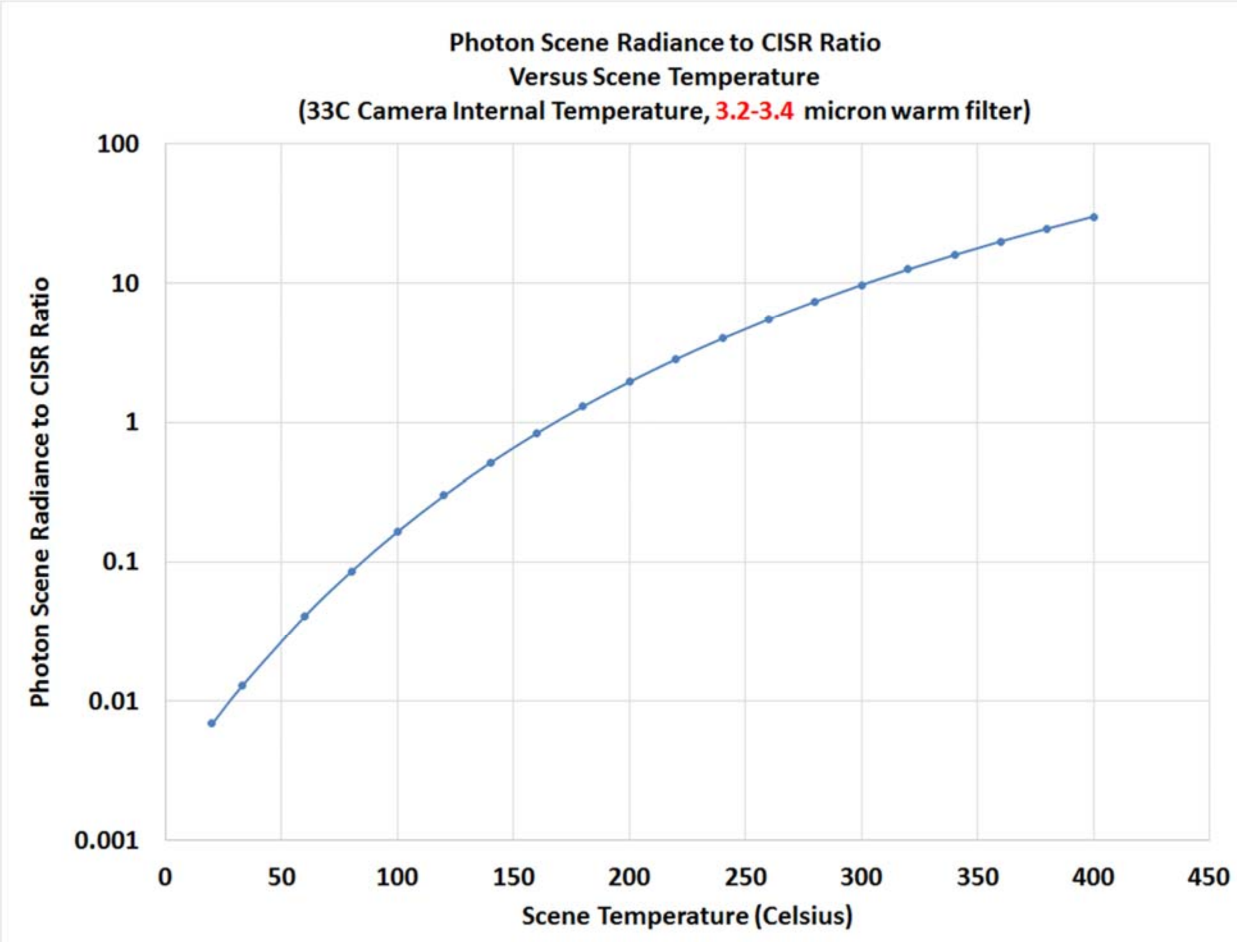


**Figure 16a. Photon scene radiance to CISR radiance ratio versus scene temperature for 33C optics and methane warm filter in the 3.2-3.4-micron band, all in a 3-5-micron camera. Above 300C, the scene radiance is 10X the CISR, and the camera will be very insensitive to the warm filter effects, i.e. the camera will act like it has a cold filter in the 3.2-3.4-micron band.**

The radiance ratio curve is more favorable for a carbon dioxide gas imaging warm filter (4.2-4.4 microns) in a 3-5-micron camera with a 33C camera internal temperature, as shown in Figure 16b. The scene radiance is ~9% of the CISR at 33C scene temperatures:

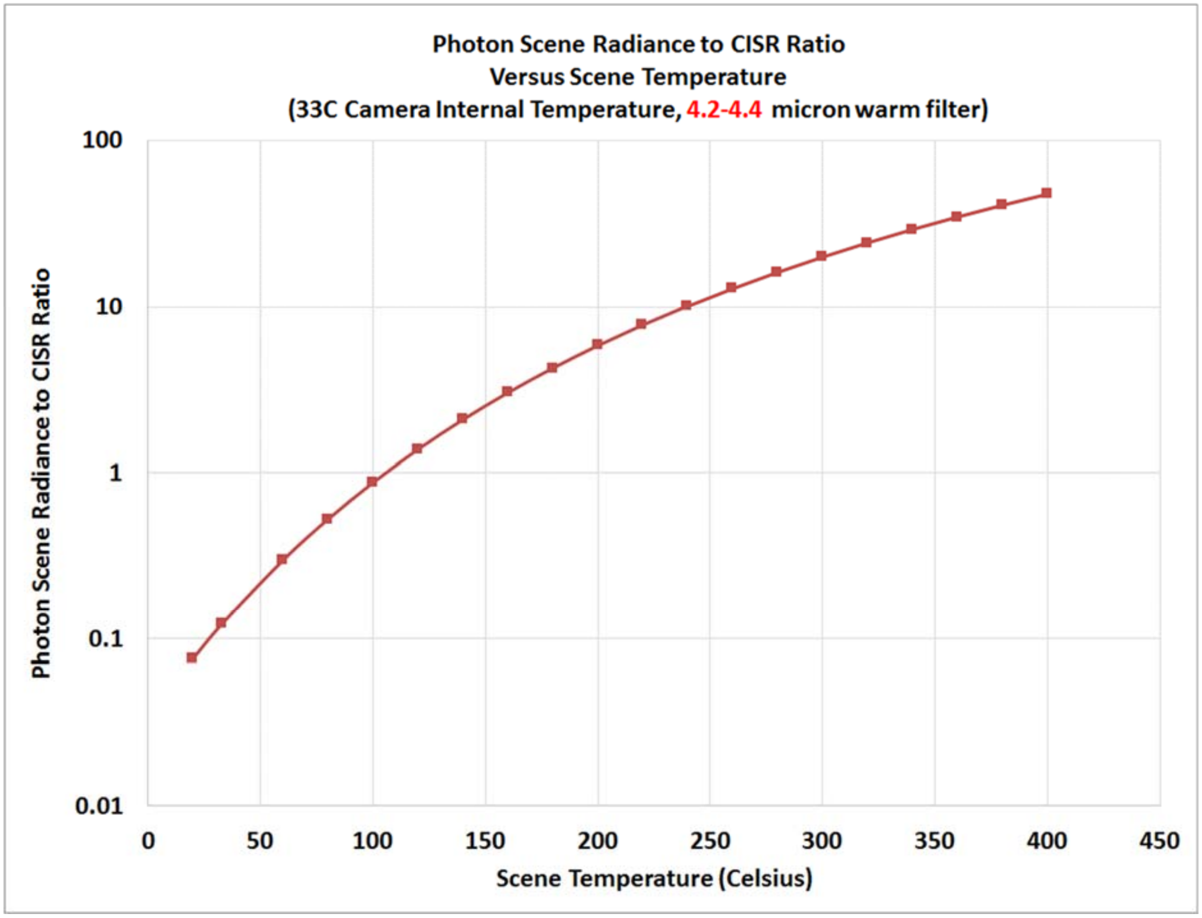


**Figure 16b. Photon scene radiance to CISR radiance ratio versus scene temperature for 33C optics and carbon dioxide warm filter in the 4.2-4.4-micron band, all in a 3-5-micron camera. Above 240C, the scene radiance is 10X the CISR and the camera will be very insensitive to the warm filter effects, i.e. the camera will act like it has a cold filter in the 4.2-4.4-micron band. This warm filter camera configuration is much less sensitive to drift effects when imaging an ambient scene that the methane filter configuration plotted in Figure 16a. There is a lot more in-band radiation when you go from 3.3 microns center wavelength to 4.3 microns.**

If one has a narrowband cold-filtered camera, and one is imaging an ambient temperature scene (something close to the camera internal temperature), the integration time can be greatly increased over a what a 3-5-micron camera with a similar narrowband warm filter can support. The cold-filtered methane camera used in these experiments was operated at **150 msec** integration time as its standard factory preset. A 25C scene imaged by this camera at 150 msec has a mean of about 7000 counts. The warm-filtered camera could not be set to integration times much longer than about 25msec without saturating on a 25C scene. Figure 17a shows the warm-filtered camera imaging a stream of propane against a 20C blackbody. The image is noisy, and the gas is hard to see. Figure 17b shows an image of the 30C blackbody and propane stream with the cold-filtered camera at the same integration time. The image is less noisy. Figures 17c-d show the same gas stream with a 100C blackbody backlighting. The gas is much easier to see and there is less difference in image quality between the two camera types. The major difference is the much larger pedestal of CISR in the warm-filtered camera.

***At the maximum useful integration times of these two cameras, the cold-filtered methane camera will be about 6 times more sensitive for a 25C scene relative to the warm-filtered methane camera.***

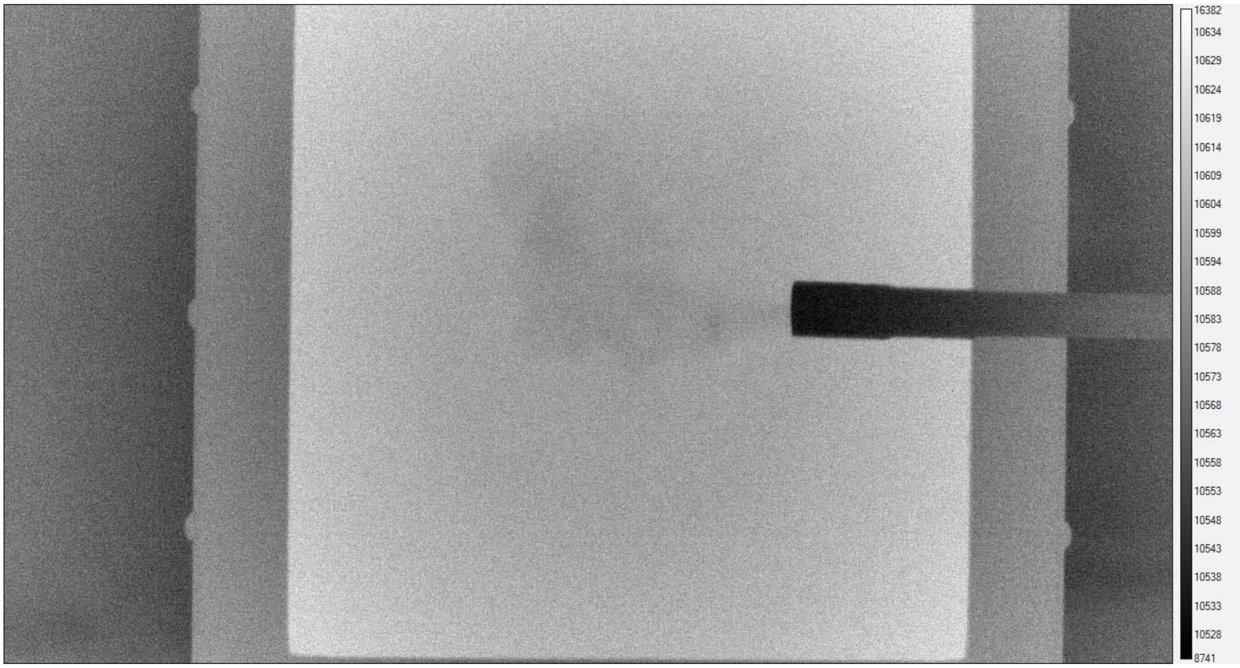


**Figure 17a. Warm-filtered methane camera imaging a 30C blackbody. Integration time is 20msec, and the mean digital count value on the blackbody is ~10,600. A 30msec integration time resulted in a saturated camera. The propane gas is hard to see with the temporal noise. Since the propane gas cloud is dynamic, frame averaging to reduce temporal noise won't help.**

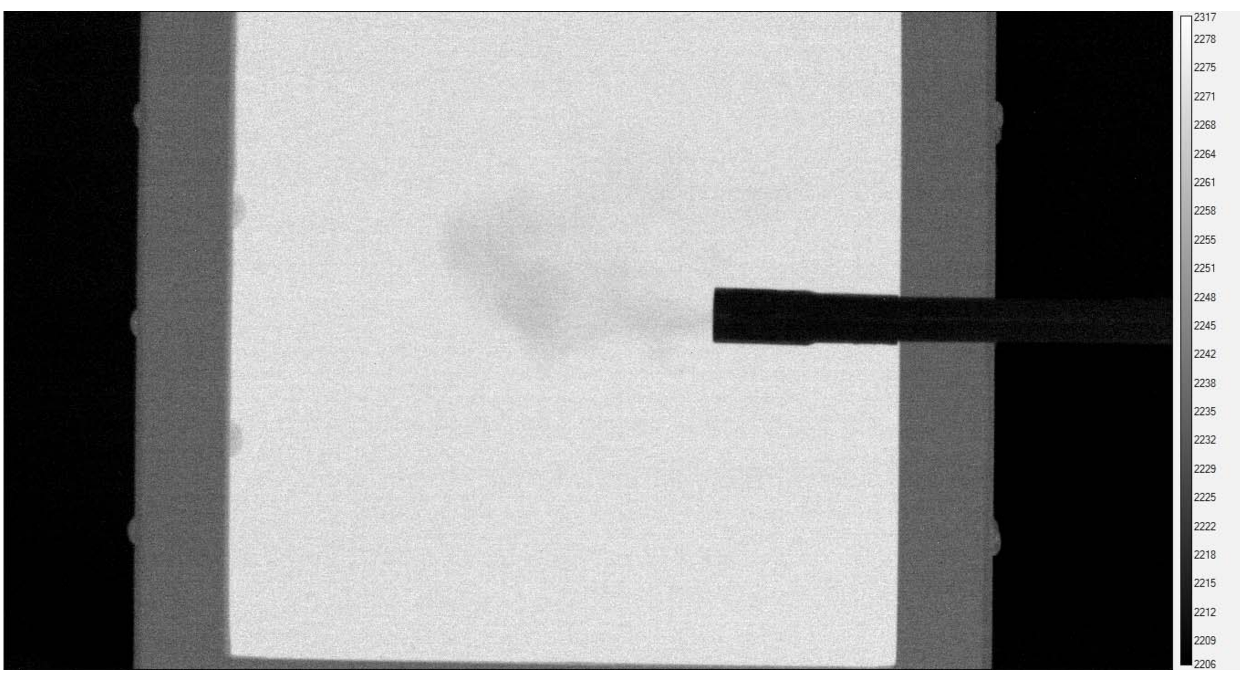

**Figure 17b. Cold-filtered methane camera imaging a 30C blackbody. Integration time is 150msec, and the mean digital count value on the blackbody is ~7100 counts. The camera is ~6 times more sensitive than the warm-filtered version. The gas is more visible than to the warm-filtered camera, though there is little contrast in radiance between the gas and the blackbody.**

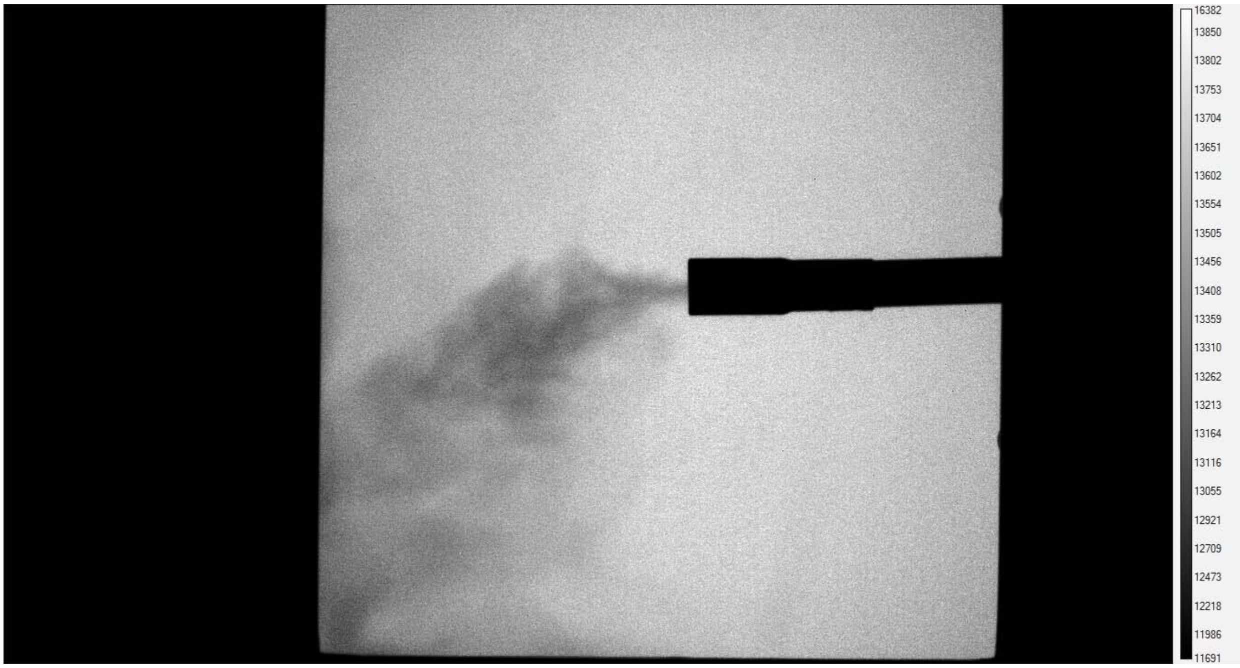


**Figure 17c. Warm-filtered methane camera imaging a 100C blackbody backlighting an unlit propane torch emitting propane gas. Integration time is 20 msec.**

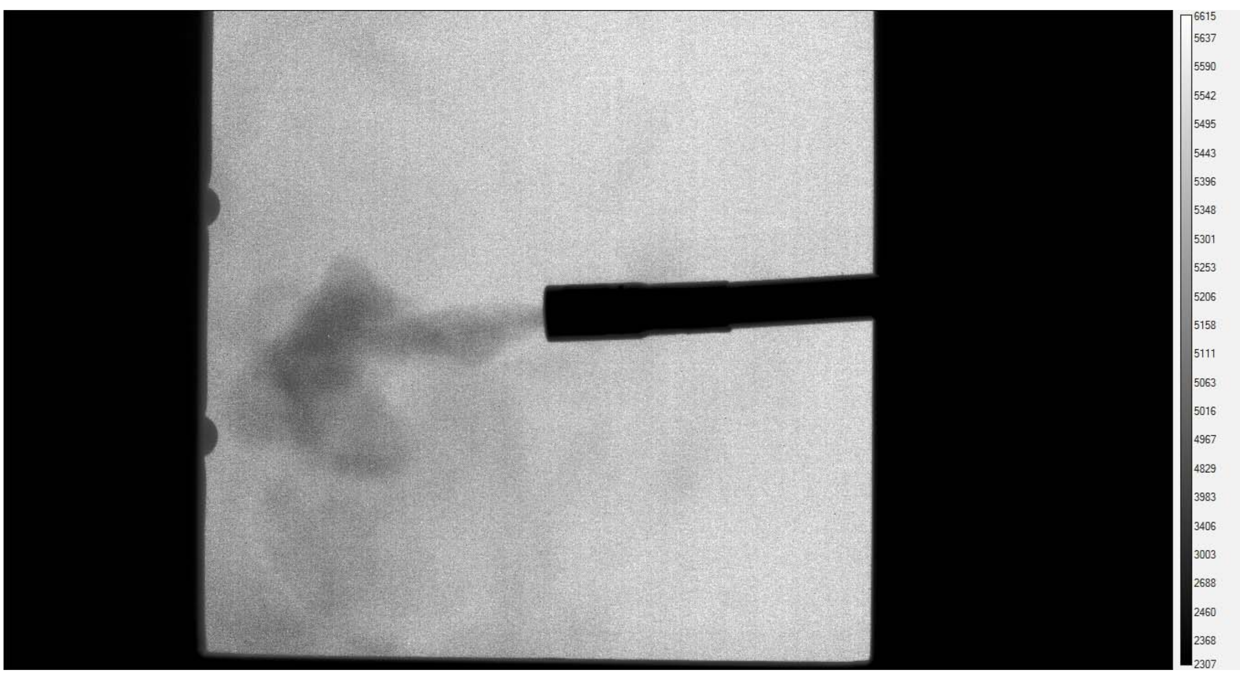


**Figure 17d. Cold-filtered methane camera imaging a 100C blackbody backlighting an unlit propane torch emitting propane gas. Integration time is 20 msec.**

## Non-Uniformity with No NUC applied:

Another troublesome problem which affects both imaging and radiometric applications is the non-uniformity introduced by a warm filter in a narrow band. Figures 18a-b show a uniform 23C scene imaged with the two cameras with their NUCs turned off. The narcissus effect from the warm filter is very apparent, with a ~3,000 count dip in the center relative to the corners:

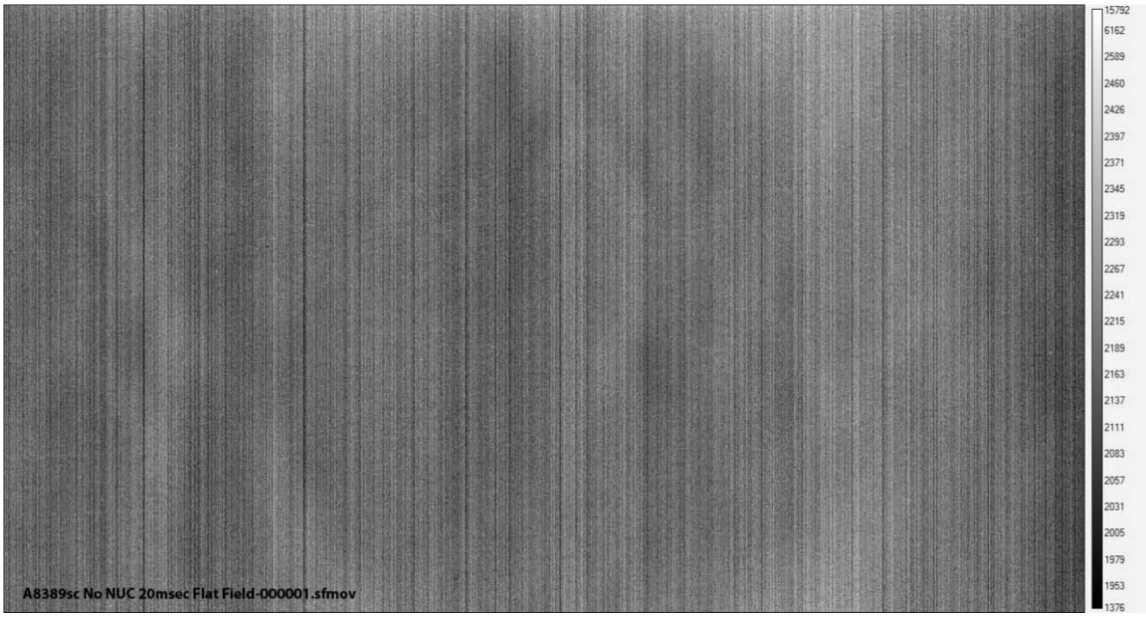


**Figure 18a. Cold-filtered camera looking at a 23C ambient uniform scene at 20msec integration time. The NUC is turned off. The image mean is 2190 counts and the standard deviation is 97 counts. There is no circularly symmetric non-uniformity structure.**

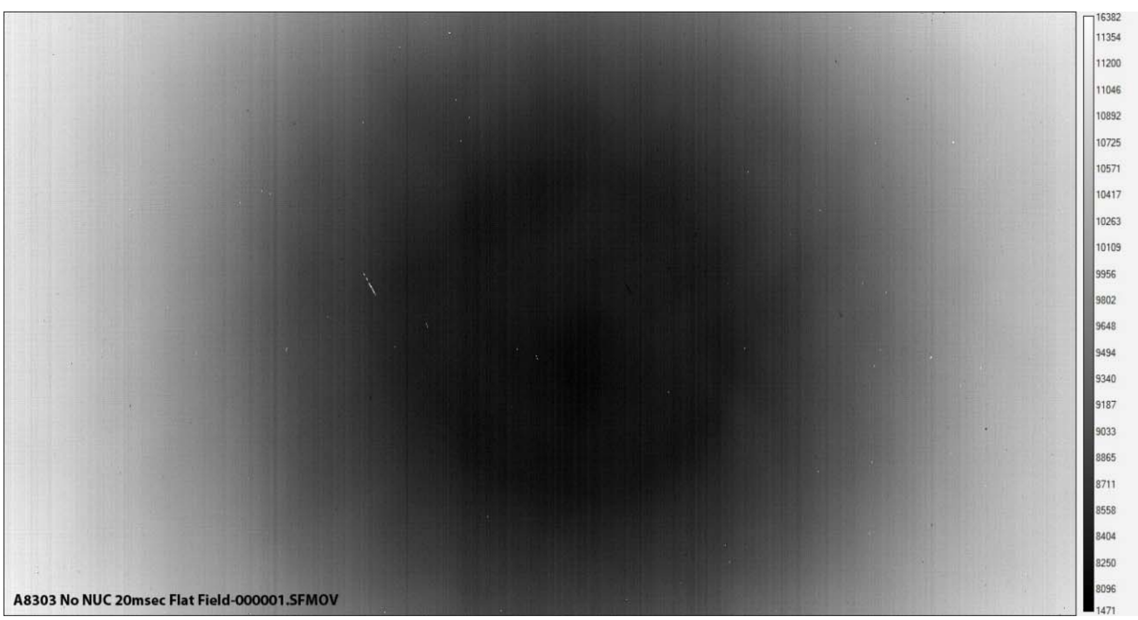


**Figure 18b. Warm-filtered camera looking at the same 23C ambient uniform scene with a 20msec integration time. The NUC is turned off. The image mean is 9447 counts and the standard deviation is 920 counts. The**

**huge standard deviation is caused by the cold narcissus reflection of the cold stop by the warm filter. There is also copious CISR that is causing the much higher image mean compared to the mean of the image in Figure 18a.**

## External Offset Update Time Drift

The two camera configurations exhibit very different behaviors when allowed to operate for hours. The warm-filtered camera non-uniformity is much more affected by temperature drifts in the optics compared to the cold filtered camera, as shown in Figures 19a-c. This drift issue manifests itself as a cold "bullseye" pattern in the center of the image – the so-called narcissus effect where the focal plane is viewing a reflection of itself in the reflective bandpass filter mounted (in this case) behind the lens about 5 mm from the warm window of the dewar.

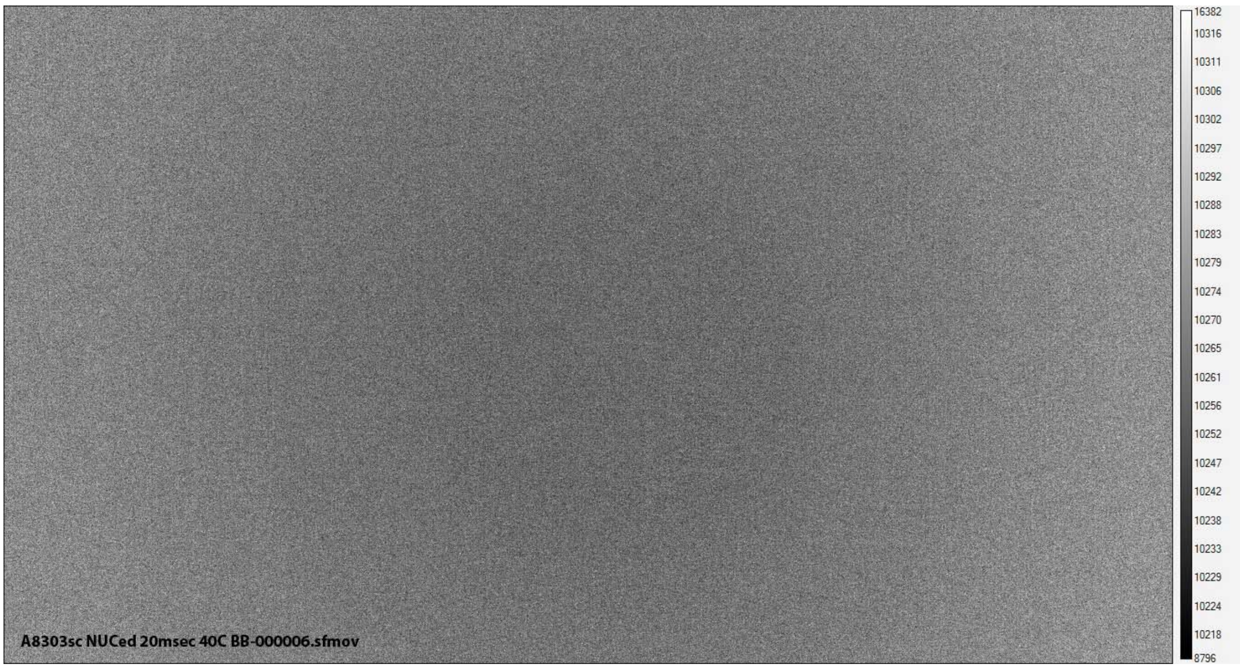


**Figure 19a. Image of a 40C blackbody after the warm-filtered camera has had an external offset update applied. 20msec integration time. The mean is 10,265 counts, and the standard deviation is 11.5 counts.**

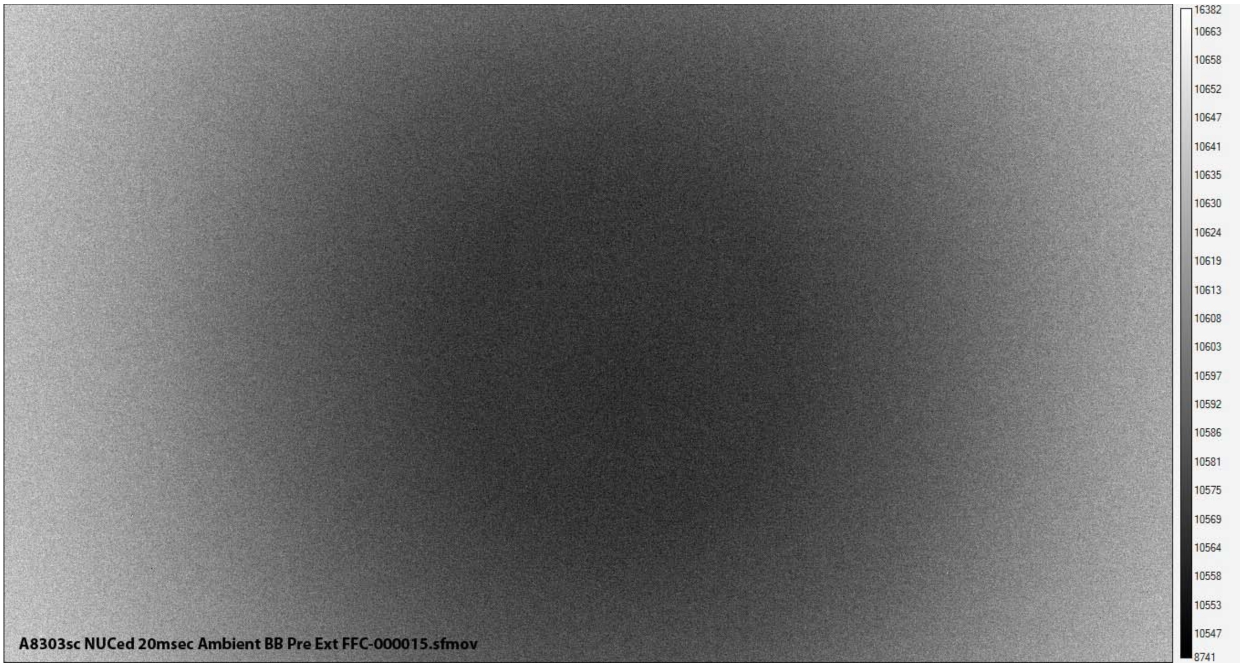


**Figure 19b. Ambient uniform source being viewed 105 minutes later, after the system has drifted and the narcissus has reappeared. The mean is 10595 counts, and the standard deviation is 22 counts. 20 msec integration time.**

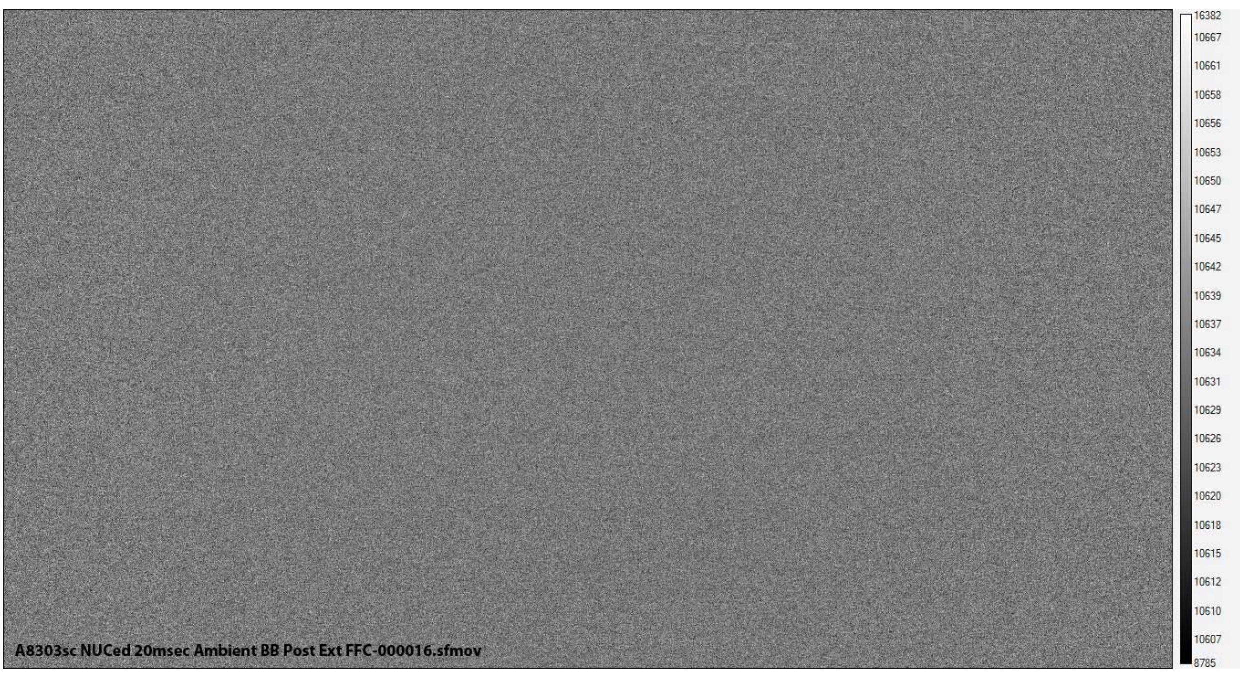


**Figure 19c. Ambient uniform source right after external offset update which cleans up the fixed-pattern noise. The standard deviation has dropped to 8.6 counts, and the mean is 10,634 counts. 20 msec integration time.**

The cold-filtered camera is much more stable over time, as shown in Figures 20a-c.

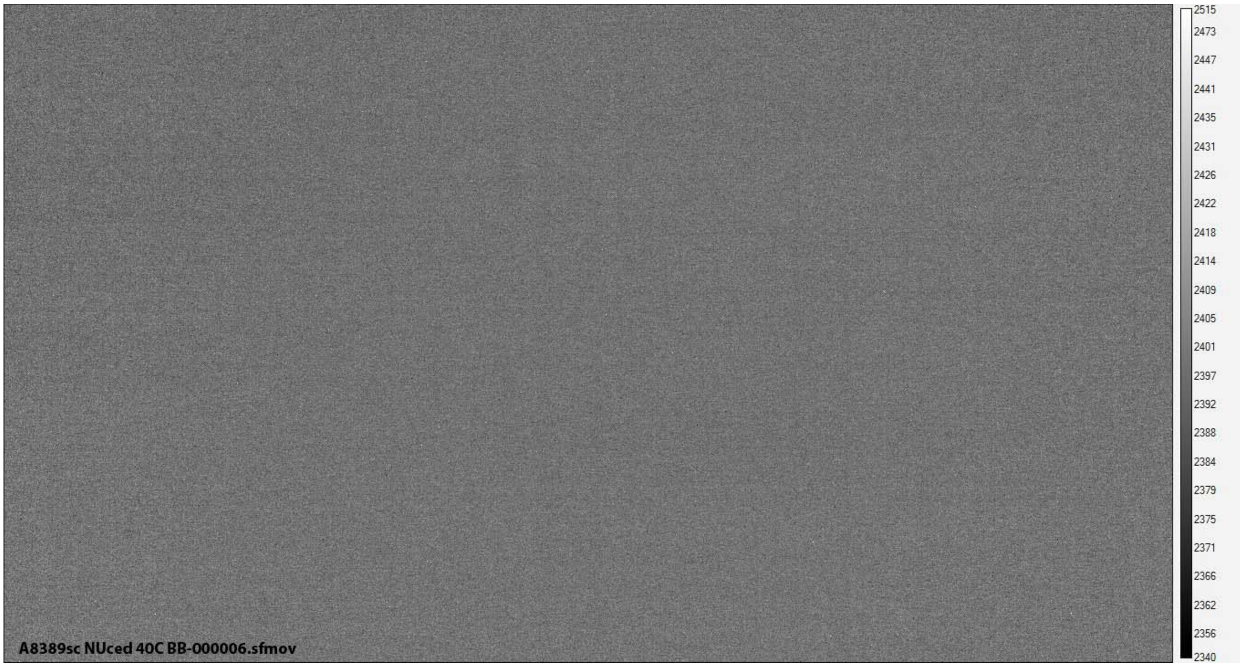


**Figure 20a. 40C blackbody after cold-filtered camera has had an external offset update applied. 5.3 counts standard deviation and 2397 counts image mean. 20msec integration time.**

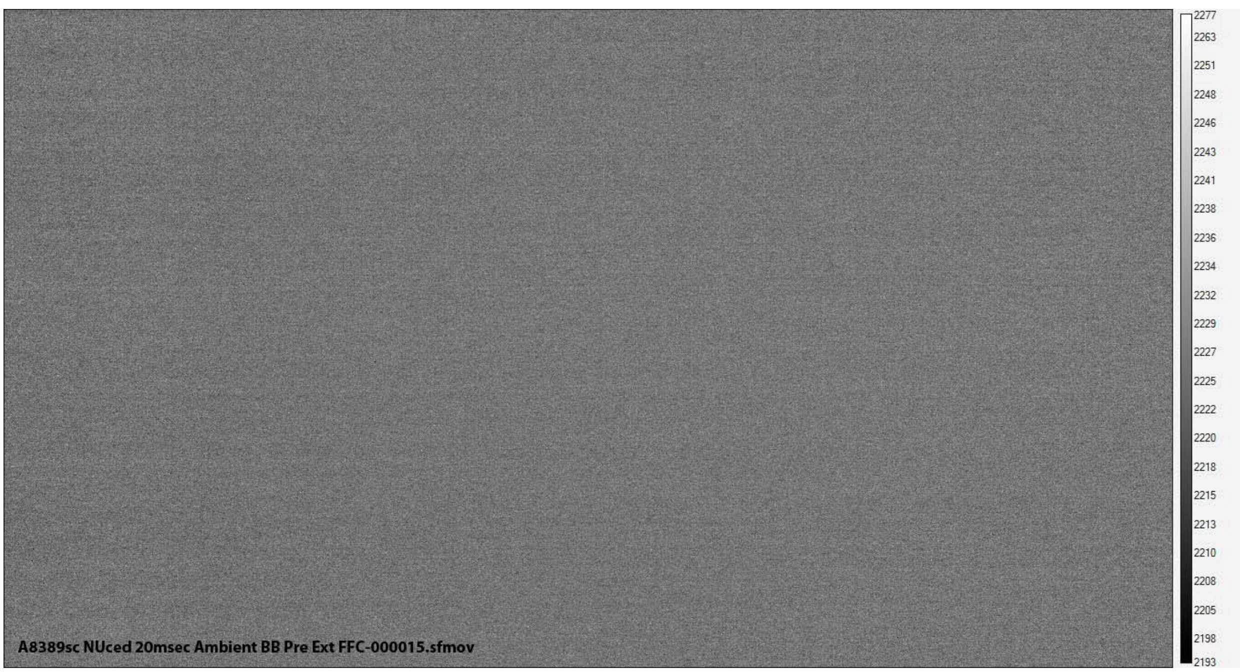


**Figure 20b. Ambient uniform source being viewed 105 minutes later, after the system has been allowed to drift. There is virtually NO change in the image uniformity! The standard deviation is 3.2 counts, and the mean is 2226 counts. 20 msec integration time.**

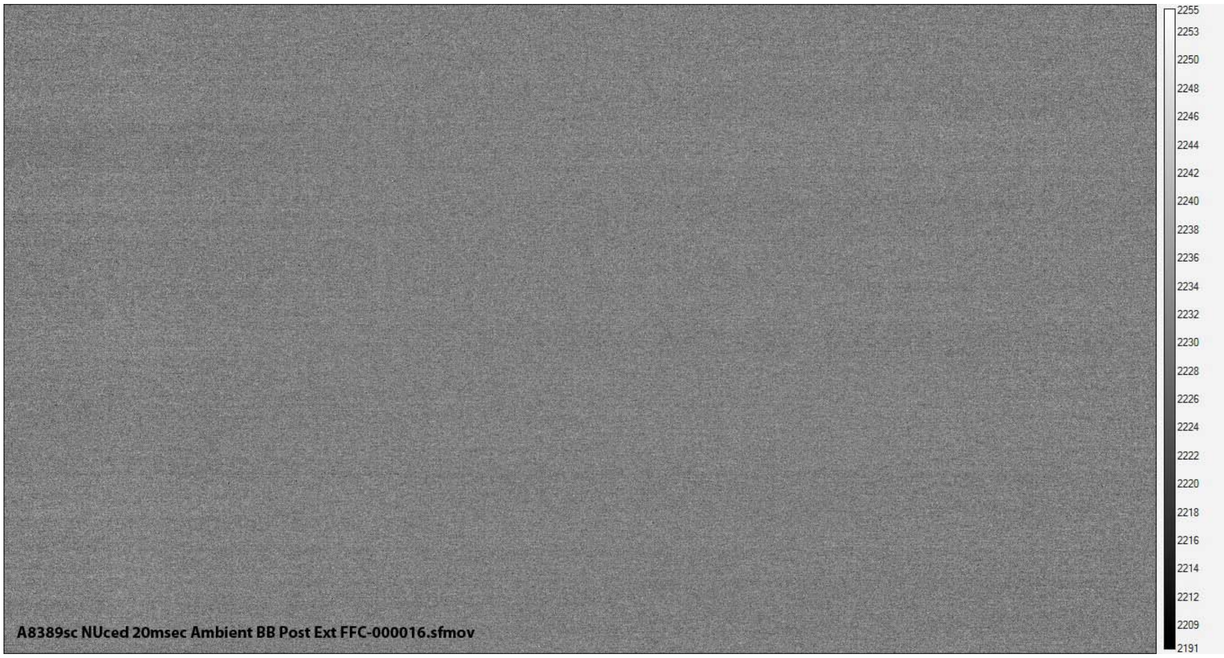


**Figure 20c. Ambient uniform source right after external offset update which slightly cleans up the already very low fixed-pattern noise. The standard deviation has dropped to 2.9 counts, and the mean is 2231 counts. 20 msec integration time.**

## Summary

Here are some key takeaway points from this chapter:

- IR bandpass filters can be helpful to make a material opaque, or see through a material, or to "see" some gasses.

- Using a warm filter to image ambient temperature scenes (with accordingly long integration times) is generally a bad idea due to camera internal self-radiation or "CISR" which causes level offsets that eat into the dynamic range of the camera and also non-uniformities due to narcissus effects.

- Warm filters work about as well as cold filters when the target of interest has orders of magnitude more emitted in-band radiation compared to the CISR and the integration time can be set to something short. This condition is met with ND filters used for high-temperature calibrations.

- For high temperature targets, it is always preferable to limit the sensor exposure by shortening the integration time, rather than using a less transmissive filter for exposure control. Put another way, an ND2 filter and 1msec integration time will work better than an ND3 filter and 10 msec integration time. The shorter integration time will limit the CISR effects like level offset and narcissus-induced non-uniformity.

- Almost all the FLIR Niceville science cameras support warm filter use in either 4-position filter wheels or "behind the lens" holders, depending on the camera model. The exceptions are the RS series cameras for range work.

- Custom cold filters can be incorporated into most cooled, photon-counting FLIR Niceville science cameras. The camera is built around this filter, and it can't be changed once the camera is assembled without a risky and expensive rework.

- Frequent external update offsets may be required to improve camera performance when using warm filters as the camera internal and warm filter temperatures change.

# Chapter 6 : Image Noise in Infrared Camera Systems


**Noisy InGaAs camera image of man in low-light conditions, Canyonlands National Park, Utah**

# Contents

Chapter 6 : Image Noise in Infrared Camera Systems .......... 6—279
Introduction .......... 6—281
Spatial Noise .......... 6—283
Temporal Noise .......... 6—290
Scene Noise .......... 6—291
Dark Current Noise .......... 6—291
Read Noise .......... 6—292
Electronics Noise .......... 6—292
Stray Light Noise .......... 6—292
InGaAs Camera Noise .......... 6—293
Scene Noise .......... 6—294
Dark Current Noise .......... 6—294
VisGaAs Dark Current .......... 6—296
Readout Noise .......... 6—296
Electronics Noise .......... 6—297
Dark Current .......... 6—297
Lattice-matched InGaAs, Extended InGaAs and VisGaAs .......... 6—298
3D Noise .......... 6—299
Temporal Noise Components .......... 6—299
Spatial Noise Components .......... 6—300
Summary .......... 6—301

## Introduction

Image noise is one of the most asked-about specifications of infrared camera system performance. The terminology that is used to describe noise is very prone to being misused, which can create a lot of confusion. This chapter will talk about the causes of image noise, how it is measured, and how it affects image quality and radiometric accuracy. Along the way, the terminology used to describe aspects of image noise will be defined.

There are two main classes of noise that affect infrared imaging systems:

- Temporal Noise
- Spatial Noise

Temporal noise is noise that changes with time. Temporal noise affects each pixel individually in an imaging sensor array, but it can also simultaneously affect all the pixels in the whole image ("frame bounce"). Temporal noise is the most noticeable noise in a system that has a good non-uniformity correction. When a well-NUCed camera is pointed at a very uniform IR source, the image looks a lot like the "snow" that appeared on older TV sets when they were tuned to a dead channel. Each pixel's digital count value is fluctuating randomly with time around a mean value. If the fluctuations in the pixel count values are not correlated with each other, then there is no fixed spatial texture to the pattern. This is pure temporal image noise and it is shown in Figure 16 below. The image is a single frame from a sequence of 64 frames taken while the camera is pointing at a black-painted metal plate with a special roughened surface at ambient temperature. This is a very, very uniform scene, especially since the plate surface itself was located only 1cm in front of the lens, and thus is way out of focus.

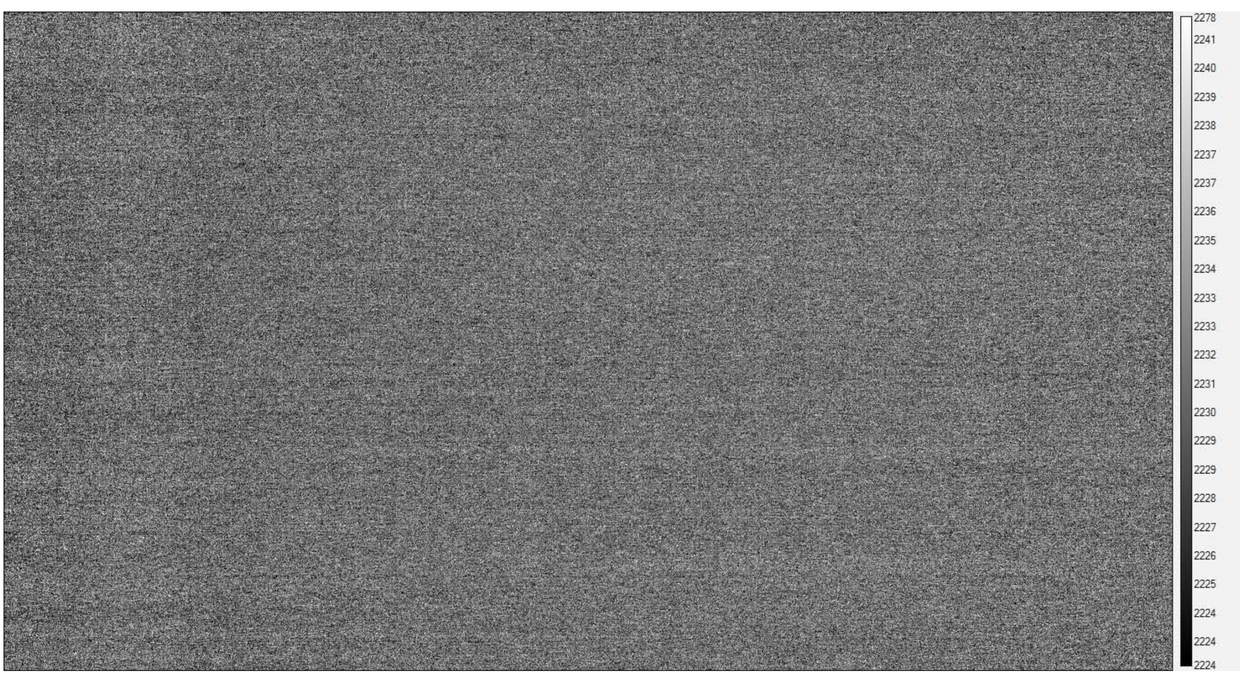

**Figure 16. A well-NUCed midwave IR camera image of a 23C ambient uniform scene. The image texture is caused by random temporal noise of the detectors and can be described as "Snow" or "Salt and pepper". If we neglect bad pixels, the image has a narrow distribution of digital count values, in this case varying from 2187 to 2278 counts. This image is one of a series of 64 frames taken at 30 frames per second. The integration time is 20msec.**

Another hallmark of random temporal noise is that the digital count values will follow a normal distribution, and the histogram of the whole image will have a gaussian distribution, as shown in

Figure 17:

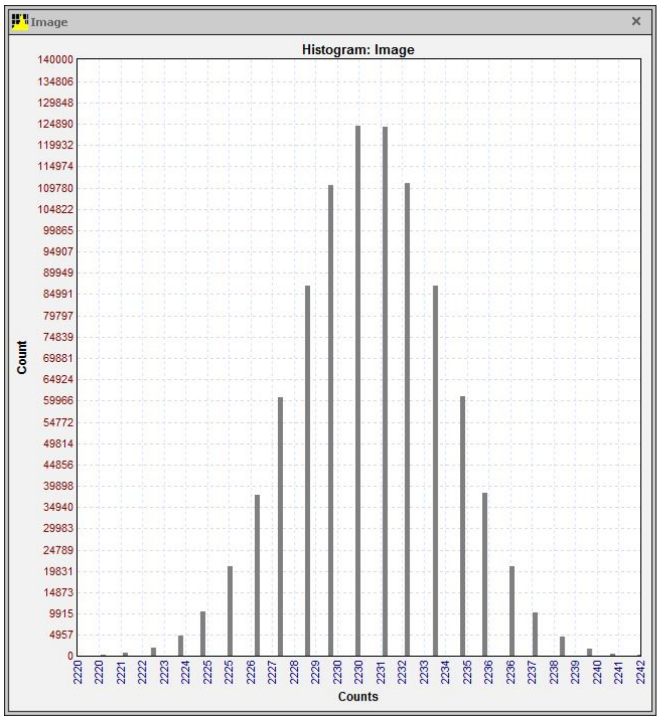


**Figure 17. Histogram of a single frame of a 64 frame NUCed image sequence showing a normal distribution of digital count values for all the pixels in the image. The standard deviation is 2.9 digital counts. This image is very well NUCed – the spatial noise is low.**

A single pixel exhibits random fluctuations in its count value over time. This is due to the randomness of photons arriving in any given time interval. This noise has various names: shot noise, quantum noise, scene noise. Figure 18 shows a time plot of a single pixel over the 64 frames. The pixel's count value fluctuates between 2225 and 2236 counts. Every pixel in the array is doing the same thing, with the exception of bad pixels, which may be open, shorted, or exhibiting non-random behavior ("twinkling").

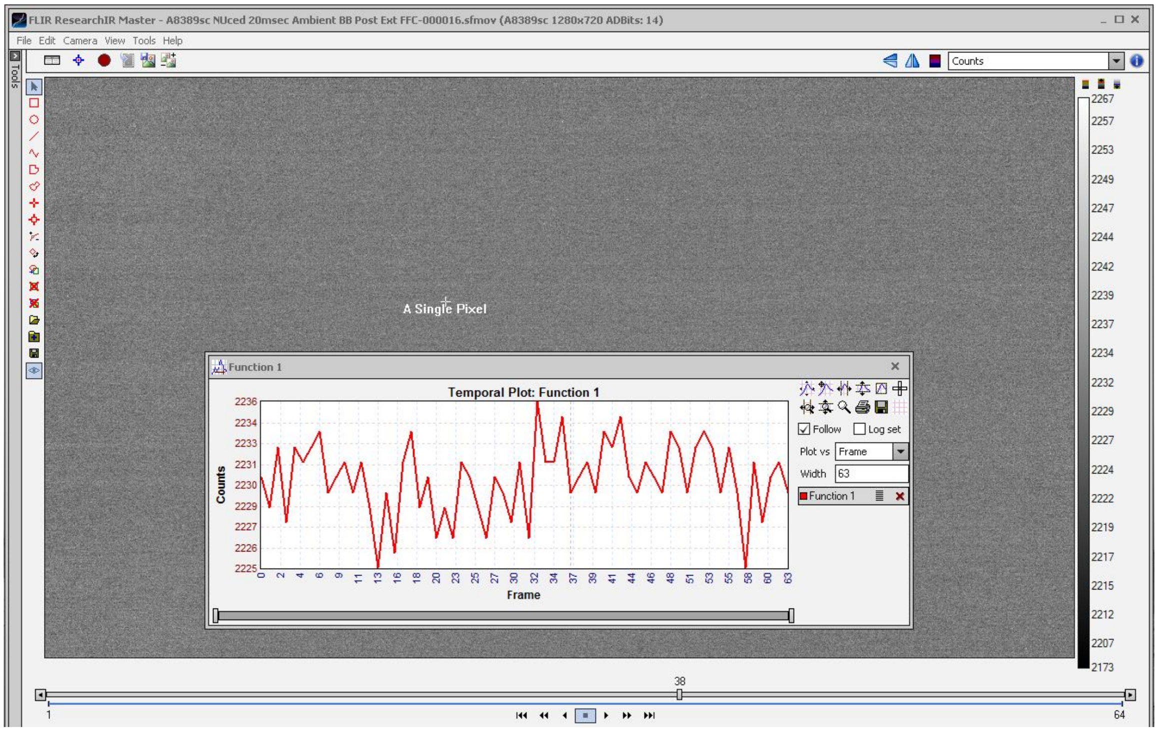


**Figure 18. Time plot of a single pixel over 64 image frames. The pixel value stays between 2225 and 2236 counts.**

## Spatial Noise

Spatial noise, also called fixed-pattern noise, is noise that does not change with time (although it may change on slow time scales due to thermal drift of the optics). A good example of spatial noise is the irregular image appearance of a thermal infrared camera that is pointing at a uniform IR source like an area blackbody or an integrating sphere exit port with the non-uniformity correction (NUC) turned off. An example of this is shown in Figure 19, which is a single frame from a 16-frame sequence. The histogram of this image, shown in Figure 20, shows a wide distribution (~3000 counts at the base) and an asymmetric shape. The spatial noise in the image has a number of components to it. The coldshield shading makes the corners darker – this is a radial gradient in the image caused by the coldshield geometry and having a $\cos^4(\theta)$ shape to it. The column to column variation is caused by differences in the ROIC column amplifier offsets. There are some bad pixels – they are either shorted and look white or are open circuits and look black. The standard deviation is 515 counts. This image has a lot of spatial noise in it.

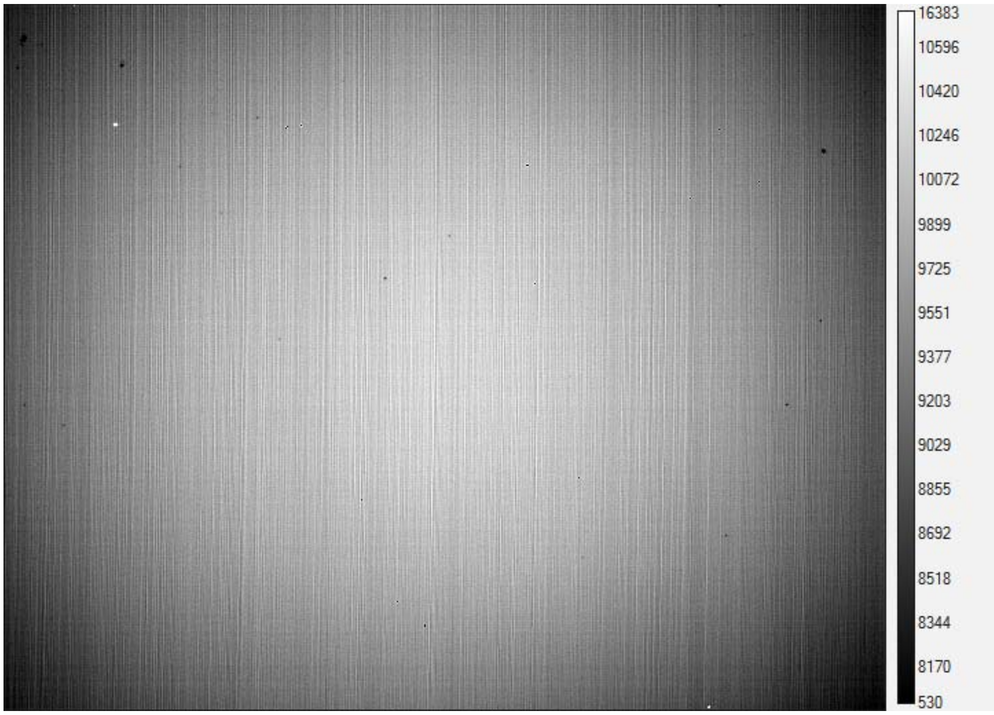


**Figure 19. Midwave IR camera with the NUC turned off pointing at a 25C uniform blackbody source**

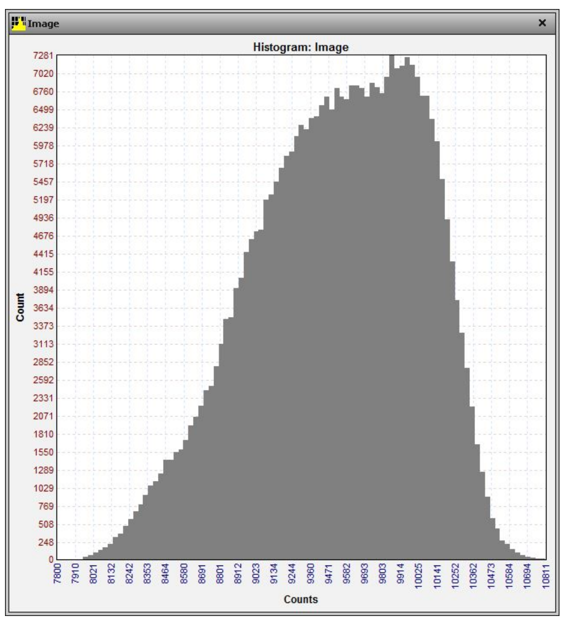


**Figure 20. Histogram of image made with a midwave IR camera with the NUC turned off pointing at a 25C uniform blackbody source. The histogram distribution is asymmetric with a standard deviation of 515 counts and is ~3000 counts wide at the base.**

Figure 21 is the same image with a NUC applied. The spatial noise decreases, and the image statistics changes drastically. The standard deviation of a single frame is now 6.0 counts.

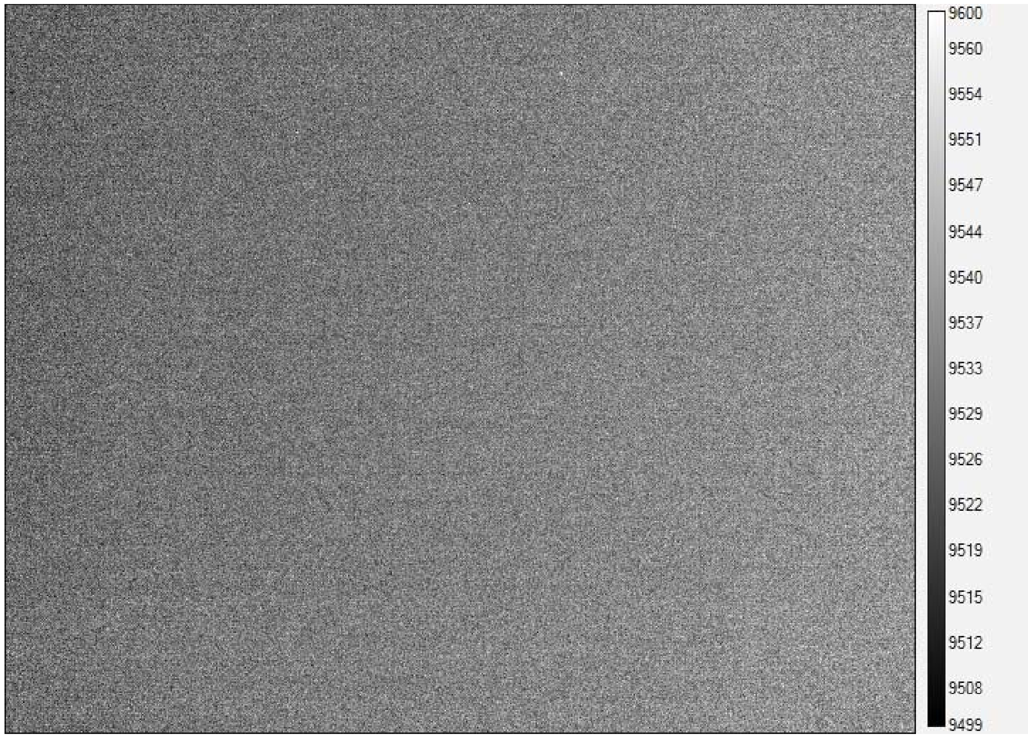


**Figure 21. The same image as Figure 20, but now with a high quality NUC applied.**

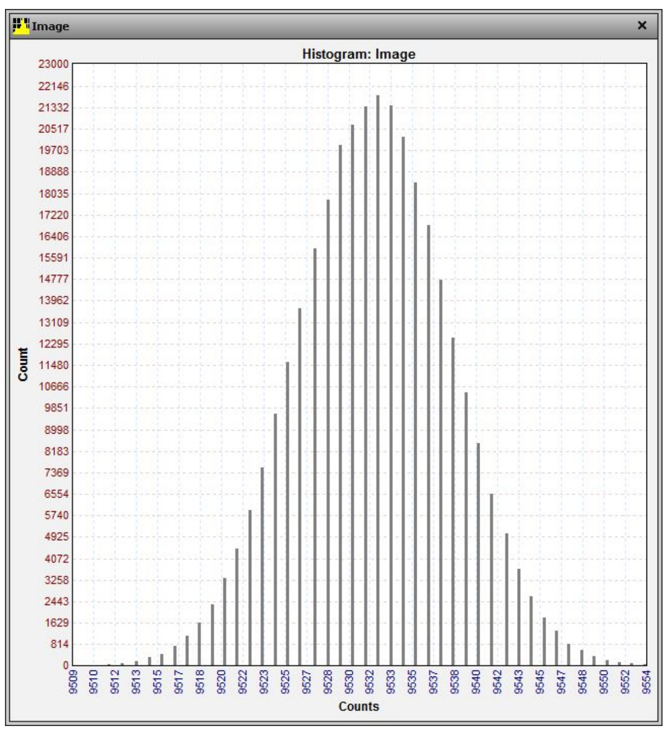


**Figure 22. The histogram of the image in Figure 21. The standard deviation is now 6.0 counts.**

I deliberately set the length to be 16 images for this image sequence. ResearchIR software has an image filter set which includes a frame average with a maximum depth of 16 frames. I applied that to the 16-frame sequence, and get the frame-averaged image shown in Figure 23:

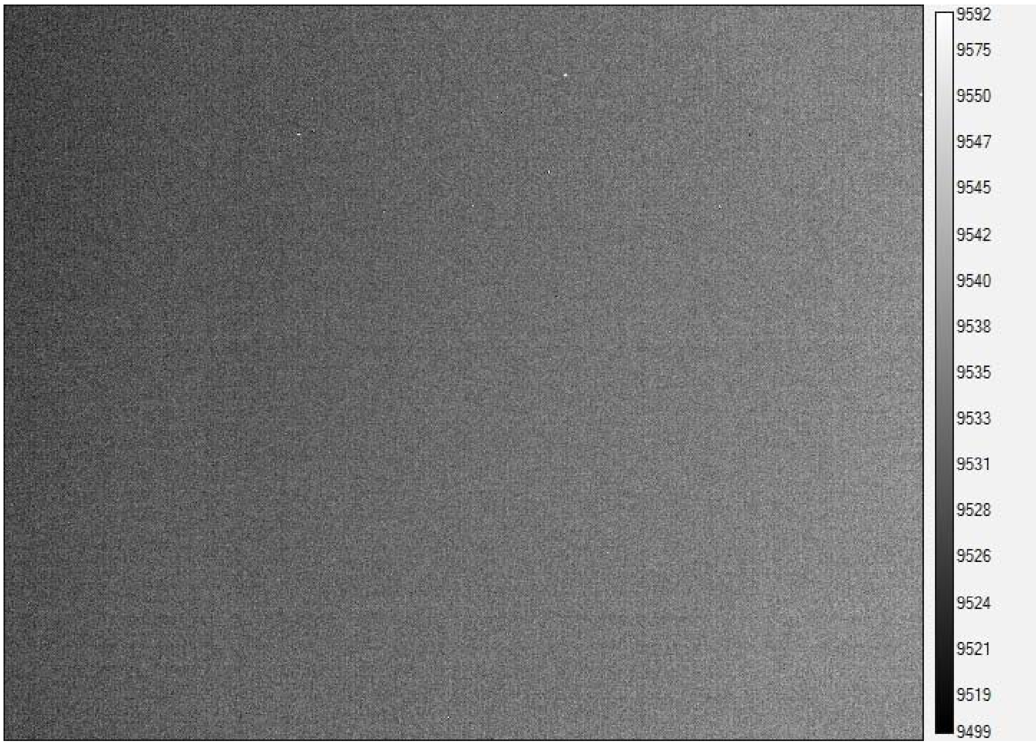


**Figure 23. The average of the 16-frame image sequence. The temporal noise component has dropped down considerably.**

The histogram of this image has a standard deviation of 3.5 counts, down from 6.0 counts for a single frame from the sequence. Frame averaging works well out to long image sequences. Eventually, though, if one averages for hundreds of frames, one starts to get optics temperature drift effects and other noise sources, like "1/f" or flicker noise begins to affect the image quality. Flicker noise has a power spectrum that scales with the inverse of frequency. We have found that a length of 64 frames is a good value for averaging to reduce temporal noise, or to measure noise equivalent delta T.

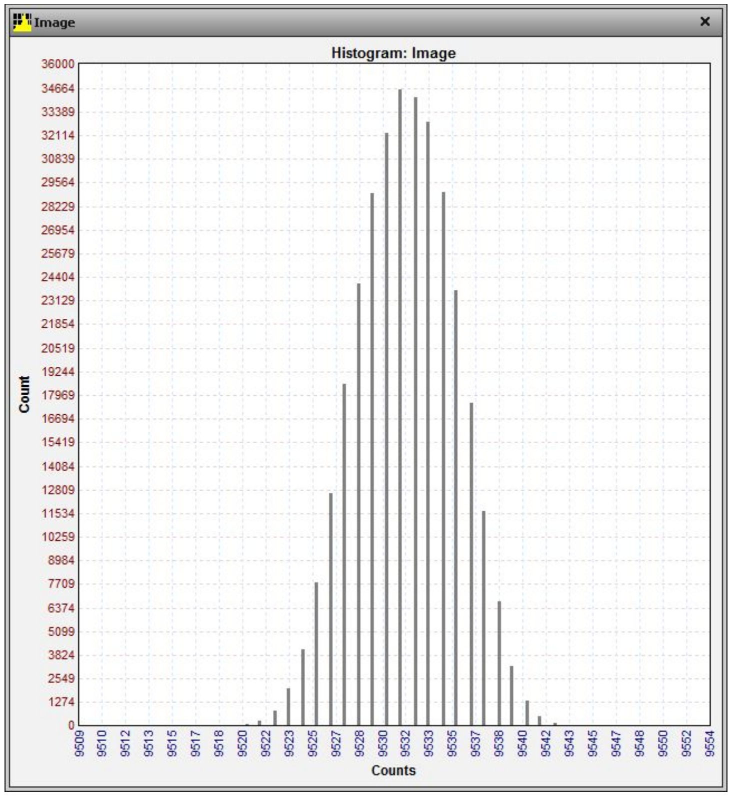


**Figure 24. The histogram of the 16-frame average image in Figure 3c. The standard deviation is now 3.5 counts, down from 6.0 counts for a single frame from the sequence.**

Here is another example of frame averaging at work. The following images in Figure 25, Figure 26 and Figure 27 were taken from a 64-frame sequence of images of a 25C laboratory blackbody taken with a 640x512, 25-micron pixel InSb camera with a 25mm lens. The camera had been previously NUCed with a 20C and 30C blackbody sources.

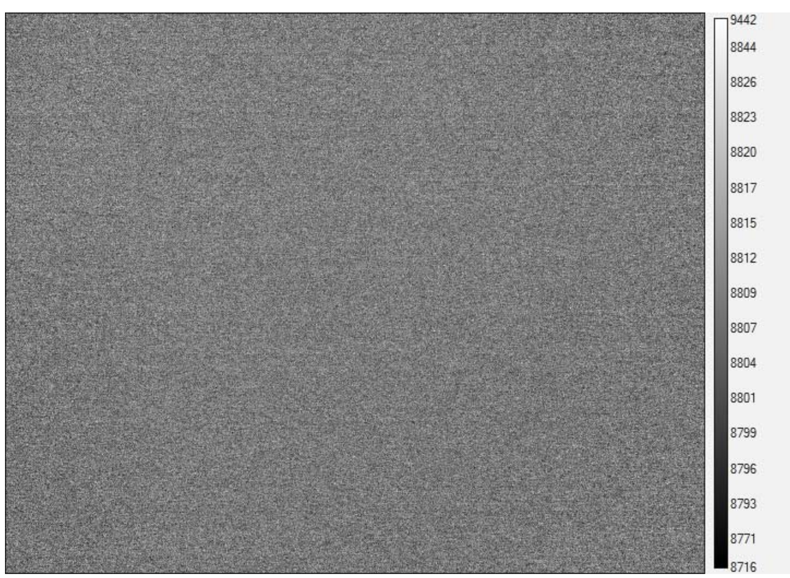

**Figure 25. This is Frame 1 of the sequence. The standard deviation is 3.6 counts.**

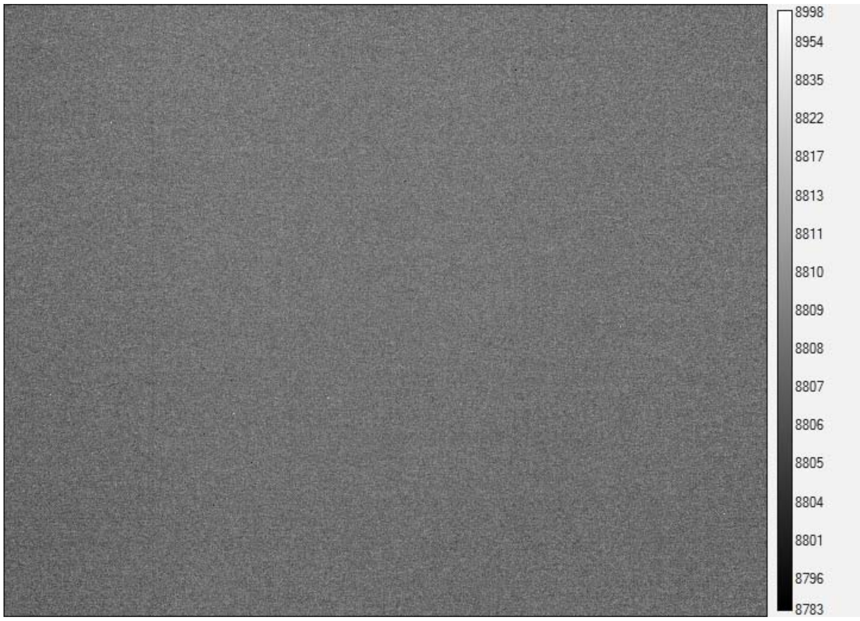


**Figure 26. This is the average of the first 16 frames in the sequence, and one can see that the noise has noticeably decreased. The standard deviation is 1.0 counts, ~4 times smaller than the single frame standard deviation. The noise should decrease as the square root of the number of the frames in the sequence. 16 frames averaged together will lower the noise by a factor of 4 over a single frame, which is close to what was observed.**

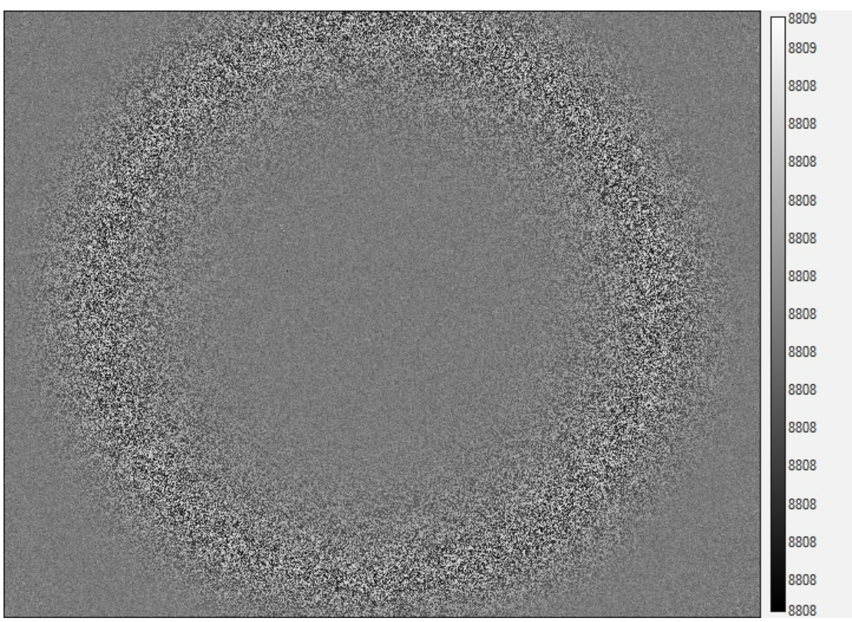


**Figure 27. This is the average of the 64 frames in the sequence. The standard deviation is 0.15 counts, which is ~24 times smaller than the single-frame image standard deviation. The AGC makes the image look "bad" but it is actually very, very uniform. I would not have expected the standard deviation to be this low, but then there are "wings" to the distribution that make it non-gaussian. There is a circular artifact in the image which is probably caused by reflections in the lens and/or cold shield.**

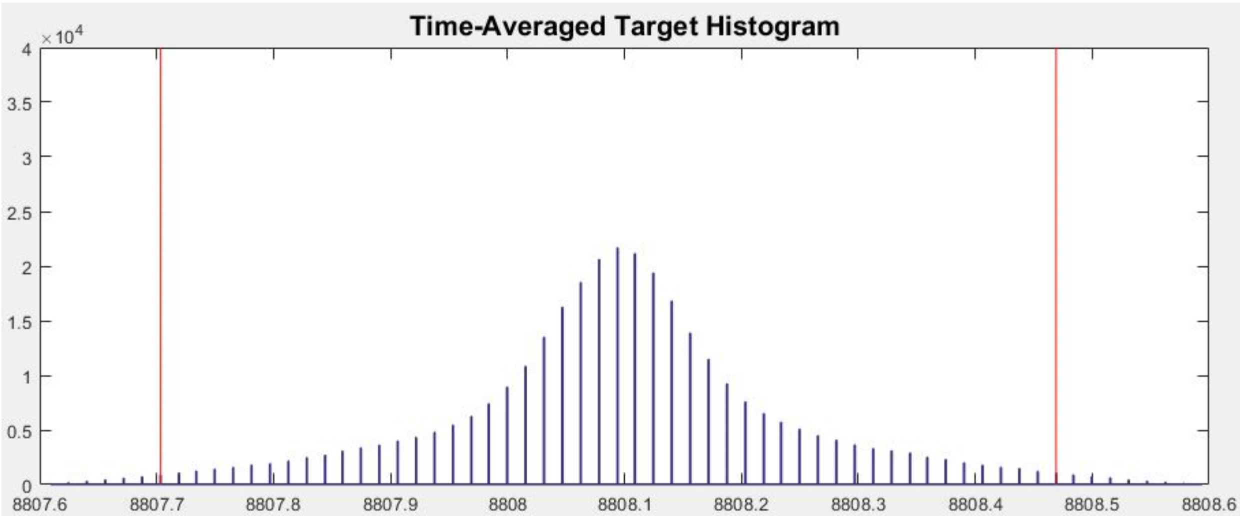


**Figure 28. The histogram of the 64-frame average image.**

## Temporal Noise

Most midwave science camera applications are happening in the so-called background-limited regime where the image noise that you see on the computer screen is dominated by random fluctuations in the scene radiance – this is called shot noise, quantum noise or scene noise. There are other noise sources at work here, some of which are talked about all the time, like read noise, and some which few camera users ever consider, like stray light in the DDCA.

All the noise sources in a typical IR camera system are assumed to be both random (in terms of their statistics) and independent of each other (the readout noise does not change with scene noise, for example). Therefore, these noise sources add like vectors that are orthogonal to each other. If they were not independent then they would just sum up, but instead they are on average out of phase with each other by 90 degrees). This noise adding method is called adding in quadrature – you sum the squares of each noise source and then take the square root of the sum. This method is also called RSS, which stands for root-sum-square: you sum the squares and take the root of the sum. You can think of it like the Pythagorean Theorem where two perpendicular sides of a triangle add in quadrature to give you the hypotenuse: $C^2 = A^2 + B^2$. In this case, we have an n-dimensional vector that is composed of n random and independent noise components. It turns out that if one of those noise sources is substantially larger than the others, then it will dominate. Key point: the scene noise typically dominates the system noise for many midwave imaging applications. This domination effect is just like the case of a right triangle with one long side and one short side. The hypotenuse is going to be very close to the length of the long side and the short side's length has little effect on the hypotenuse.

You thus have various noise sources that are "RSSed" together to get the system noise:

1. Scene Noise
2. Dark Current Noise
3. Read Noise
4. Electronics Noise
5. Stray Light Noise

Here is a brief description of each noise source in the typical midwave IR InSb camera system:

## Scene Noise

The scene noise is typically about 2000 electrons for a 0403-based sensor at mid-scale. The scene noise is the square root of the mean photoelectrons generated by radiation from the scene. Why the square root of the mean? Since these are photons which are generated independently of each other, we assume a Gaussian or normal statistical distribution here; you can read more about the normal distribution on Wikipedia – it is a cornerstone of statistical theory. A 2000 electron scene noise is what you get when there are 4 million photoelectrons in the integration capacitor in each pixel unit cell. We always tell customers to run the integration time to put the mean of their scene at midwell, i.e. about 8000 counts, so it is a common situation to have about 4 million photoelectrons in each pixel's integration capacitor. For an A6700 camera with the 0403 readout, midwell is about 4 million electrons and the square root of that is 2000.

## Dark Current Noise

Customers who are attempting to model the performance of their camera often ask us what the dark current is in our InSb detectors. This value is not something we can readily disclose to customers, as it is highly proprietary. The dark current noise is the square root of the electrons of dark current integrated during the integration time of the sensor. It is the square root of the "dark charge", the dark charge being how many electrons of dark current are collected per frame. Like scene noise, dark current noise is the square root of the mean; we assume a normal statistical distribution. By the way, no one else but me calls it dark charge, but I think it is a useful concept. You can have higher dark current (like if you ran your InSb at 87K instead of 77K) and not get much dark charge at all if you are imaging a hot scene at a short integration time. The dark current noise is generally small for InSb running at 77K temperature unless you are running at a really long integration time. If the sensor is imaging a hot scene and the integration time is 10 microseconds, then almost no dark current is integrated, and the dark noise is tiny and can be completely ignored.

The only time FLIR cameras ever run InSb at tens of milliseconds integration time is when there is a narrow band cold filter and/or a very high f/number coldstop. The GasFinder filter for methane imaging is a narrow filter and that is why a special GasFinder A8303 camera we built has to be run at 115 msec to get to midscale digital counts when it is pointed at a 25C scene temperature. The special A8303 gas camera is also f/4, which further limits light from the scene.

Dark current matters in that case. You cannot run a 3-5-micron InSb camera with an f/2.5 coldshield at 20 msec even if you point the camera into a bucket of liquid nitrogen, because the pixels will be flooded with radiance emitted from the lens itself and the warm window. Try it yourself!

## Read Noise

Read noise is short for readout noise. The camera sensor consists of InSb detectors hybridized onto a readout IC, which is a very complex mixed-signal integrated circuit. It is the most challenging design for any part of an infrared camera. Every time the sensor collects a frame of image data, there is some noise associated with the storage of the photoelectron charge on the integration capacitor in each unit cell in the readout, and its transmission as a voltage to the output of the readout. The read noise for typical readouts we make is about 500-800 electrons. It is 642 electrons for the 0403 ROIC used in the A6700 series.

## Electronics Noise

The typical RMS electronics noise is ~800 electrons for a 0403-based camera. This value is computed by disconnecting the sensor from the electronics and just measuring the noise from the electronics itself. An RMS noise of 800 electrons translates into about 1.7 digital counts. This noise is partly due to the digitizer that converts the voltages from the sensor into digital counts values.

## Stray Light Noise

Stray light noise is caused by leakage of light into the dewar, typically around the electrical feedthroughs that get signals in and out of the sensor. This is typically a small amount of light and just adds to the scene noise, though it is not correlated with the scene radiance itself. The light is emitted from the ceramic materials on the motherboard upon which the sensor is mounted. Those ceramics might be at 40C for a warm camera body. For an integration time around 2 msec, the noise from stray light is a small effect. At short integration times, it becomes a very small effect. The only time it becomes important is the same regime where dark current becomes important: long integration times, which you will only get to if you have either a narrowband cold filter, a high f/number coldstop, or both.

## InGaAs Camera Noise

The A6261sc camera has a focal plane array built on the ISC1202 ROIC. Currently, the InGaAs science cameras built by the Niceville division of FLIR use a readout integrated circuit called the ISC1202 or 1202 for short. The 1202 has three gain states with well capacities and readout noise values shown in Table 1:

| ISC1202 ROIC Gain State | Well capacity (e-) | Readout Noise (e-) | Conversion Gain (electrons/count) |
|---|---|---|---|
| Low | 1.44E6 | 397 | 96 |
| Medium | 9.57E4 | 38 | 6.6 |
| High | 1.91E4 | 11 | 1.6 |

**Table 1. 1202 ROIC characteristics**

In this section, I will discuss the noise behavior of the 1202 with FLIR's InGaAs detectors hybridized to it. The noise performance varies dramatically with the gain state that the camera is operated in. Most InSb focal planes are designed with fairly large wells and a single on-chip gain state. InGaAs ROICs are different, because SWIR scene dynamics are so different. For example, the range of ambient outdoor scene illumination in the SWIR band changes by 6 orders of magnitude from day to night. Unlike MWIR scenes, SWIR scenes can have zero in-band radiation, so the readouts are typically designed to operate down at extremely low light levels. In order for them to give meaningful image data, their readout noise values have to be quite low, and they use very different ROIC unit cell architectures compared to MWIR InSb ROICs.

There are a variety of noise sources in an InGaAs camera. They are as follows:

1. Scene Noise: The square root of the number of electrons in the well of a given pixel.
2. Dark Current Noise: Dark current is a leakage current which increases with detector temperature. The current is integrated during integration time alongside of the photocurrent from the scene. It is present even when no SWIR radiation is striking the detectors, hence the name "dark current". The square root of the integrated dark current is the dark current noise.
3. Readout Noise: The readout IC in the focal plane has intrinsic noise caused by switching of current
4. Electronics Noise: The focal plane outputs voltages from each pixel by row and column, all synchronized with the pixel clock in the camera electronics. These voltages are transmitted to the

digitizer inputs through a "flex cable". A certain amount of noise is present even without a focal plane being there. This noise is usually measured in digital count units and is about 1.7 counts typically.

These noise sources are (generally) uncorrelated with each other, and so they are added in quadrature, or by the root-sum-square or RSS method. The resulting total noise can be thought of as a vector which has four orthogonal components.

$$\mathbf{Noise_{total} = RSS(shot\ noise,\ read\ noise,\ dark\ current\ noise,\ electronics\ noise)}$$

$$\mathbf{Noise_{total} = Sqrt(Scene^2 + Dark\ Current^2 + Read^2 + Electronics^2)}$$

## Scene Noise

In darkness, the shot noise is 0. No photons, no noise. For a 1202 at midwell, there will be 7.2E5 electrons, which translates to 849 electrons. This noise source will typically dominate over the other noise sources, unless the dark current is very high due to a high sensor operating temperature, or a long integration time, or both. If the scene has enough background to get to midwell at 1msec integration time, then the system will definitely be background limited.

## Dark Current Noise

The dark current specification for FLIR's A6261sc InGaAs camera sensors is about 1.2e-14 amps at their 30C operating temperature, which is 7.6E4 electrons/sec. For a short integration time, the dark current is very low. At the minimum integration time of 50 microseconds for the 1202 ROIC, the integrated dark current is only ~4 electrons. The noise on that dark charge[19] is the square root of 4, which is only 2 electrons. This is negligible compared to all other noise sources.

At a long integration time like 50 msec, the dark charge is 4000 electrons and the noise is 63 electrons. That is a significant amount of both noise and well fill, particularly for the smaller well capacities of medium and high gain modes of the 1202. In order to image low background scenes, the dark current really needs to be reduced as small as possible, which would mean cooling the focal plane to colder temperatures. The dark current drops by a factor of ~2 for every 7C decrease in temperature. The A6261sc cameras are currently operating at a 30C setpoint, which limits imaging performance for very low backgrounds. We have operated the A6261sc at 200msec integration time in medium gain mode and got reasonable image quality considering the very low backgrounds. Figure 14a was taken in Canyonlands National Park on partially overcast no-moon conditions with an A6261sc camera running at 200msec integration time in

[19] I call integrated dark current "dark charge" but no one else does, though they ought to, since saying "dark current" does not really describe it!

medium gain mode with an f/1.4 SWIR lens.  There are very low levels of man-made light in these imaging conditions and all the exposure "knobs" have to be turned up all the way.


**Figure 14a.  Low-light InGaAs camera image of man in no-moon conditions, Canyonlands National Park, Utah. Integration time is 200msec.**

Figure 14b is an image of the same scene, but with an integration time of 15msec. In this case, the signal is lower by a factor of 200/15 or 13.3.  The man is barely visible.  If an observer is watching a video of this scene, and the man moves, it will be far easier to detect him, but if he holds still, he will be very hard to detect, especially by an automated video analytics system.


**Figure 14b. Low-light InGaAs camera image of man in no-moon conditions, Canyonlands National Park, Utah. Integration time is 15msec.**

## VisGaAs Dark Current

The A6262sc VisGaAs cameras have about 2X the dark current of the A6261sc InGaAs cameras. The thinning process that goes into making VisGaAs tends to generate surface defects that act like carrier traps and increase dark current.

## Readout Noise

Read noise values for the A6261sc Readout IC are 230 e- in low gain, and 29 electrons in medium gain. High gain read noise measures out to 30 electrons, which give it no advantage over medium gain. We don't like to use high gain for another reason: since we are running the InGaAs at 30C and the integrated dark current is too high for such a tiny well, dark current fills up the wells too quickly. The readout noise is fairly insensitive to integration time and sensor temperature. If the camera sensor is in high gain mode, the scene background is very low, and the dark current noise is very low, then readout noise will dominate over the other noise sources. The electronics noise is always less than the readout noise by design. If it wasn't, then the designer is using the wrong digitizer!

## Electronics Noise

Electronics noise is 1.7 counts. One can convert this to electrons by multiplying times the conversion gain of the focal plane in each gain state, as shown in Table 1 below. The electronics noise values are:

| ISC1202 ROIC Gain State | Conversion Gain (electrons/count) | Electronics Noise (e-) |
|---|---|---|
| Low | 96 | 163 |
| Medium | 6.6 | 11 |
| High | 1.6 | 2.7 |

**Table 2. Conversion gains and electronics noise levels for the 1202 ROIC**

## Dark Current

We often get asked about dark current for both InGaAs and InSb detectors. Customers want to know what the dark current value is, often because they are modeling the camera system and want to understand the various sources of noise in the system. But they also want to know other performance specifications, so in this white paper, I will describe some of those other specifications as well for InGaAs, InSb and SLS detectors.

Dark current means an electrical current that flows in a photodiode in the absence of any stimulation by EM radiation, i.e. the photodiode is in the dark. Dark current is a phenomenon associated with photodiodes that are being reversed biased. Reverse bias means that a voltage is applied across the diode in a reverse polarity from what would make the diode conduct like a conductor. In a perfect photodiode, there would be no dark current at all. In real photodiodes, there is always some current flow, which is also called ***leakage current.***

The dark current is sensitive to the temperature of the diode junction. At higher temperatures, there are more thermally induced carriers in the conduction band of the diode, and there will be increased dark current. As the reverse bias voltage is increased, there will also be more generation of dark current.

In an InGaAs photodiode, the typical rule of thumb is that the dark current will double for every 7 C increase in junction temperature. For InGaAs camera applications in low light, it is important to reduce the dark current to small values for two reasons. One is that the dark current fills up the integration capacitor and reduces the dynamic range of the camera at a fixed integration time. I call this time-integrated dark current “dark charge”. The other is that that additional well fill causes additional shot noise. If you take the square root of the dark charge in electrons, then that is the shot noise from the dark current.

The dark current density for FLIR InGaAs detectors at a 20C operating temperature has a typical value of <2 nA/sq. cm. A selection of four test reports for production sensors showed dark current values of 0.98, 1.35, 1.27, and 1.22 nA/sq cm at 20C.

A dark current density of 2 nA/sq. cm translates to 28,000 electrons per second of dark current in a 15-micron pitch pixel, which is the pixel pitch of FPAs built on our ISC1202 ROIC. At 30C, the sensor operating temperature of the A6261sc camera, this dark current value increases by $2^{(10/7)}$, since the dark current doubles with every 7C increase in temperature. This 30C dark current value works out to 75,000 electrons/second.

Dark current values can vary quite a bit from sensor to sensor because we set the maximum dark current spec value to be pretty high. FLIR specifies a maximum dark current value of 10 nA/sq. cm at 20C. I have measured an A6261sc engineering camera and got a dark current value which was 13.88 nA/sq cm at 30C operating temperature. This translates to 5.2 nA/sq. cm at 20C. We are generally delivering sensors with dark currents that are much lower than the 10 nA/sq. cm at 20C specification.

## Lattice-matched InGaAs, Extended InGaAs and VisGaAs

Our InGaAs is ***lattice-matched*** material, which means that there is a particular ratio of indium, gallium and arsenic to give the crystals the same lattice constant as the indium phosphide substrate that the InGaAs material is grown on. This reduces strain in the interface which means that one tends to get fewer defective pixels caused by crystal defects and dislocations that one does in a different alloy that is not lattice matched.

Lattice-matched InGaAs has a cutoff wavelength of about 1.7 microns at room temperature. FPAS with detector made of other alloys of InGaAs have been fabricated. They have decreased bandgap energy, and therefore longer wavelength cutoffs, and the material is called **Extended InGaAs**. These sensors have significantly higher dark current since it is easier for thermally induced carriers to occur when the bandgap energy is lower. They also have lower operability, because there tend to be more crystal dislocations.

VisGaAs focal plane arrays are built with lattice-matched InGaAs, and the InP substrate is thinned down to increase the QE of the FPA to shorter wavelengths of radiation, down into the visible band, which is how the name was derived. The typical lattice-matched InGaAs sensor will have a steeply falling QE curve below 950nm wavelength.

VisGaAs dark currents are typically 2X higher than InGaAs dark currents, that is, 4 nA/sq cm at 20C. VisGaAs has higher overall quantum efficiency (~70%) since the InP substrate has been thinned down, removing a layer that tends to absorb some radiation before it enters the active InGaAs layer. VisGaAs FPAs tend to have lower operability (99% instead of 99.5%).

**Dark Current Scenarios:**

1) For an InGaAs camera imaging during the day, there is typically plenty of signal and well fill will be 50% or greater, even at ~few msec integration times, and so the dark current noise is negligible compared to the scene noise. This is the background-limited case.
2) In low-light situations where one is forced to image at long integration times, the dark current will fill up the well and add significantly to the system noise.

## 3D Noise

The previous discussion of noise has divided it into two types of noise: spatial and temporal. A complete analysis of image noise includes spatio-temporal noise components, and these values make up what is known as the **3D Noise Model**. These are time-varying noise components that have spatial correlation. The most common manifestation of this type of noise is in uncooled camera images, where the signal levels of individual rows and columns of pixels can all "move together" up and down in a random manner.

In all, there are 8 components in the 3D noise model. These components are:

| Noise Component | Pixel Variations | Row Variations | Column Variations | Frame Variations |
|---|---|---|---|---|
| **Temporal** | Sigma_thv | Sigma_tv | Sigma_th | sigma_t |
| **Spatial** | sigma_vh | sigma_v | sigma_h | Sigma_s = 0 |

### Temporal Noise Components

The most dominant noise component is sigma_thv, which is known as random 3-D noise. It is the temporal variation in individual pixel signal levels, with no correlation across time or array position. This noise component is predominantly caused by shot noise from the scene, that is, random arrival times of uncorrelated photons on each pixel. An analogy is what happens if one pours a bucket of ball bearings onto a chess board. In any given time interval, the number of ball bearings hitting a square of the chess board will vary, and the variation has a magnitude which is the square root of the mean number of bearings per time interval per chess board square. The same thing happens with photons; though one might have a flux density of 10,000 photons/sec/pixel incident on an imaging sensor, in any 1-second interval, the number of photons can vary, with the typical 1 sigma variation being the square root of 10,000 or 100. A typical photons/sec/pixel value would be 10,000 +/- 100 photons, but sometimes, the pixel could receive 10,200 photons, or 10,300 photons with those 2-sigma and 3-sigma events being increasingly rare.

This is the noise component that is the simplest to measure, and it is measured at the factory for all InSb and SLS science cameras. Its value is usually expressed as NEdT or noise equivalent delta T. The measurement involves the following steps:

1. A test camera has its lens installed and the camera is pointed at a 25C blackbody
2. The test camera's integration time is adjusted so that its mean array count value is ~8000 counts
3. A sequence of 100 frames is taken of the 25C blackbody
4. A proprietary script is run to compute the sigma of each pixel in the array over the 100-frame sequence. These sigma values are averaged to get an array temporal noise value expressed in digital counts.
5. The camera is then pointed at a 20C and a 30C blackbody in succession, with 100 frames being acquired at each blackbody
6. Each of these sequences is frame averaged, and the resulting average frame is itself averaged to get a digital count value.
7. The difference in these two count values is computed, and this value is divided by 10C, the temperature difference between the two blackbodies. This value is the responsivity of the array in counts/degree.
8. The temporal noise is divided by the responsivity to yield the noise in temperature units. A typical value for NEdT might be 20 milliKelvin, written as 20mK for short. This is the same as 0.02C, since K and C units have the same magnitude, just a different offset.

The sigma_tv and sigma_th components are variations in array row averages and array column averages that change from frame to frame. The former is sometimes called image "streaking" and the latter is called image "rain". These noise components are typically smaller than sigma_thv unless there is a problem with the camera.

The final noise component in the set is called sigma_t, also known as "image flicker" or "frame bounce", where all the pixels in an image sensor move together, which leads to frame-to-frame variations in the image intensity. This noise is generally small compared to sigma_thv, and it is often ignored. I have seen frame bounce be an issue when a focal plane detector bias power supply had an unwanted oscillatory behavior.

### Spatial Noise Components

Sigma_vh is random fixed-pattern noise, that is, noise that does not change with time and is not correlated with other pixels. The non-uniformity correction is designed to minimize this noise component. The best approach to minimize fixed-pattern noise is to start with a high-quality 2-point NUC, one that was generated with a pair of very uniform blackbodies, as the factory NUCs are. The camera should have the lens intended for use installed and then the camera should be powered up while in a stable temperature environment and allowed to stabilize its internal temperature. Finally, right before image data is taken, an external offset update should be performed using a uniform radiance blackbody (often an ambient temperature metal plate). Right after the external offset update, the fixed-pattern noise should be extremely small, though it will tend to slowly appear over time.

Sigma_v and sigma_h are variations in row averages and column averages that change from frame to frame. They manifest as fluctuating horizontal and vertical lines in the image.

The sigma_s noise component is the sensor array mean value. It is always zero by definition and is included in the table solely for mathematical completeness.

## Summary

IR camera systems tend to be operated in conditions where the dominant temporal noise source is scene noise, also known as shot noise or quantum noise. Scene noise is random fluctuations in the background of an infrared scene. Background is the "pedestal" or level of IR radiation emitted by the scene even in the absence of a target. There is also background generated by the camera optics and camera interior. The scene background can be quite low in situations where the camera is pointed into a clear sky in cold dry environments. In that case, the camera optics may be the biggest source of background and with it, shot noise. Other sources of temporal noise include dark current, readout noise and electronics noise. When integration times are long, as they can be in SWIR cameras, dark current noise can be significant. Read noise and electronics noise can become significant in situations where the background and dark current are both low. That can include looking at cold scenes or targets while using very short integration times to stop motion.

Spatial noise is always present but can be mitigated by high-quality non-uniformity corrections.

Radiometric accuracy is determined in part by the signal to noise ratio of the target of interest. When the target has a radiance that is close to the bottom of the radiometric range in use, measurements can get very noisy.

# Chapter 7 : MTF and Optics in Infrared Camera Systems

# Contents

Chapter 7 : MTF and Optics in Infrared Camera Systems ........ 7—302
Introduction ........ 7—304
What is the Size of the Blur Spot? ........ 7—305
RS8313 Range Camera ........ 7—307
Modulation Transfer Function ........ 7—308
MTF ........ 7—311
Comparison of Different Sensors ........ 7—312
Spot Size Effect ........ 7—315
Summary ........ 7—321
Bibliography ........ 7—321

## Introduction

One of the most noticeable aspects of infrared camera images is that they tend to look a bit "soft" to the eye when compared to visible-light photos of the same scene. If one digitally zooms in on small details in the image, the details can be quite blurry looking. This has implications for IR camera systems designed to detect small objects, like video trackers. An infrared image of a distant aircraft may look very fuzzy when the image is zoomed in. The image softness also has implications for temperature measurement. Small objects at a certain temperature above the ambient temperature of the scene can often look colder than they actually are. This so-called "spot-size effect" is discussed later in this chapter.

There are several reasons for this image softness, but the main reason is due to diffraction. Diffraction is an optical effect where light waves propagating through an aperture spread out in angle on the other side of the aperture. The longer the wavelength of the light, the more angular spreading one gets for a fixed aperture size. The smaller the aperture, the more the light waves spread out in angle. Typical visible and infrared camera systems have optics with similar f/numbers, around f/2.5. But midwave IR radiation has a wavelength that is ~8 times longer than visible light. For comparable f/numbers, the MWIR wave spreading is 8 times worse than it would be for a visible-light system with the same aperture size. Lenses produce diffraction effects on the images they form on focal plane arrays. But lenses also inevitably have aberrations which tend to misdirect light rays and soften the image further. We will consider both diffraction and aberration in this chapter.

Another image-softening effect is the relatively low sensor resolution of IR cameras compared to visible-light cameras. Until about 10 years ago, commercially available thermal IR camera sensors were substantially smaller than 1 megapixel. In 1999, a typical camera had a 320x256 pixel sensor which is only about 0.08 megapixels, and the pixel size was 30μm square. When the imagery from this type of camera is displayed on a 15-inch monitor, the image looks inherently soft. There just aren't that many pixels to go around. The analog video data was typically run through a digital image filter to soften the square edges of the pixels and make the images more appealing to the eye. Within a few years, image formats for standard IR cameras had grown to 640x512, a four-fold improvement that made a substantial improvement to the appearance of thermal IR video. The 25μm pixel size was still fairly large. An image of a point source forms a **blur spot** on the sensor due to optical effects in the lens. For these 640x512 cameras with 25μm square pixels, the blur spot size was comparable to the pixel size for the cameras with f/2.5 coldshields. Image sharpness is partially driven by the size of the pixels used in an imaging sensor. When the blur spot is of comparable size to the pixel, then the sharpness of a point-source target depends on the alignment of the image on the grid array of pixels. For this particular case of the 25μm pixel camera, if a point source image is centered on a pixel, then not much energy will fall onto adjacent pixels. But if the image of the source is centered on the intersection of four pixels, then the point source will appear bigger and less bright than in the first case. When the pixel size is comparable to the blur spot size, and the system is pointed at a moving point source like a distant jet exhaust plume, the apparent brightness of the plume will vary as the image of it moves across the focal plane array, producing an oscillation in the apparent radiant intensity.  This sampling effect may affect the performance of a tracker algorithm and it further complicates any efforts to make radiometric measurements of the target. For long-range tracking cameras like the FLIR RS8300 series, the blur spot is many times bigger than a sensor pixel, so this effect is much less pronounced, as there are always a number of pixels irradiated by the spot produced by a point source.

The analysis of an imaging system's angular resolution starts with the principal of superposition. Every point in the object space is mapped to a point in image space.  Each point in object space that is imaged onto a focal plane array becomes spread out in a manner described by the **point-spread function** of the optics.  The PSF describes the shape of the blur spot. The final result in image space is a linear superposition or summation of all the PSFs from all the points being imaged in the scene.  The focal plane array then samples the image which is focused upon it. The sampling is done at the spatial resolution of the detectors on the focal plane array.  Smaller pixels can sample the image at higher spatial resolution, which is desirable.  More recent science camera sensors have pixel pitches that are 14μm, 12μm or smaller.  The advantage of using these small pixels is that more of them can be fit into a fixed piece of sensor real estate, which lowers sensor manufacturing cost.  As well, the effect of the detector dimensions on image sharpness becomes less pronounced.  This leads camera and optics designers to improve the optics to take advantage of the smaller pixels.  The downside of small pixels in an IR camera focal plane array is the reduced charge handling capacity of the unit cell amplifier circuit, more issues with crosstalk of signal between pixels, and the greater degree of difficulty in designing the readout IC with such tight space limitations.

## What is the Size of the Blur Spot?

The blur spot size is determined by various optical effects.  The primary effect for thermal IR cameras is diffraction, followed by optical aberrations in the lens.  For long-range target imaging, atmospheric turbulence will also cause blurring.  This is the effect that makes stars appear to twinkle and makes them look blurred out, since we observe them through many kilometers of air path that is full of constantly moving convection cells that distort the phase of incoming spherical light waves due to variations in refractive index on the boundaries of the cells.

If we ignore aberrations in the optics and turbulence effects, then the blur spot size can be calculated by analysis of the diffraction pattern produced by a circular aperture.  A monochromatic point source of IR radiation (like a narrowband laser) propagated through a perfect lens (no aberrations) will make a diffraction pattern of concentric rings with a central bright spot known as an Airy pattern, as shown in Figure 1. The central bright spot is called the Airy disk, and this central disc contains 84% of the IR radiation in the Airy pattern.

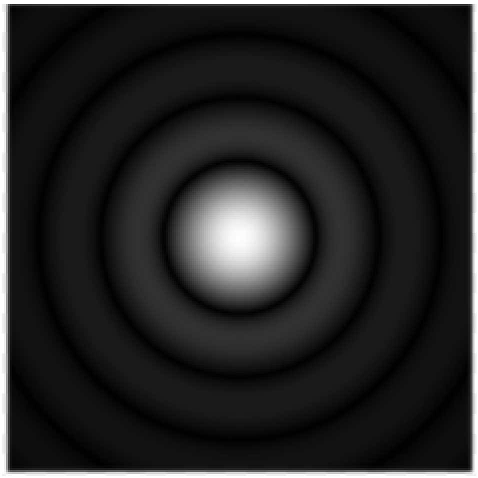

**Figure 1. Diffraction pattern or Airy pattern from a monochromatic plane wave coming through a circular aperture**

If we assume monochromatic light, infinitely small pixels and lenses with no aberrations, two different point sources will form Airy patterns and the two sources become indistinguishable from each other when they get close enough in the image, as shown below in Figure 2.

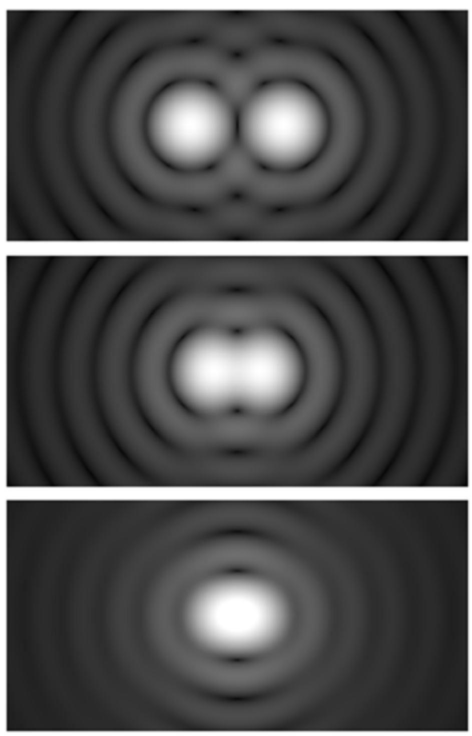

**Figure 2. Two monochromatic Airy patterns at various spacings. The top image has them spaced apart by twice the distance to the first minimum. The middle image shows one Airy disk that is exactly the distance to the first minimum of the second pattern (the so-called Rayleigh criterion), and the bottom image has half the distance.**

The first minimum in the Airy pattern (roughly the radius of the Airy disk) is given by the following expression:

**Airy Disk radius = 1.22 * Wavelength * f/number**

For a typical 3-5um InSb camera, an intermediate wavelength value of 4µm can be used and the f/number for a typical InSb camera coldshield is f/2.5.

Inserting these values into the expression above gives a radius of 12.2µm and a diameter of 24.4µm. This blur spot size is well matched to a 25µm pixel. It means that for a so-called diffraction-limited lens, the majority of the energy from a point source can be contained within a single pixel. A diffraction-limited lens is one where the effects on the blur spot size caused by aberrations in the lens are negligible compared to the effect of diffraction. In a real camera system, the source radiation is not monochromatic. For a 3-5µm camera imaging a point source like a distant jet exhaust, the actual Airy pattern will be more complex than the monochromatic patterns described above. The light waves with 3um wavelength will diffract less than the 5µm light waves. The patterns from these different wavelengths will overlap and spread out more than the 4µm wavelength example above. In the end, the blur spot size for a non-monochromatic target will depend on the relative spectral distribution of the target within the 3-5µm camera bandpass.

## RS8313 Range Camera

For a FLIR RS8313 range camera, the situation is quite different from the 25µm pixel, f/2.5 science camera example mentioned above. In the RS8313, the f/number of the optics is f/5 and the sensor pixels are only 14µm square. The diffraction blur spot diameter is 48.8µm for 4µm wavelength radiation. This blur is 3.5 times the pixel size, so the optics drives the resolution of the system more than the pixel size. Another way to state it is that each pixel's detector on the focal plane array is oversampling the optics-induced blur, which means that the image sharpness of the system is much more dependent on the optics than on the detector size. In the case of f/5 optics in the midwave IR band, the diffraction effect is substantial.

The correct way to model the diffraction blur is more complex than this simple model result shown in Figure 3. The diffraction patterns for all the wavelengths between 3 and 5µm are added together with the appropriate spectral weighting for the target spectrum. There is additional blur from aberrations in the lens, so the actual blur spot is even bigger than the diffraction blur alone.

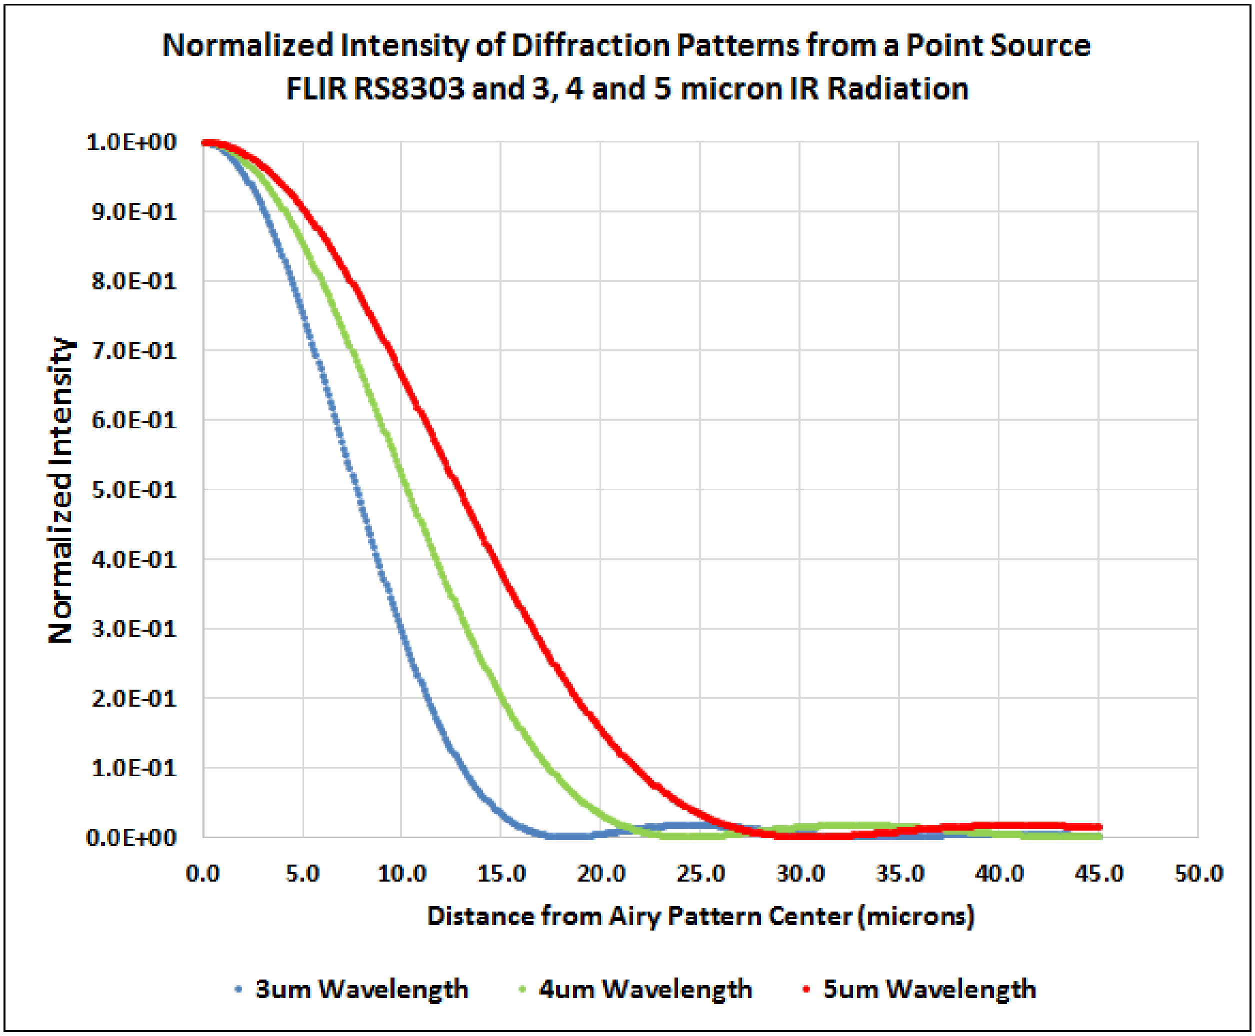


**Figure 3. Three Airy patterns for three different wavelengths, all normalized to 1 at zero distance from the center of the pattern. At a mid-band wavelength of 4µm, the first minimum of the Airy pattern occurs at a radius of 24.4µm. This radius is linearly proportional to f/number which in this case is f/5. Because of the longer wavelength, thermal IR camera systems have more much diffraction compared to visible-light systems at the same f/number.**

## Modulation Transfer Function

While the point-spread function is useful and interesting, there is a more powerful way to describe an optical system's resolving power or ability to discern small details in a scene. The modulation transfer function (MTF for short) is a graph of the normalized contrast in an imaging system as a function of spatial frequency. Spatial frequency is a characteristic of an image that is periodic in position across space, and it can be thought of as the inverse of position space. The mathematical description is that the MTF is related to the Fourier transform of the point-spread function from position space to spatial-frequency space. The point-spread function by itself is a difficult parameter to use in a trade study of camera system design. It is much easier and useful to consider optical systems imaging resolution performance in the context of spatial frequencies, because MTF curves for each component of an optical system can be multiplied together to determine the total system MTF. It is not possible to do a simple multiplication of point spread functions in position space and end up with the point spread function for the system. Instead, one needs to do convolution integrals. In spatial frequency space, one can simply multiply the modulation transfer functions together to get the total system MTF.

A recent measurement of an RS8300-series camera was done using a Santa Barbara Infrared scene projector (AKA collimator) and a half-moon target that is back illuminated by a 100 °C area blackbody, as shown in Figure 4. The half-moon target is rotated so that its edge is slanted by an angle of around 5 angular degrees off vertical. A series of images is taken of this so-called slant-edge target and a MATLAB script called sfrmat.m (written by Peter Burns) is used to process the image data into an edge-spread function with sub-pixel sampling. I mention this MATLAB script because it is easily found on the web with a search on the name of the script and Peter Burns – he has generously made it available to all. The first derivative of the edge-spread function with respect to position gives the line spread function. The absolute value of the Fourier transform of the LSF is the spatial frequency response or MTF.

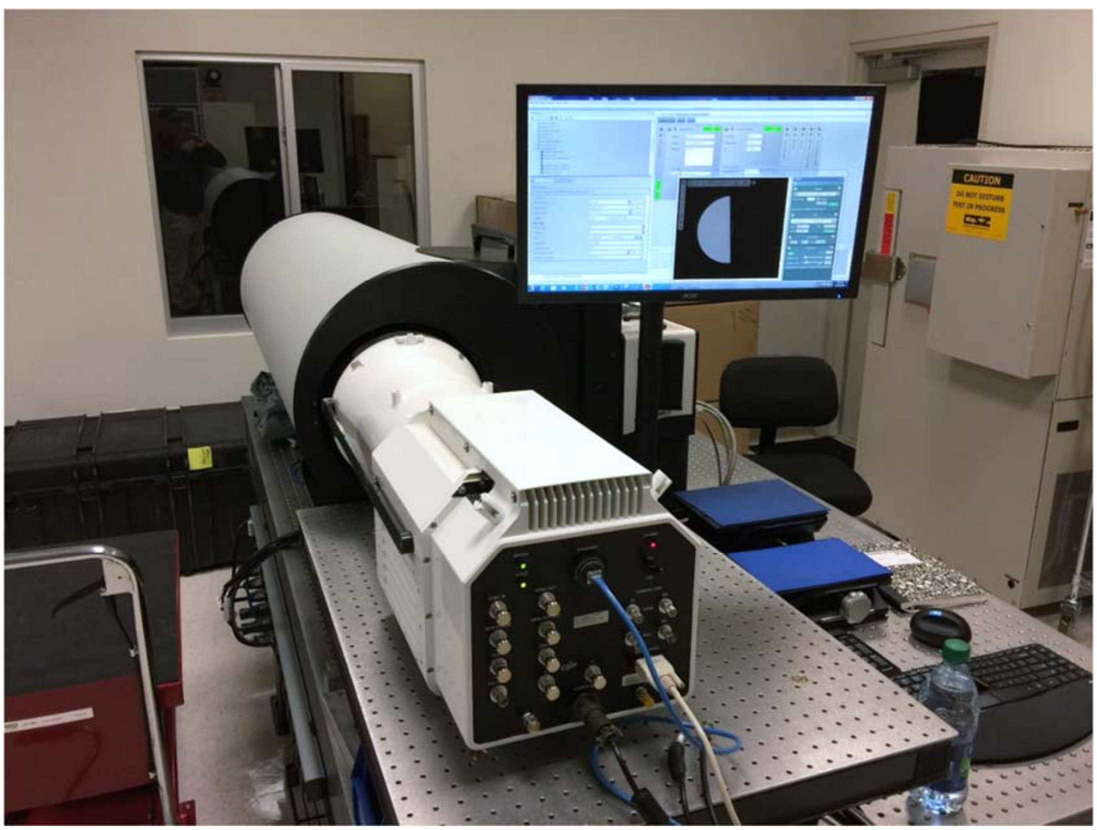


**Figure 4. RS8313 looking into collimator at slant edge target**

The camera was zoomed in to the maximum focal length of 1200mm. At that high magnification, only a small section of the target close to the edge transition between hot and cold is visible in the IR image.

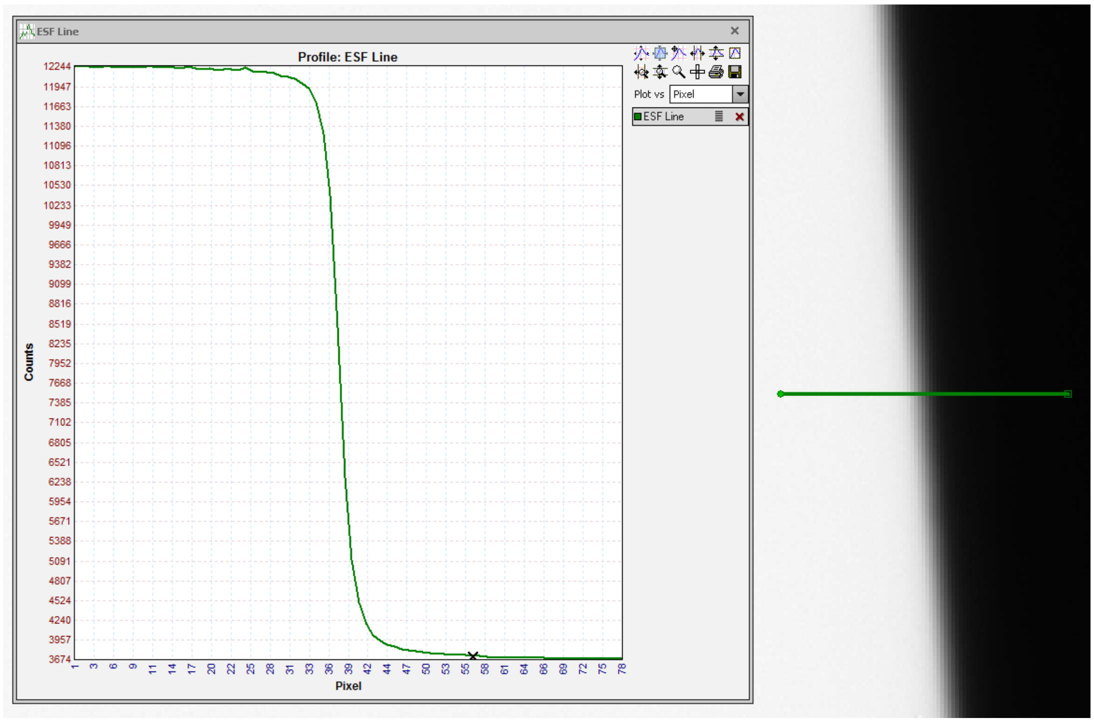


**Figure 5. Edge spread function along a single line profile on the slant-edge target image**

The ***edge spread function*** looks like the line profile in Figure 5, except that it is super-sampled by combining the edge-spread functions for many rows of the image with the correct phase shift added. The resulting function is even smoother than the above line profile.

Taking the first derivative of the edge spread function with respect to the horizontal pixel coordinates yields the ***line spread function*** or LSF, which looks like the plot in Figure 6:

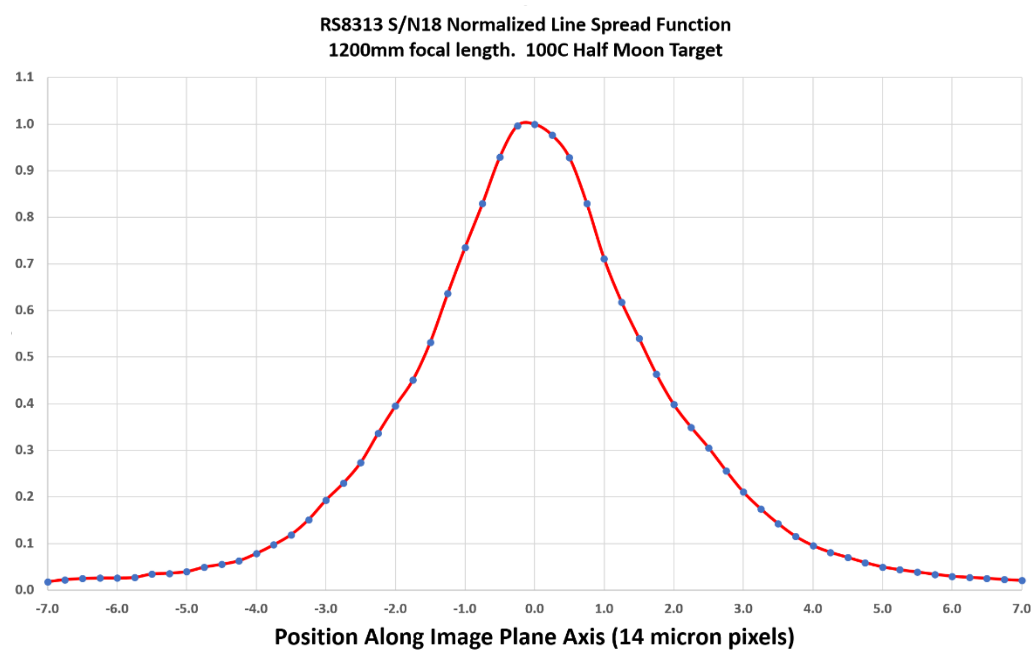


**Figure 6. Line spread function on RS8313 image**

Figure 6 is equivalent to a horizontal line profile sliced perpendicularly to the image of a very thin, bright line target, like a thin slit that is back illuminated. The X axis is in units of pixel pitch, which is 14 microns. The Y axis is in normalized intensity units. An interpretation of the plot is that a very thin line target blurs out to be about 10 pixels wide.

Imaging a knife-edge target does not result in a perfectly sharp-looking image of an edge. Instead there is intrinsic blur in the transition zone from dark to light, even at the optimal focus. Focus error makes the blurring worse. This blurring has a particularly noticeable effect on images of small targets. This can be summarized as follows:

**When even a sharply defined, high contrast target subtends a small enough angle as seen by a camera (target less than about 20 pixels across), then the blurring of the edges of the target in the image reduces the peak brightness, in turn reducing contrast and making the target radiance closer to the background.**

For a camera like the RS8313 at full zoom, the angular resolution of the system is excellent from an object-space point of view. Each pixel in the sensor subtends an 11.7 microradian angular field of view at full zoom, which is a focal length of 1200mm. A 1-meter target at 1000-meter range will produce an image with 86 pixels across it, which is excellent spatial sampling.

## MTF

The MTF of the camera system is computed by taking the Fourier transform of the LSF. The absolute value of this is the MTF. Figure 7 shows the MTF curve derived from the LSF shown in the previous section, where the X axis is expressed in units of cycles per milliradian. This is a very useful unit of measure for spatial frequencies, since it is referenced to angles in object space.

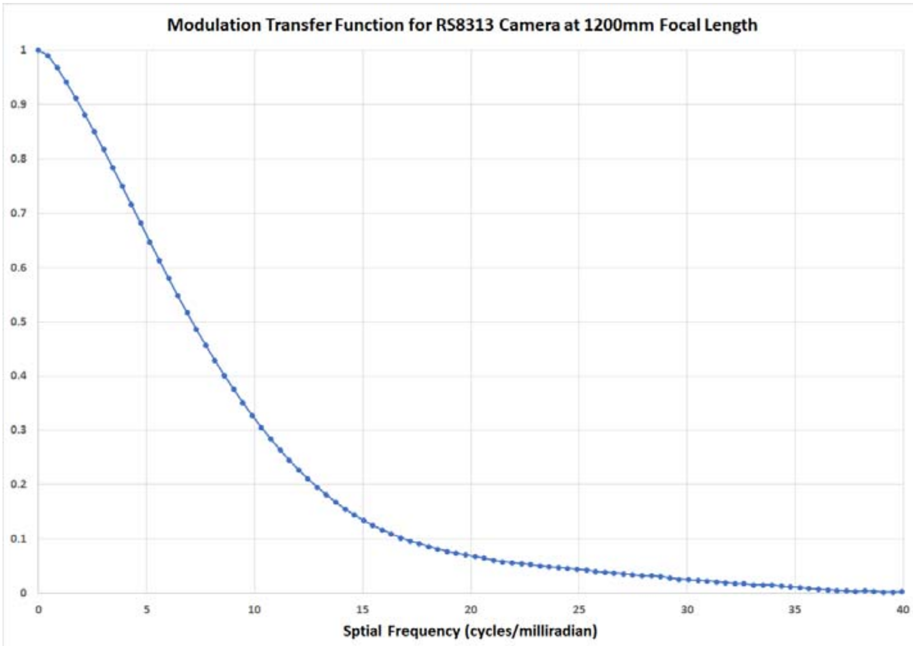


**Figure 7. Modulation transfer function for RS8313 camera at full zoom**

## Comparison of Different Sensors

I now consider the FLIR RS8313 optics imaging the scene onto focal plane arrays with different pixel sizes. The standard RS8313 camera system is built on the SC8313 camera with 14µm pixels. The FLIR X8500 camera has 12µm pixels and the X6900 has 25µm pixels. The X8500 is 1280x1024 and is desirable for high resolution imaging with decent frame rates (181Hz). The X6900 is 640x512 but has a very high frame rate of 1000Hz. As a point of comparison, one can graph relevant MTF curves for both the existing 14µm pixel case and then for a RS6900 camera which would contain the X6900 camera with 25µm pixels. The plot in Figure 8 show the MTF curves in object space which means that the spatial frequencies are defined in units of cycles/milliradian, projected out into the scene. The blue and green curves show the sampling MTF, also known as the detector footprint MTF. For the 25µm pixel pitch sampling MTF curve (green), the MTF drops to zero at a spatial frequency of 48 cycles/mrad. This is the object-space spatial frequency where a cycle's wavelength is the same as the detector size, and contrast drops to zero because every pixel gets the same energy from the scene. The 14µm pixel pitch sampling MTF curve (blue) is much higher because the pixels are smaller. Smaller pixels sample the image at higher spatial frequencies.

The red and purple curves are the system MTFs for the 14µm and 25µm pixels, respectively. At lower spatial frequencies, like 10 cycles/mrad, the percentage difference in the two MTF value is small. At a high spatial frequency like 30 cy/mrad, the percentage difference becomes profound.

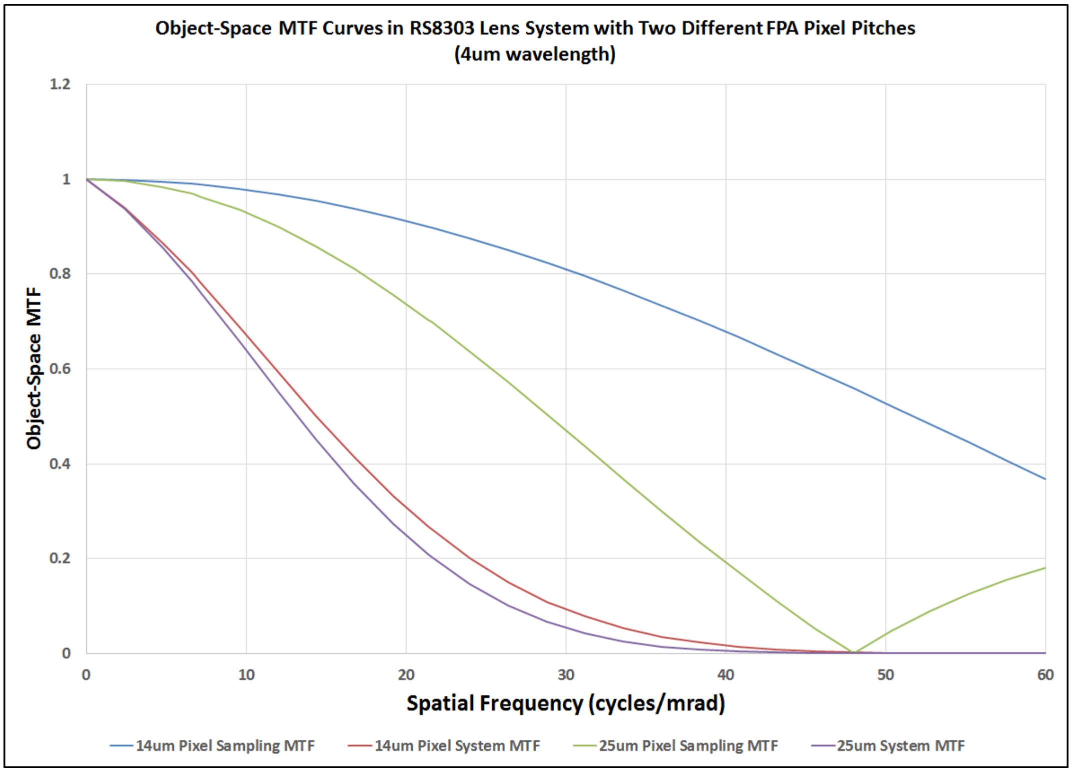


**Figure 8. Object-space MTF curves for two different cameras used with the RS8313 lens at full zoom (1200mm focal length)**

Figure 8 shows a plot of the ratio of the MTFs for the two different pixel pitches. The ratio is pretty close to 1 out to about 15 cycles/mrad, then it starts to climb. By 35 cycles/mrad, the MTF for the 14µm pixel system is 2.2 times higher than for the 25µm case. This 14um system will transfer contrast at 25 cy/mrad 2.2 times better than the 25µm pixel system would. BUT it should be noted that when the spatial frequency gets out to 35 cycles/mrad, the MTF for the 14µm pixel system is only 2%. The high-frequency features in the scene will have to have very high contrast to make a decent-looking image with this factor of 50 penalty. For the 25µm case, the modulation transfer is 2.2 times worse, or less than 1% transfer.

The upshot of this analysis is that from the perspective of imaging resolution performance, it is always desirable to have small pixels that oversample optics blur. Making pixels smaller should always improve system MTF, although there are other sources of blurring and crosstalk in the focal plane array that become more pronounced for very small pixels, notably ROIC crosstalk and diffusion of charge carriers in the InSb from one pixel to another. The disadvantage of smaller pixels is less sensitivity for a given photon flux density and a tendency for the well capacity to scale with pixel area, which means that small pixels tend to have less dynamic range.

An excellent software package for the modeling of system MTF is NV-IPM. IPM stands for integrated performance model. This software was written by engineers at the Night Vision and Electronic Sensors Directorate at Fort Belvoir in Virginia. IPM allows the user to model a camera system from the target all the way through to the perception of a human observer. It is a vastly more comprehensive model compared to NVThermIP, the earlier modeling software. It is also possible to model the images of a scene by giving the model an image as an input. The model "plays" the image through the system MTF and shows the user what the target would look like when imaged by the camera system at a specified range.

Figures 9a-d show examples of simulations of a very simple scene imaged by two variants of Niceville cameras: the RS6800 and the RS8313 at full zoom, which corresponds to a focal length of 1200mm. The input image is a three-bar target with 100C of thermal contrast created in Adobe Photoshop and intensity scaled using ResearchIR. The bars are 7 meters high and the range to the bars is 5 km through a typical rural atmosphere with medium levels of turbulence. The atmospheric modeling is done by MODTRAN code, a radiative transport model developed by the Air Force Research Laboratory. Figure 9a is the RS6803 simulated image:

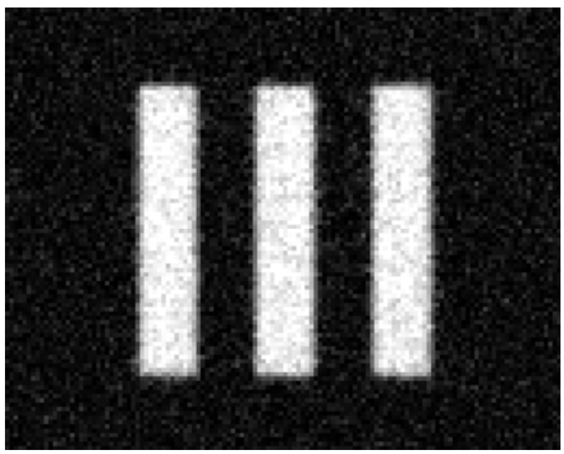

**Figure 9a. RS6803 simulated image of bars through 5km of air**

Figure 9b is the RS8313 simulated image:

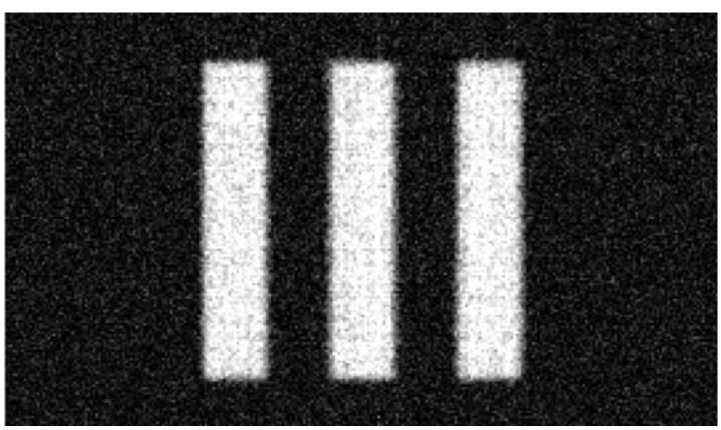

**Figure 9b. RS8313 simulated image of bars through 5km of air**

The target is clearly reproduced in both images. Now let's look at the same target imaged through 15km of atmosphere by the same cameras. Figure 9c is the RS6803 simulated image, which is very pixelated but has slightly higher contrast than the RS8313 simulated image in Figure 9d. One sees the effects of larger pixels: lower spatial resolution at the same focal length, but higher image contrast because more IR radiation is collected per pixel. At 15km range, the spectrally weighed atmospheric transmission in the 3-5-micron band is reduced to about 6%. The 100 °C of contrast at the target is reduced to 0.06 or 6% of its initial value due to absorption through 15km of air.

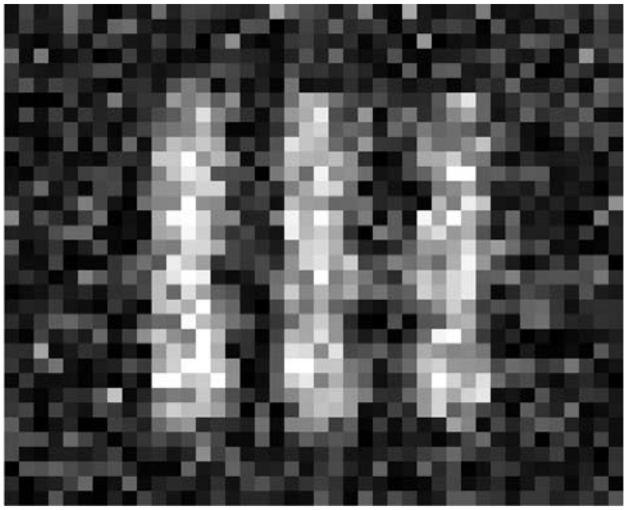

**Figure 9c. Simulated image of 100 °C contrast bar target viewed by an RS6803 through 15km of air**

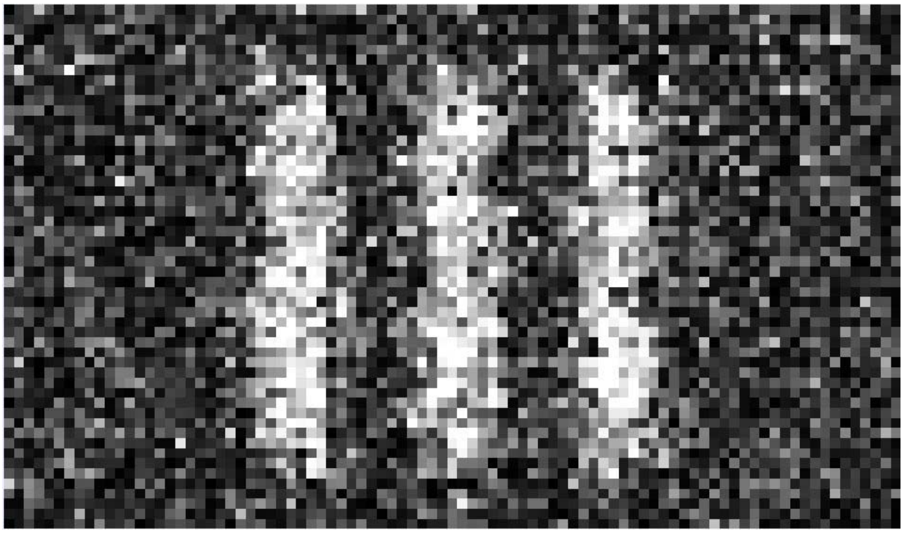

**Figure 9d. Simulated image of 100 °C contrast bar target viewed by an RS8313 through 15km of air**

## Spot Size Effect

The spot size effect is caused by diffraction and aberrations in the lens blurring the image of a circular target. The blur is the two-dimensional convolution of the point spread function of the optics with the geometrical image, as shown below in Figure 10. Note the sloped sides of the pattern in the simulated image. When those slopes intersect, there will be a loss of contrast in the center of the spot.

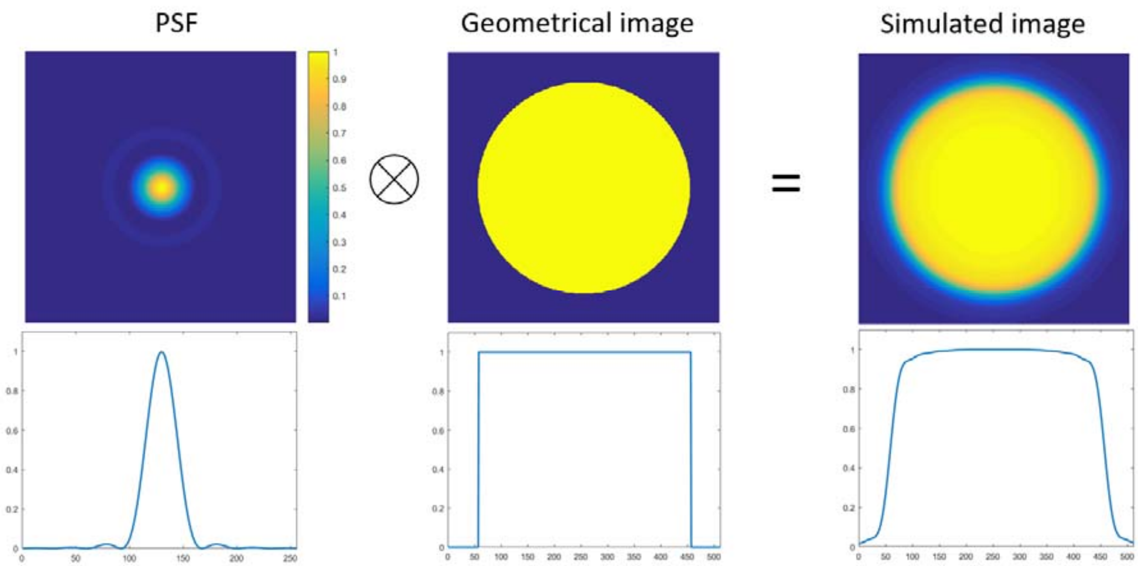


**Figure 10. Convolution of point spread function and geometrical image to produce simulated blurred image**

A good way to understand the spot size effect is to look at a target of known size and temperature and compare what the camera measured to "ground truth". Then vary the size of the target. We have found through experimentation that the effect occurs for targets that are smaller than about 20 pixels. A spotmeter placed in the center of a hot target will start to drop in value as the target gets to be less than about 20 pixels across. This is a rule of thumb only, and it can vary for different lens types.

Figure 11a-d shows images of a cavity blackbody set to 65 °C with an aperture wheel in front of it. The circular aperture diameters vary from 16mm to 1.6mm.

The first image shows a case where the spot size effect is negligible at the center of a target which is about 38 pixels across. Note the sloping sides of the line profile. There is not a sharp transition at the aperture edges – instead there is a softening of the image. The aperture is 16mm in diameter which is about 38 pixels at the 30.5cm range. The spotmeter in the center reads 65.7 °C, very close to the indicated temperature of 65 °C.

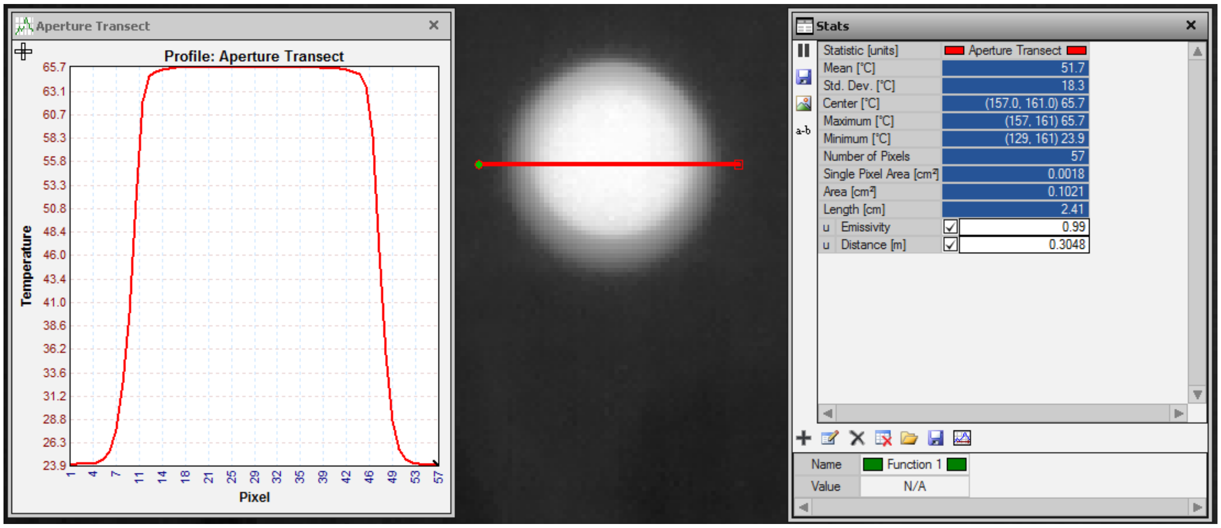


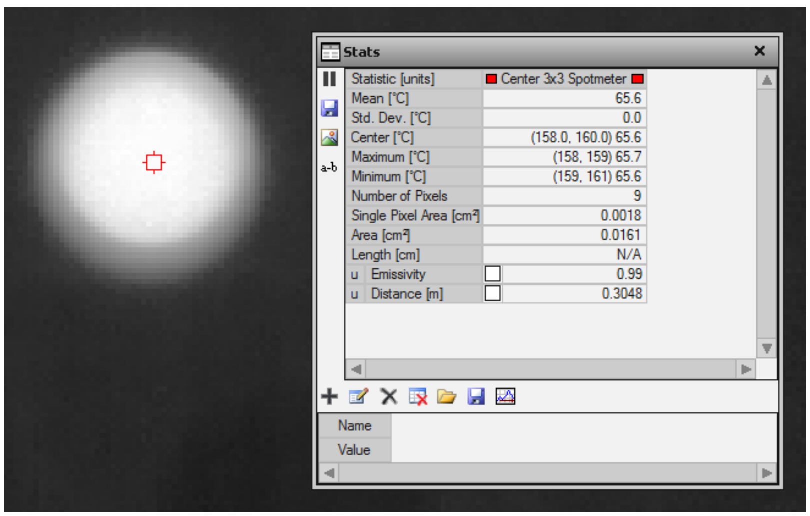


**Figure 11a. FLIR A325 uncooled camera image of a 16mm aperture in front of a 65 °C cavity blackbody imaged at 30.5cm working distance. Note the blurriness of the aperture edges – this is not due to bad focus, it is the intrinsic softness in the image caused by diffraction, lens aberrations, and pixel size. The maximum measured temperature is 65.6 °C**

In the next set of images in Figure 11b, the aperture is 6.4mm in diameter. This diameter is about 15 pixels in the image. The spotmeter temperature has dropped to 65.1 °C, down from 65.7 °C. The spot-size effect is starting to affect the measurements.

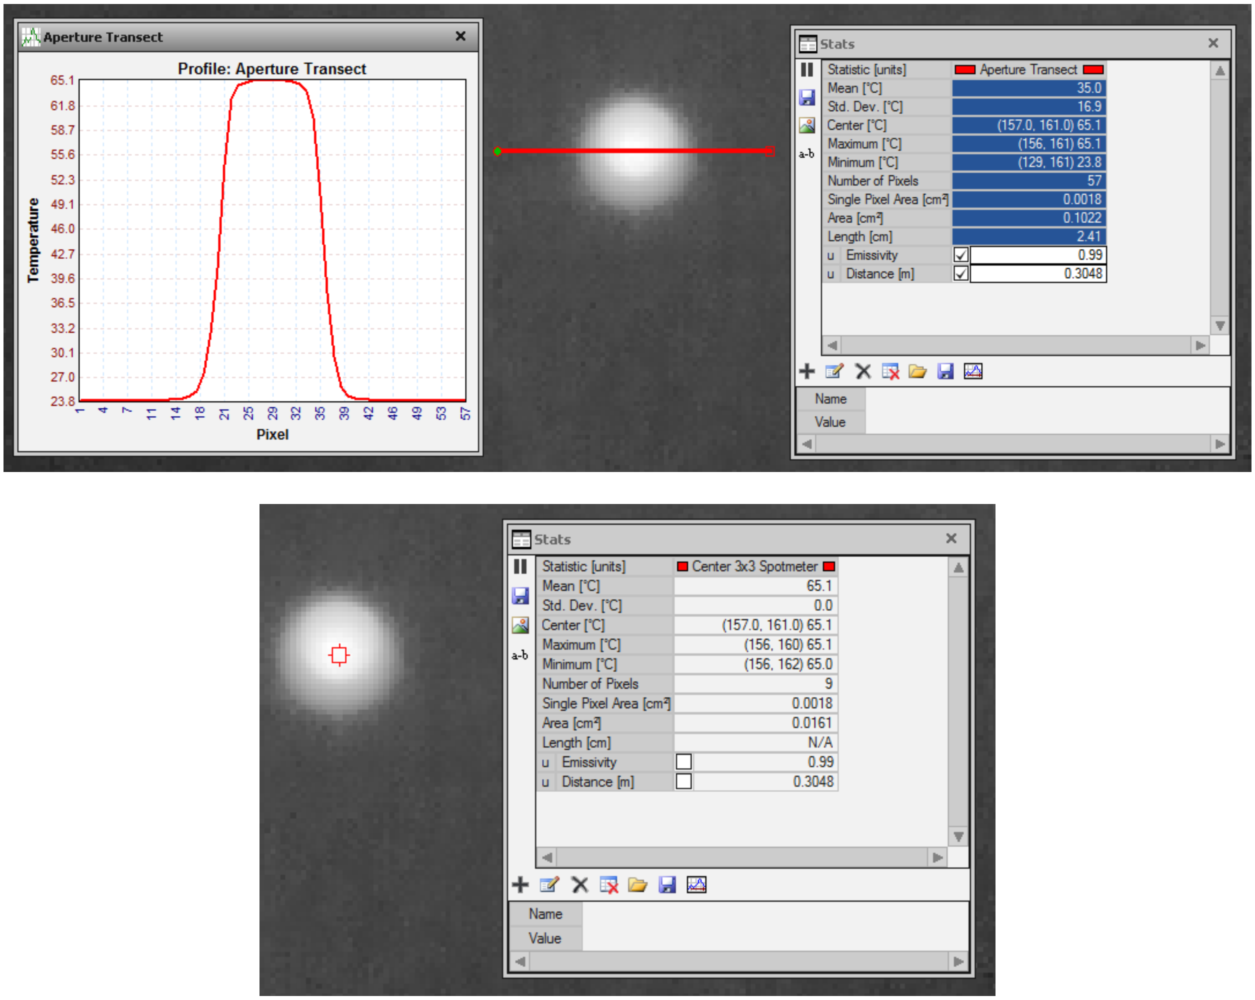


**Figure 11b. The smaller 6.4mm diameter aperture lowers the maximum measured temperature slightly**

Figure 11c shows results for a 3.2mm aperture, and the spot size effect is now noticeable. There is no longer a flat “mesa top” on the line profile, and the aperture is only 7.5 pixels across. The peak temperature measurement along the transecting line is 63.6 °C, which is 2 °C colder than measurements with the 16mm aperture. The 3x3 pixel spotmeter reads 62.7 °C which is a significant error.

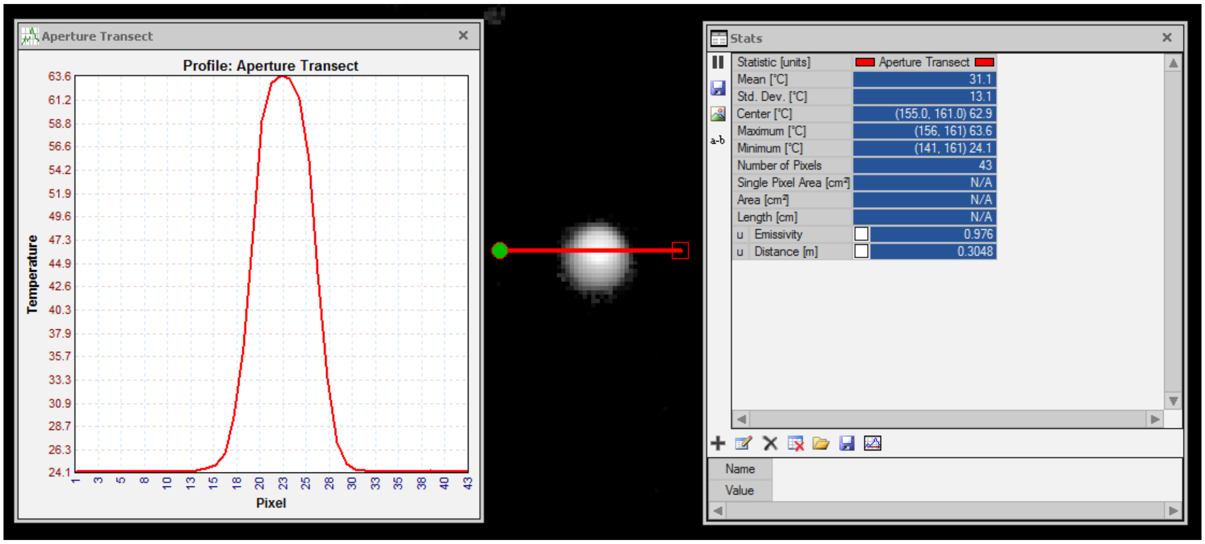


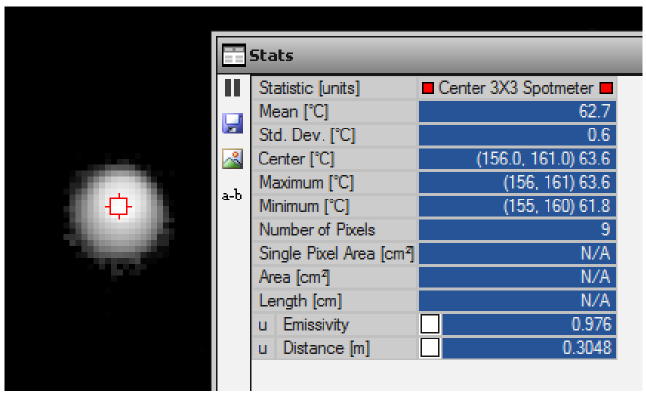


**Figure 11c. Using a 3.2mm diameter aperture lowers the maximum measured temperature to be out of the 2 °C accuracy specification**

Figure 11d shows the results from a 1.6mm aperture image. There is a very substantial loss of signal, which translates into a large temperature error. The spotmeter reads 49.1 °C at the peak.

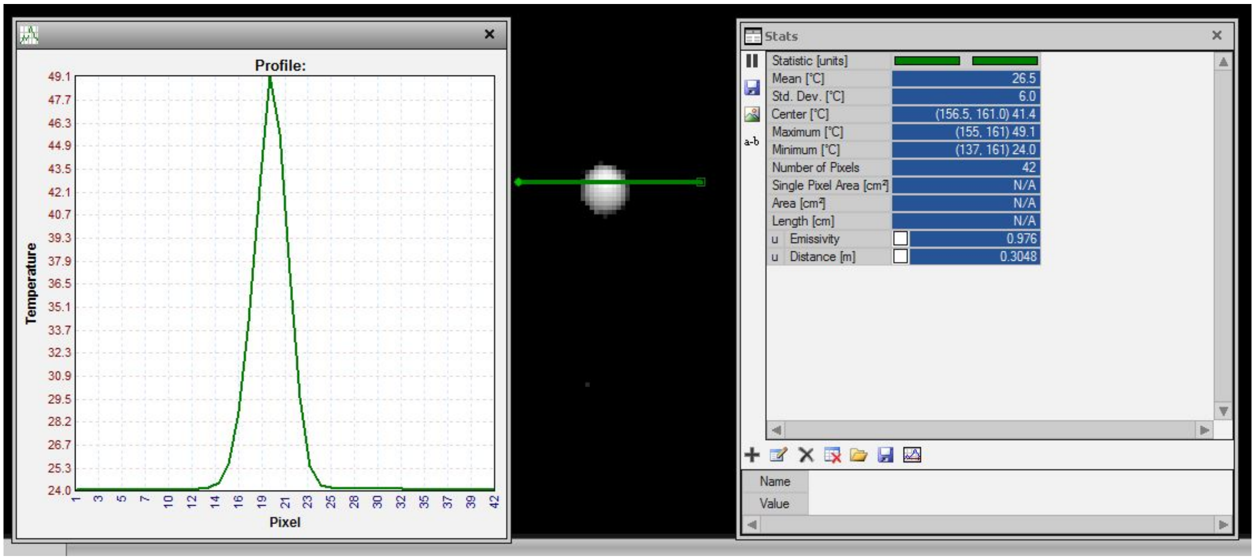


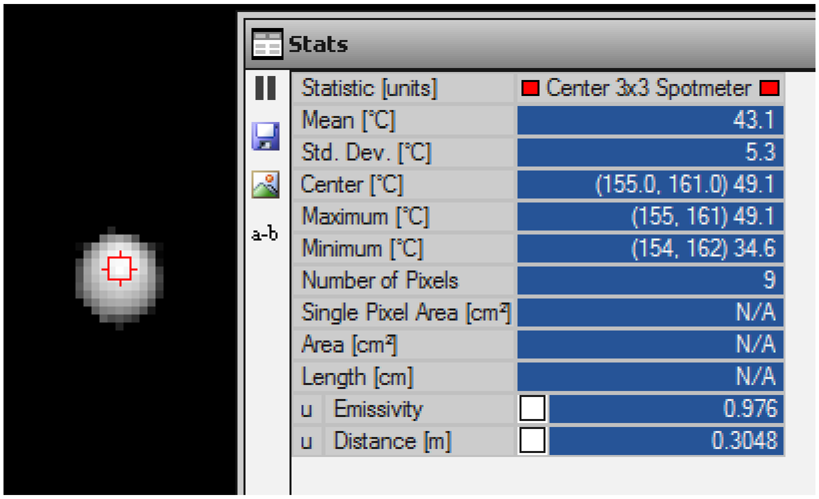


**Figure 11d. The 1.6mm diameter aperture lowers the maximum measured temperature to be way out of the 2 °C accuracy specification**

These examples above highlight the difficulty of precise radiometric measurements on targets that appear small in the image. The answer is usually to get closer to the target or use a longer focal length lens on the camera system. The general rule of thumb that seems to be true for just about any IR camera system is that a target against a high-contrast background that is bigger than about 20x20 pixels will measure out as well as if it appeared larger in the image. When the target is smaller than about 20x20 pixels, the system MTF will start to degrade the radiometric accuracy by an effect where the background gets "mixed" into the image of the target. The most common occurrence of this is a hot, small target looking colder than it really is against a cold background.

## Summary

All IR cameras images have inherent blur to them. Most of the blurring is caused by diffraction and aberrations in the optics, and some of it is due to the size of the pixels in a 2D detector array, especially for larger pixels. At longer target ranges, there is blur caused by turbulence, which is random motion of convection cells in the atmosphere, and the effects of turbulence are especially acute over multiple kilometers of range. Diffraction is typically the major cause of image blur because of the long wavelengths of IR radiation relative to visible imaging systems. Diffraction drives camera designers to use faster optics and faster coldshields; however, there are limits to how fast a long focal length lens can be because of the cost of large IR optics. Very large apertures on IR cameras also allow more phase distortion from atmospheric turbulence to affect image quality.

The MTF or modulation transfer function is a powerful and convenient metric to characterize IR camera systems. MTF is used in the design of camera sensors, coldshields and optics. Measurements of MTF in the lab are used to evaluate as-built camera systems to meet customer requirements for image resolution, and by customers to compare camera systems from different vendors.

## Bibliography

For those interested in understanding MTF and sampled imaging systems, there are two references that really stand out:

Modulation Transfer Function in Optical and Electro-Optical Systems, Glenn D. Boreman, SPIE Press (2001)

Analysis and Evaluation of Sampled Imaging Systems, Richard H. Vollmerhausen; Donald A. Reago Jr.; Ronald G. Driggers, SPIE Press 2010

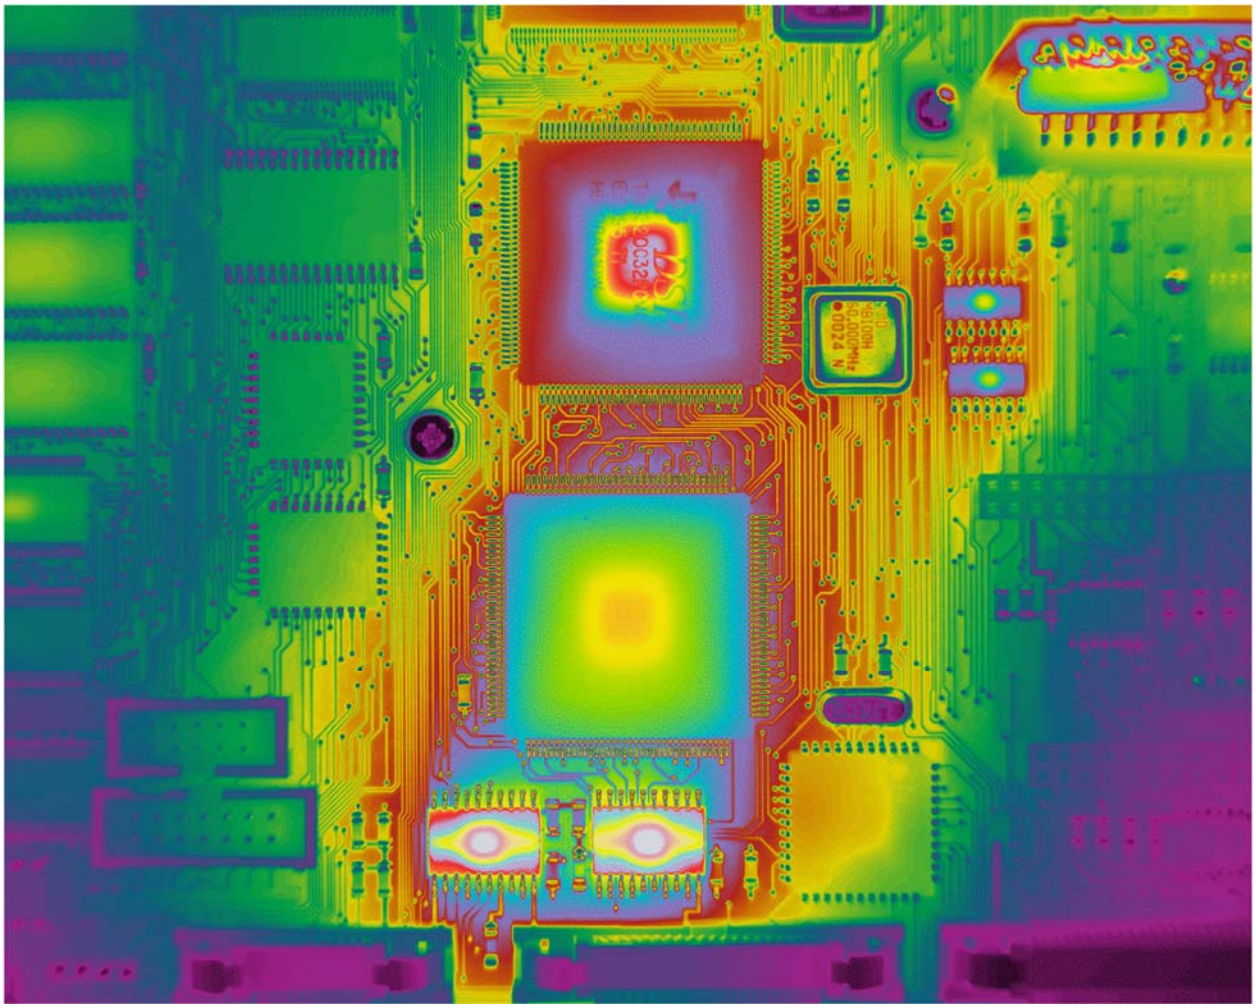

**A high-definition midwave infrared image of a powered-up circuit board with a triple rainbow palette. The palette has three rainbow stacked on top of each other with decreasing color saturation as the temperatures get hotter.**

# Contents

Chapter 8 : Displaying Infrared Image Data – AGC and Color Palettes ........ 8—322
Introduction: ........ 8—324
AGC - Displaying digital image data: ........ 8—324
Non-linear AGCs ........ 8—329
Image Palettes ........ 8—331
Monochrome Palettes ........ 8—332
Color palettes ........ 8—335
Images with high spatial frequency content ........ 8—339
Popular Color Palettes ........ 8—341
Hybrid Palettes with Both Color and Grayscale ........ 8—345
High Contrast Color Palettes: ........ 8—350
Summary ........ 8—356

## Introduction:

FLIR Science cameras output raw digital image data. For machine vision applications, the raw data is generally enough for the algorithms, but image data is most often interpreted by a human observer. These data need to be processed in order to be properly displayed to the user. Two main processing steps are involved: automatic gain control and the application of a color palette. Other image display methods might include the application of image processing filters, but these are not covered in this chapter.

## AGC - Displaying digital image data:

The typical FLIR scientific infrared camera puts out 14-bit image data in a video stream with a variable frame rate. Each pixel in the image is represented by a 14-bit digital number, which is really just a 16-bit number (two bytes) with two so-called packing bits that fill out the 14 bits to 16 bits. The brighter the scene that the pixel views, the higher the digital number. 14 bits is equal to 2^14-1 or 16,383 and the camera cannot output image pixel values higher than this. On the other end of the brightness scale, the digital numbers do not go much below about 500 counts. There is a voltage offset that is deliberately added to prevent the digitizer from receiving pixel voltage values that are below the input range.

Here is a typical midwave IR image taken with a camera that is set to a scene temperature range that extends up to around 55C. The entire 14-bit range of the camera is on display here, with just a simple linear scaling to map the 14-bit output to an 8-bit display. The typical computer screen has an 8-bit display, where each of the three principal display colors (red, green and blue) are represented by a number between 0 and 255, or 2^8-1. With a simple mapping like this, the image has low contrast:

**Fig 1a. This image has the following range of digital levels mapped to 256 display levels. We denote this as 0-16784: 256.**The image histogram shows that most of the pixels are between about 4400 counts and

12000 counts, with a spike at 16383 where the coffee mug full of hot water is saturated in the image. Figure 1b below shows the different parts of the histogram. The green lines show the display limits of the image.

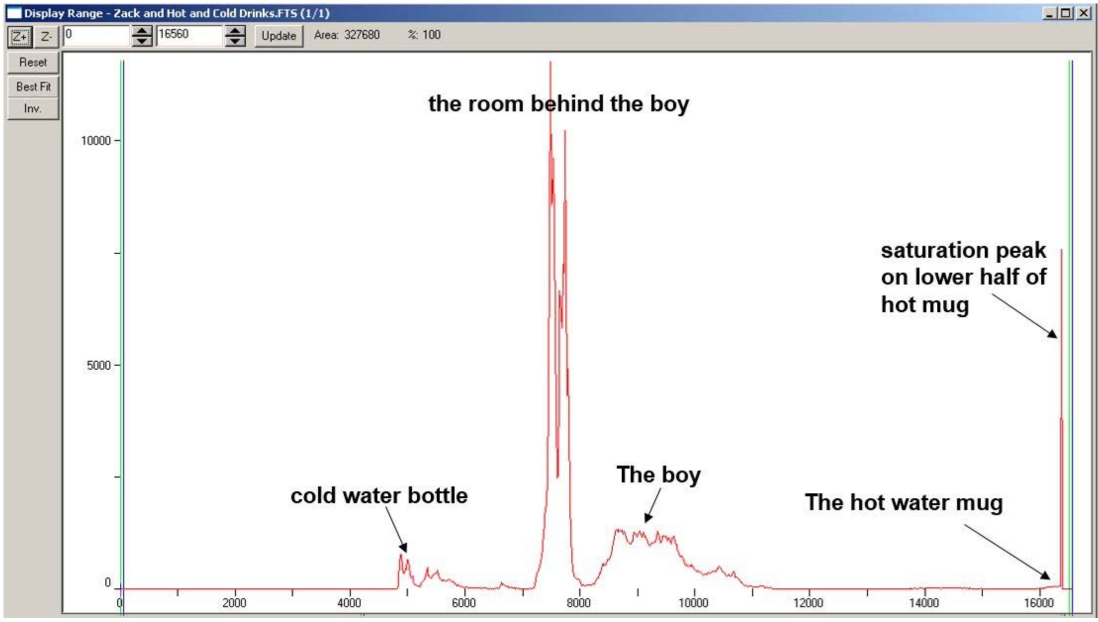


**Fig. 1b. Histogram of image in Figure 1a.**

We can improve the displayed image a little by not using up any of our 256 display levels on the 14-bit data between 0 and 4400. The mapping is now from 4400 to 16784. Note: I set the limits beyond 16383 so that one can see the saturation peak. There aren't any pixels with raw counts in the 0-4400 range so why waste levels there?

**Fig. 2a-b. The left image shows the results of limiting the display to digital levels between 4400 and 16784. The right image is the original image from 1a where all digital count levels between 0 and 16784 are mapped to 8 bits. The contrast has improved slightly, but there is still very low contrast in the scene behind the boy.**

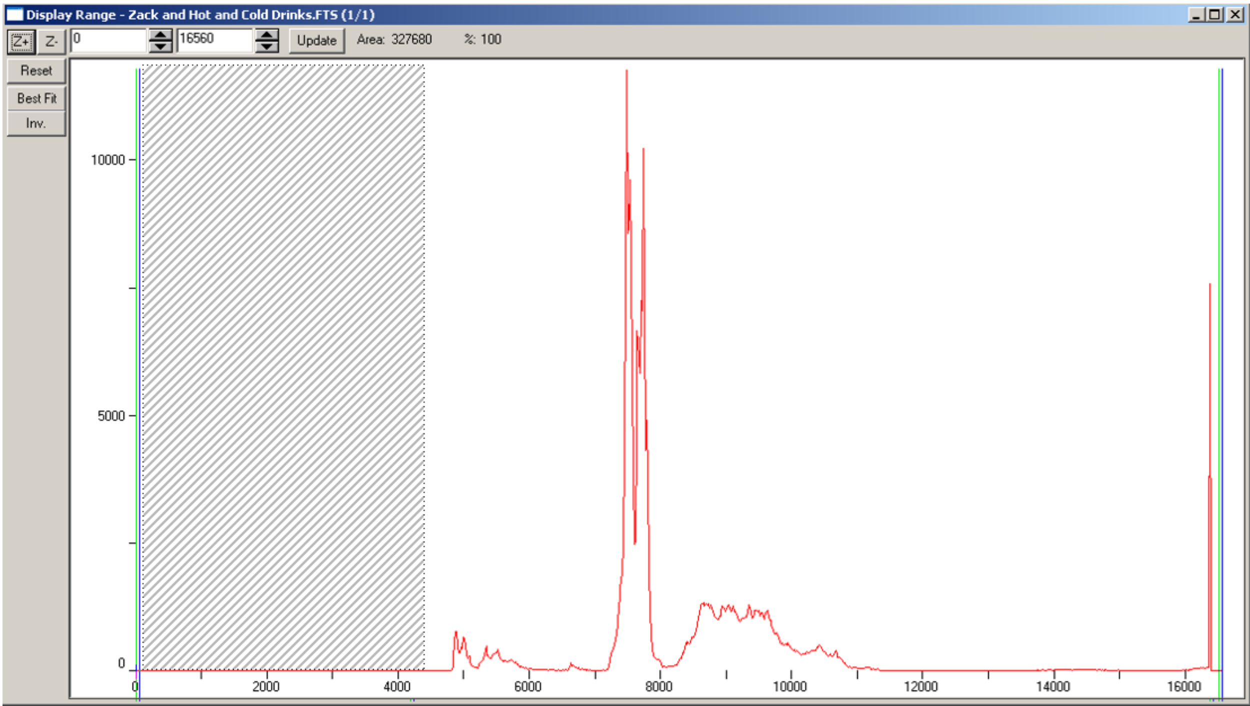


**Fig 2c. Histogram showing clipped region below 4400 counts and "desert" between 11,000 and 16,000 counts.**

We threw away the shaded area (0-4400) on the left. But we still have a gap in my histogram data between 11500 and ~16000. There are virtually no pixels with raw counts in that range and we are wasting display levels on that empty range of the histogram. We could pull in the high limit of the mapping to ~11500. But then that will blow out all pixels with counts above ~11500 – they will look dead white in the displayed image. Here is what that looks like:

**Fig. 3a. Image of boy with histogram limits of 4400 to 11,500 counts.**

Some details in the kitchen behind the boy are starting to emerge. But this image still does not look very good.

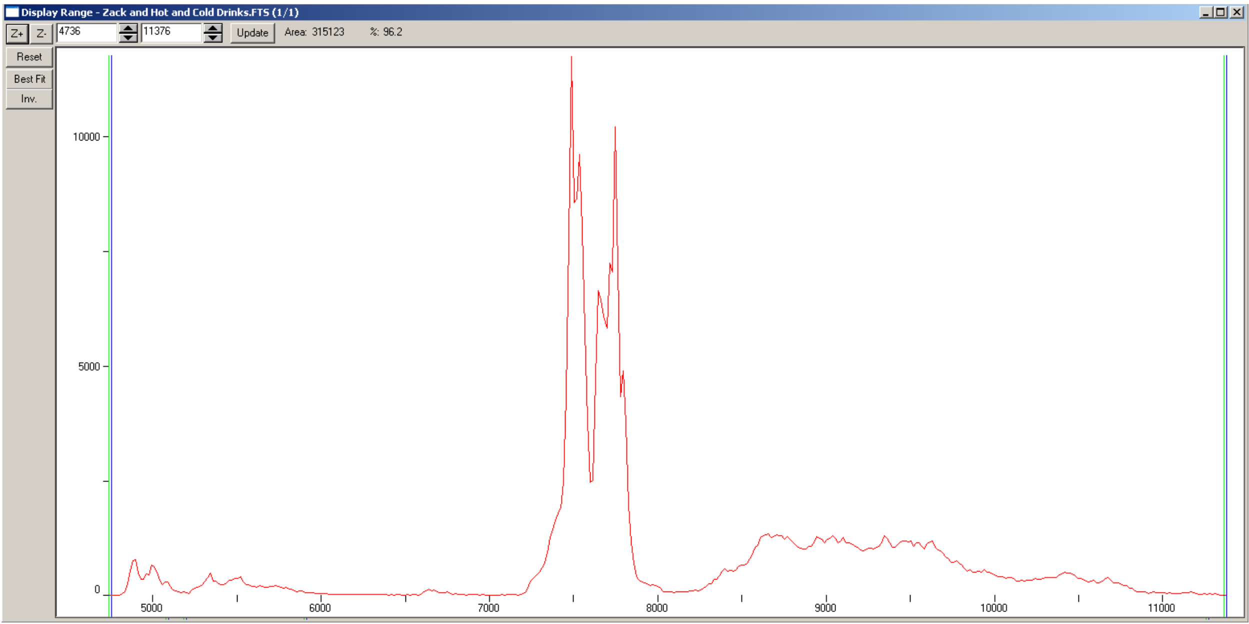


**Fig 3b. Histogram of Figure 3a image**

It is possible to use the linear AGC to enhance contrast in local areas, by moving the display limits around to highlight different parts of the image. Details on the surface of the water bottle can be shown by moving the display limits like so:

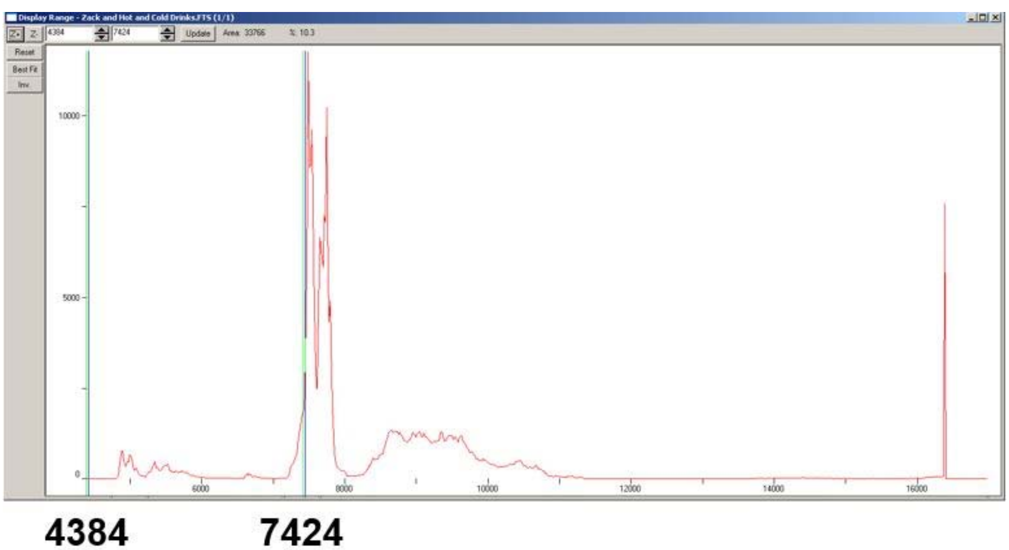


**Fig 4a. New limits on histogram to maximize contrast on just the water bottle.**

The result is that the water bottle looks good and everything else is blown out white:

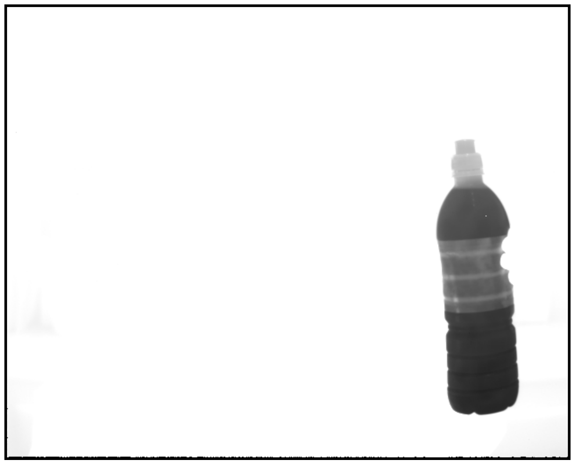

**Fig 4b. Image with maximized contrast on just the water bottle.**

This trick can be repeated to enhance various digital count ranges that encompass notable objects in the scene, but at the expense of the rest of the image:

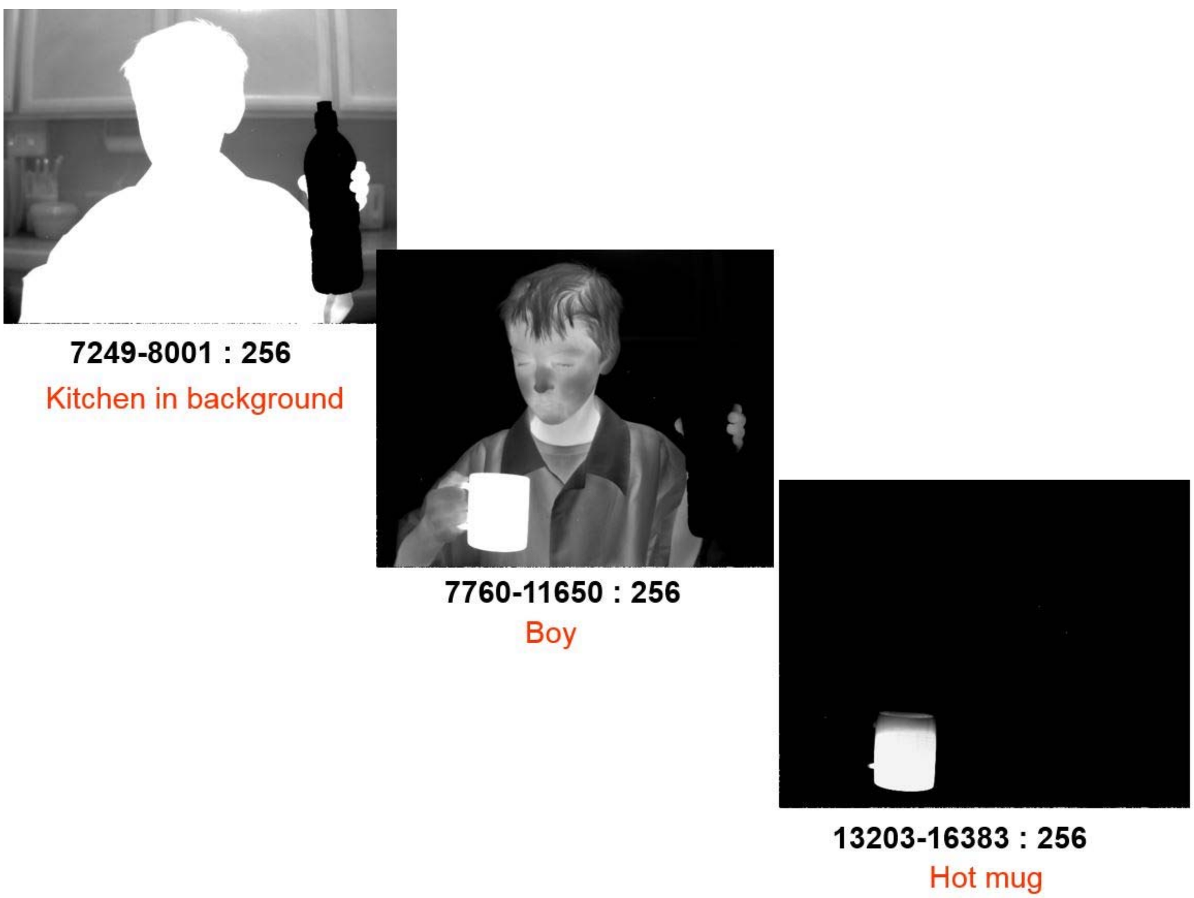


**Fig 4c. Images with maximized contrast on various features**

There is no good way to make this image work with a linear AGC so that contrast is properly enhanced all over. The best way to deal with the gap in the histogram is to use a non-linear AGC, described in the next section.

## Non-linear AGCs

A more sophisticated approach to image display is to use a non-linear AGC. The display level density is determined by the number of pixels within a range of 14-bit digital levels. When there is a low density of pixels, then few or no grey levels are assigned to that range. The mapping between input digital levels and display levels is no longer linear, which makes it harder to judge the relative radiance of objects in the scene, but the contrast is definitely enhanced. ResearchIR gives the user an option called Plateau Equalization. This is a non-linear histogram equalization with a limit imposed on the height of the peaks in the histogram. The peaks get chopped off and become plateaus for the purpose of assigning display levels. For example, in the histogram below, the maximum peak height is about 12,000. The plateau equalization algorithm might treat that as only being 8000 high.

**Fig. 5a. Image of boy with plateau equalization AGC applied.**

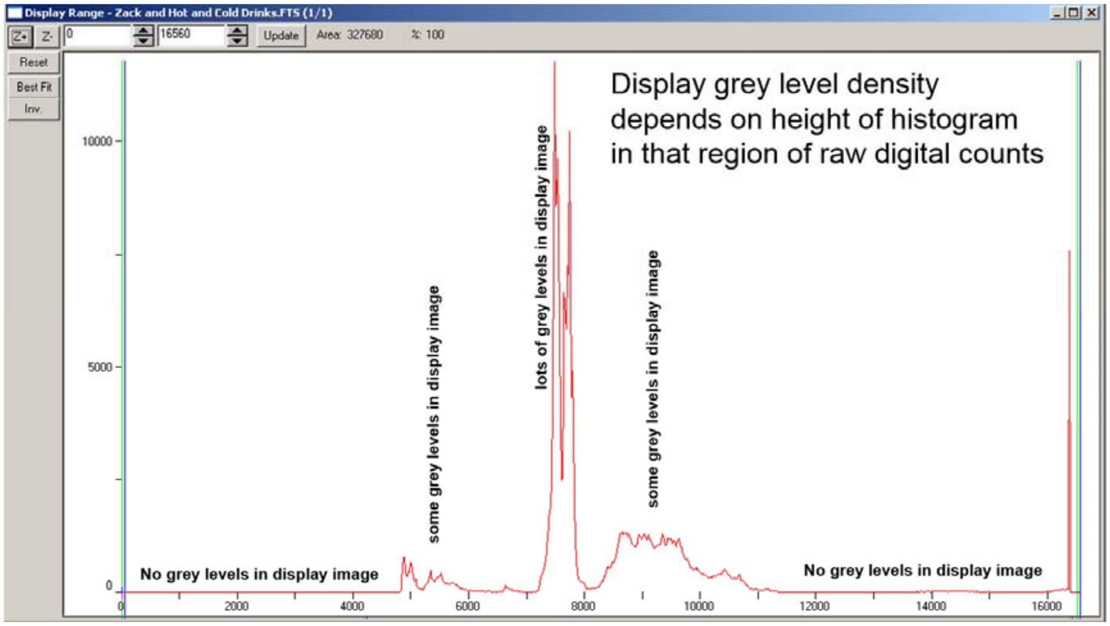


**Fig. 5b. Histogram with description of the display level density**

An even more powerful AGC method is the adaptive plateau equalization or APE. This method divides the image up into sections and then applies a plateau equalization to each section in such a way that the sections match up with each other. It produces very strong contrast everywhere but also can really enhance noise in the image!

**Fig. 6 Adaptive plateau equalization**

Finally, there is an AGC method called DDE which stands for digital detail enhancement. It is essentially plateau equalization with an additional enhancement of edge detail. A pull-down control allows the DDE value to be set from 1 to 7. The results are often useful and less overpowering than APE:

**Fig. 7. Digital Detail Enhancement with scale and DDE set to maximum values**

## Image Palettes

Another option for image enhancement is the use of image palettes other than the default greyscale or "white hot" palette. Color greatly enhances perceived visual contrast. Color palettes have been used to add color to thermal images for many years, and the practice has entered the popular vernacular, with prime examples being the thermal infrared vision of the alien in the Arnold Schwarzenegger film Predator, or the Kevin Bacon film Hollow Man. The purpose of a false-color palette on an infrared image is to increase contrast between different parts of the scene and make certain features of the image stand out. The information a standard color palette conveys is the brightness, NOT the wavelength of the infrared radiation[20].

Hollywood aside, scientific and industrial users of thermal imaging equipment can make use of color palettes to present information to others (particularly in PowerPoint presentations), to detect important thermal events, such as a pressure vessel overheating, and to monitor processes where there are very subtle temperature changes occurring. The last example brings to mind an integrated circuit package that may change surface temperature during operation – I will show an example of this later.

Color palettes are used to convey information to the human eye and brain – no additional information is "created" by applying them to grayscale imagery, and they can often make an image harder to interpret if the color palette is the wrong type! Color palettes are for human viewing only – they are not needed or wanted by machine vision systems or video analytics set up for automatic target recognition. In fact, most video analytics systems want raw 14-bit image data to work with, not the 8-bit display video that is presented to the viewer when using

[20] There are special multispectral IR camera systems where color is used to represent wavelengths, but they will not be discussed here.

ResearchIR software, or when the viewer is looking at an analog video monitor.  Since the vast majority of video monitors are set up to display 8-bit display video signals, the raw 14-bit IR video must first be mapped to 8 bits using an intensity transfer table, or AGC, as it is often called.  A grayscale or monochrome thermal IR image can be hard to interpret when the image contrast is low, because it is difficult for the human eye to distinguish between subtle shades of gray.  The application of a color palette is an attempt to increase perceived contrast, in order for a human viewer to better interpret the image. The standard color palette maps grey levels in an 8-bit image to 256 different colors.  It is certainly possible to increase the number of colors in a color palette, but it is not typically done for commercial IR camera systems. There are cases where a higher bit depth of a color palette is useful to avoid an effect called ***banding***, where the interface between adjacent display levels show hard edges that do not correspond to physical edges in the scene.

### Monochrome Palettes

The most common color palette is called **white hot**.  This palette maps the 256 display levels in an image that has had AGC applied to it over to 256 gray levels, with white being the hottest or brightest output level.  Inverting this palette gives you **black hot**, a palette that sometimes makes images look much better than white hot.  Black hot can make some marine scenes look more pleasing to the eye, as shown in Figure 8a-b:

**Figure 8a.  White hot palette applied to a midwave IR image of a cargo ship**

**Figure 8b. Black hot palette applied to a midwave IR image of a cargo ship. I find this more pleasing to the eye to have lighter-colored water.**

The ocean water looks much lighter than in white hot, and this black hot image is easier for the eye to interpret details in the water.

Another place where the black hot palette is useful is in imagery with a very hot, bright artifact in it, like an explosive fireball. The eye will tend to saturate on the fireball when it is watching live thermal IR video, and other details can be lost. For a dark-adapted eye watching live video, the fireball flash can ruin the dark adaptation. The images in Figure 9a-b show these effects. Details in the grass and the bushes behind the muzzle flash are easier to discern in the black-hot image. The black hot palette is very often used in aerial targeting pods so that the pilot of a bomber aircraft does not get momentarily blinded by the flash of a bomb going off on the ground. In general, military systems do not use color in the display of thermal IR video, except in certain cases where a little color is added to specific features in an otherwise monochrome image.

**Figure 9a. White hot palette applied to a midwave IR image of a muzzle flash from a .30 caliber rifle**

**Figure 9b. Black hot palette applied to a midwave IR image of a muzzle flash from a .30 caliber rifle. It is easier to see details in the bushes and trees and grass in the background, because the eye does not get overwhelmed by the bright white flash itself.**

## Color palettes

Colorful thermal infrared images are very popular with the public, because the color signals to the viewer that one is looking at a thermal image, especially if a familiar color palette is used. The most common colorful palette used everywhere for decades is the **rainbow** palette. This palette increases apparent contrast to the human eye with bright saturated colors of the rainbow in their natural order, with the addition of white at the top and black at the bottom, giving this color order: white, red, orange, yellow, blue, violet, black. This is a very convenient palette, since many people know the correct order of colors in the rainbow, and thus they can quickly determine what is the order of "hotness". A very useful tool to add to a color palette is a color bar which is usually presented on the right side of the image, and which contains the color gradations of the palette from hottest to coldest, with hottest at the top. Figure 10a shows this rainbow palette applied to the ship image in Figure 8. Note that there is fixed-pattern noise that is now much more apparent in the image, particularly in the areas right below the ship. This is a problem with a color palette that enhances contrast – it also enhances image noise. The color bar at the right of the image shows the digital count values that correspond to color values, with interpolated count values assumed by the viewer. In this particular case, the image is not radiometrically calibrated. If it was, then it would be possible to change the counts values to temperature or in-band radiance values for the display in the color bar. This image looks pretty good in this color palette. It is now possible to see gradients in the radiance of the structures on deck, and the structural ribs in the vessel are more apparent because the palette has "thrown contrast" across the ribs and the hull plates attached to them.

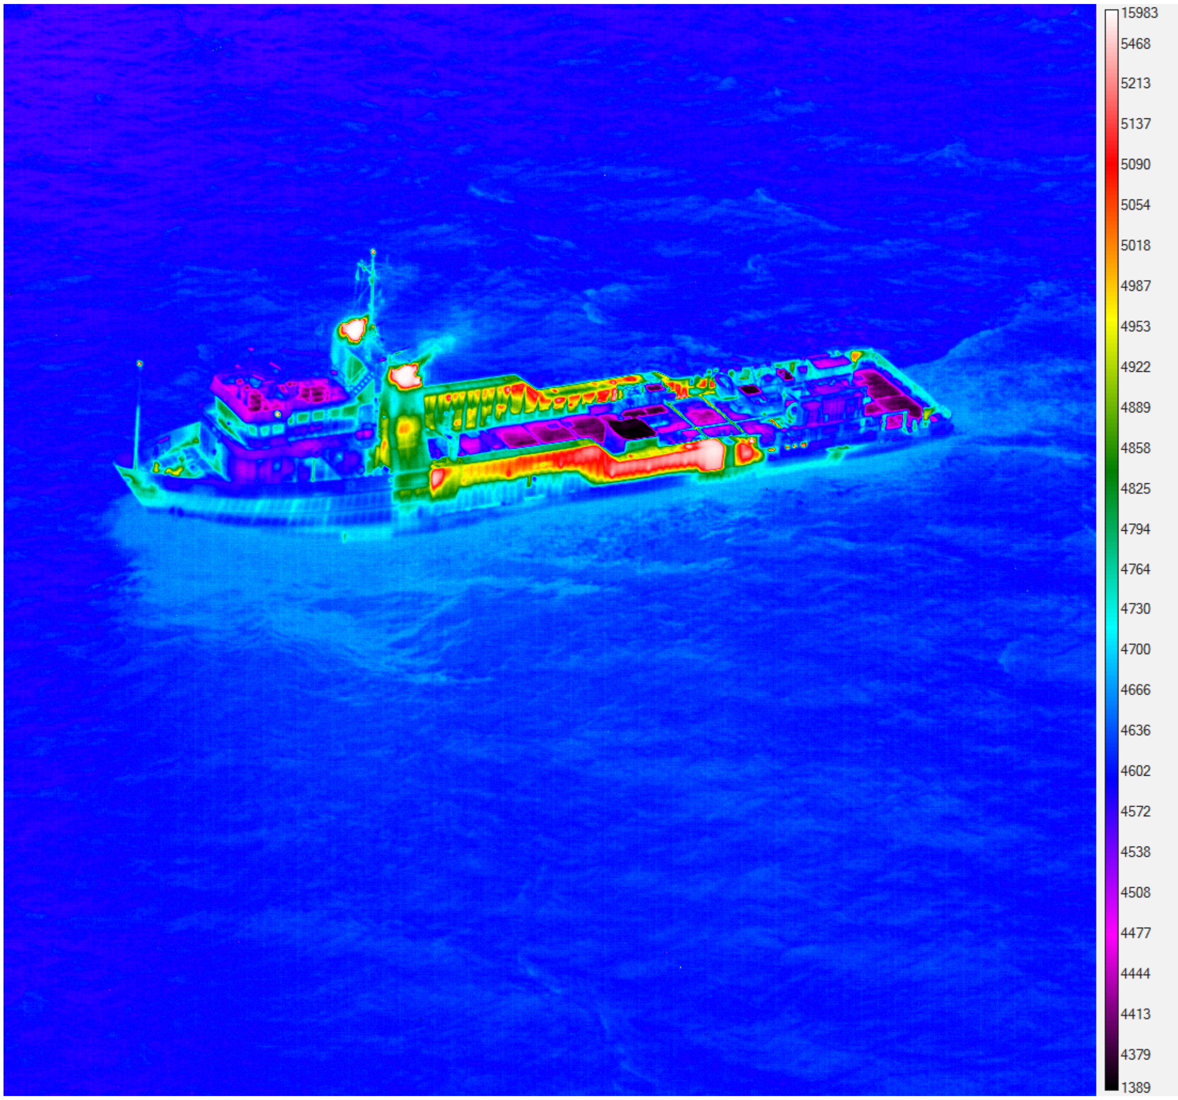


**Figure 10a. Rainbow palette applied to ship image with color bar on the right. The color bar shows corresponding digital count values on ~30 count intervals, making it possible to estimate the relative levels of pixels in the image in raw digital count units. A more useful color bar will display either in-band radiance or apparent temperature, since raw digital counts are not engineering units.**

Figures 10b-d shows the ship structure in close up, with the white hot and black hot palette images for comparison. There is no question that the structural ribs show up much better with the rainbow palette applied. With color palettes, some parts of the image may be improved, and some may be made worse. It all depends on the image and what the desired effect is.

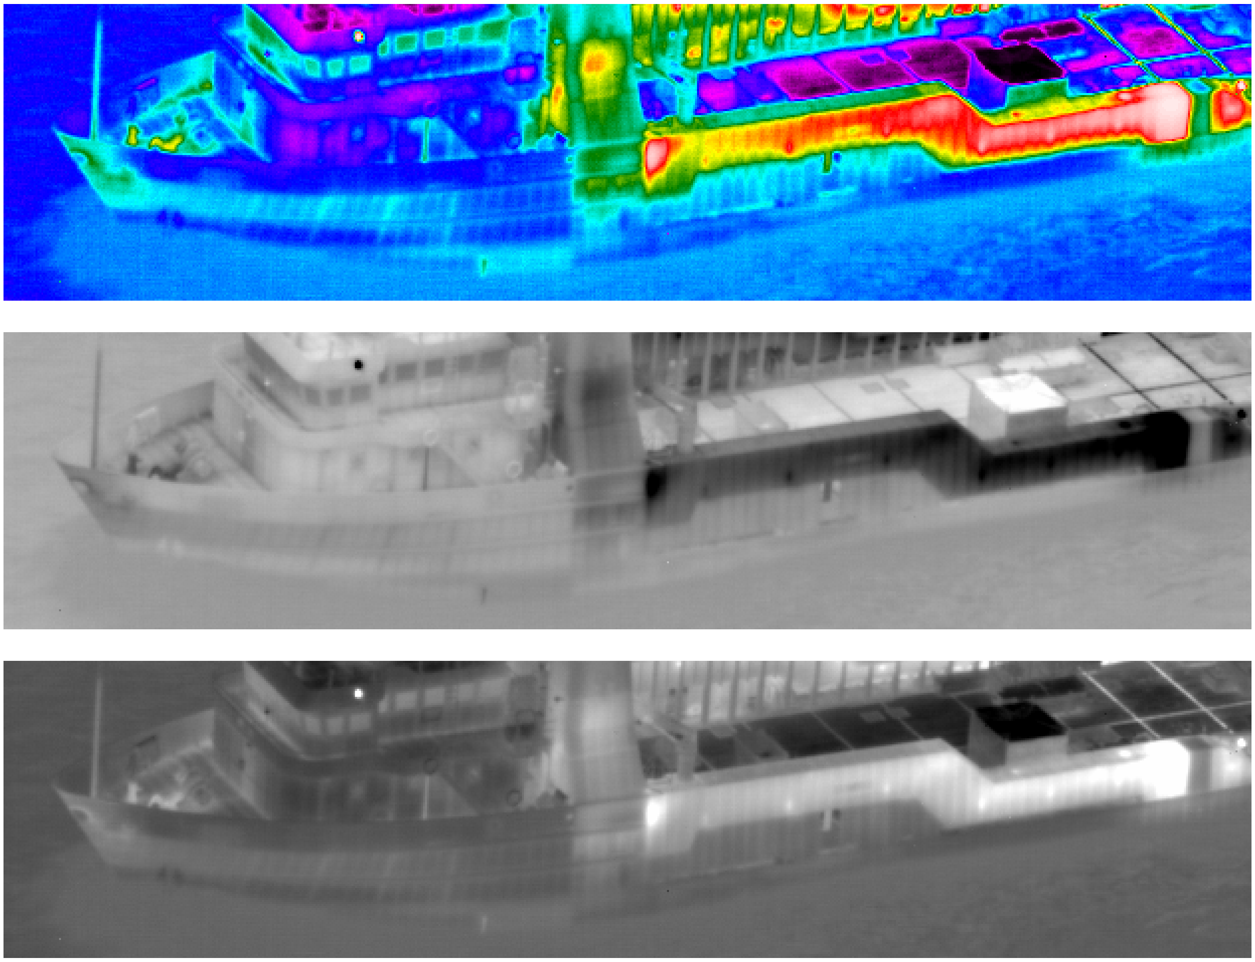

**Figure 10b-d. Closeup of the ship with three different color palettes applied.**

The rainbow palette and other like it may not work as intended on certain images. The image of the muzzle flash from the .30 caliber rifle does not look good with this rainbow palette, as shown in Figure 11a. There is just too much intrinsic contrast in the image, and the trees and bushes and grass are harder to discern. The muzzle flash pixels are clearly saturated, since one can see the fixed-pattern noise in the image where the NUC does not work anymore. This happens when all the pixels in an area of the image are at 16,383 counts before the NUC is applied, as these pixels are. I know this because I looked at the 14-bit image data with the NUC turned off, as shown in Figure 11b. The pink pixels are ones that have digital counts values over 16,382 counts.

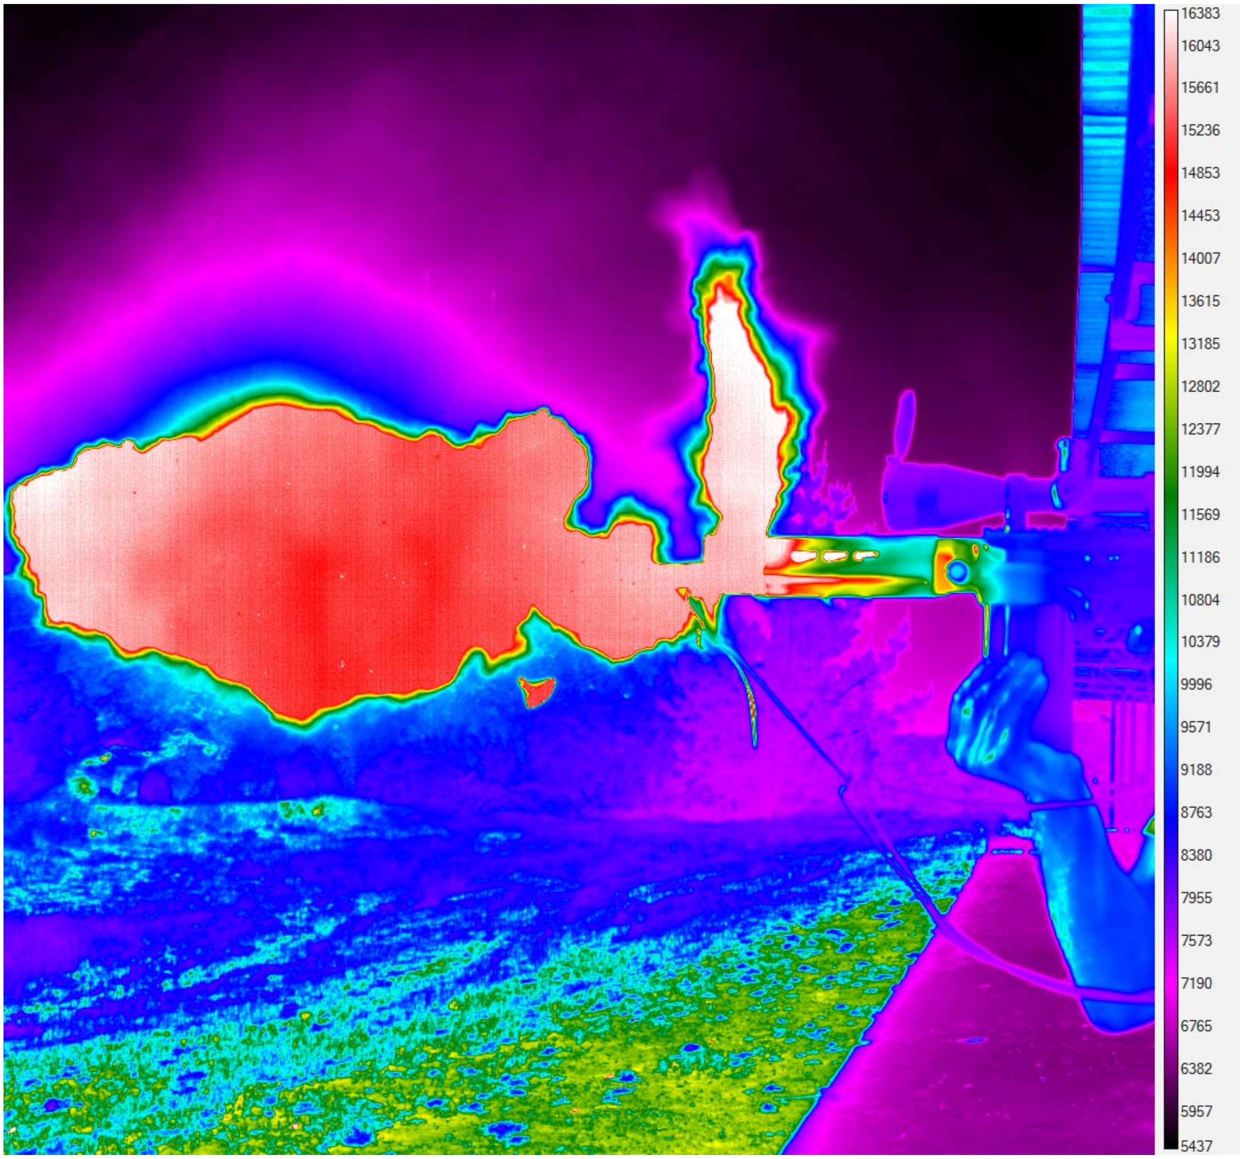


**Figure 11a. Rainbow palette applied to the muzzle flash image**

**Figure 11b. Segmentation filter applied to the muzzle flash image in the white-hot palette**

## Images with high spatial frequency content

The next image we will consider is one that does not look good when processed with any color palette that contains color in it. Figures 12a-c are generated from a midwave infrared image of Santa Barbara taken with a long focal length lens on the camera. There is an abundance of high spatial frequencies in this image, and these frequencies should not be given further enhancement over what a greyscale color palette provides. This image looks confusing to the eye when displayed with the rainbow palette: it is not even clear what one is looking at! The white-hot palette looks OK, but the overall image looks rather dark. The black hot palette seems to do the most justice to the image, though that is certainly a matter of opinion. I find that I can see details in the trees much better than in the darkness of the white-hot palette.

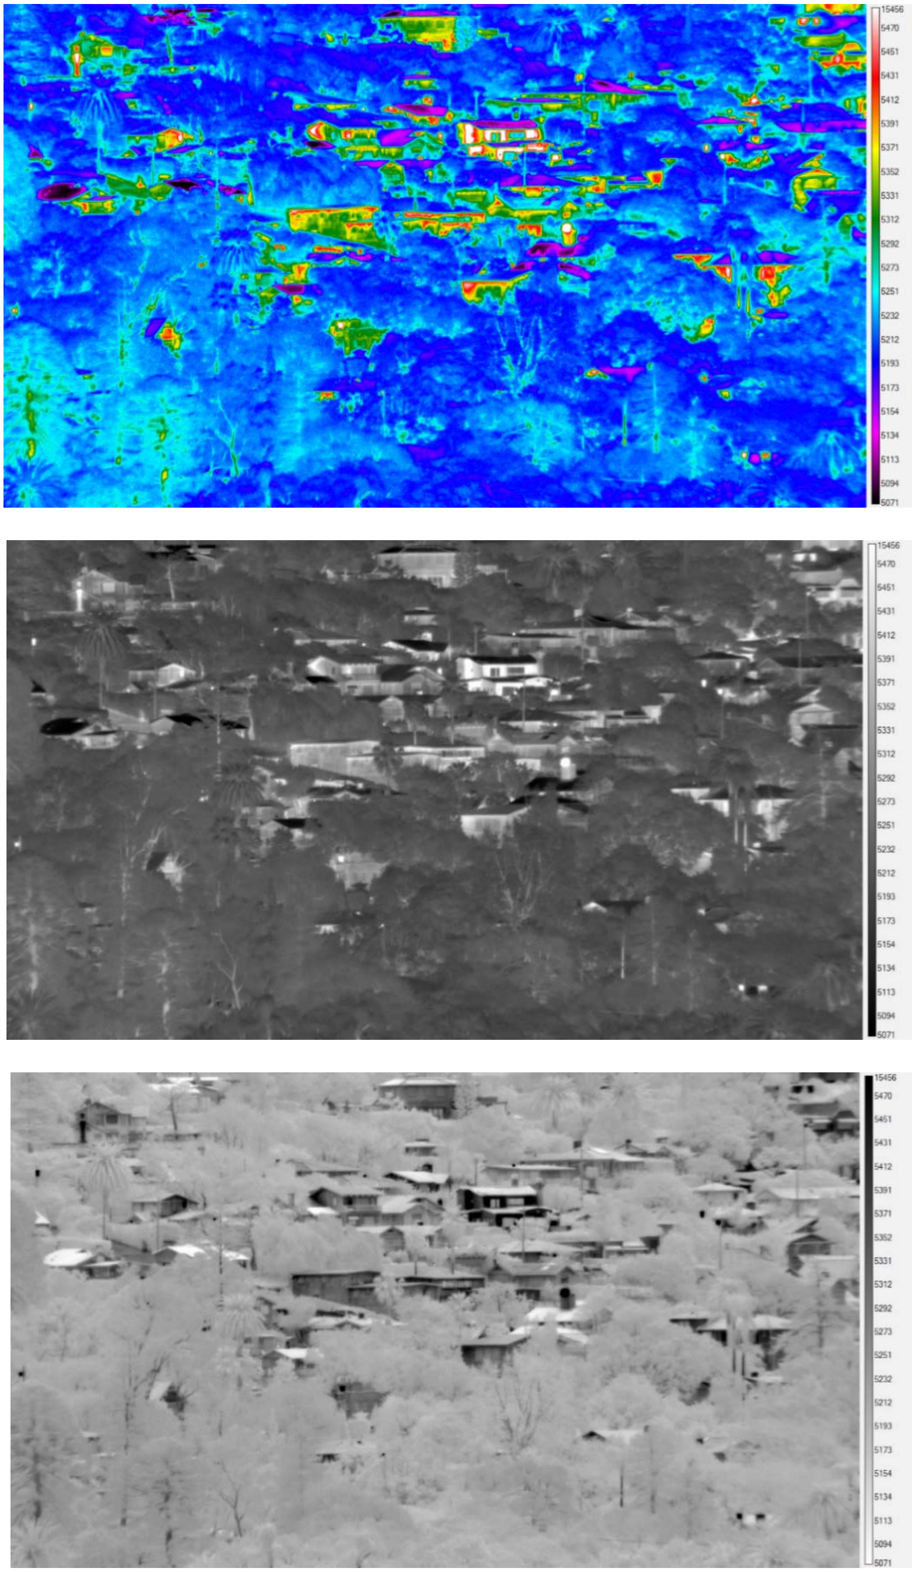


**Figures 12a-c. Rainbow, white hot and black hot palettes applied to a midwave IR image of trees and houses in Santa Barbara, California**

## Popular Color Palettes

Besides the rainbow palette, of which there are many variations, the other very popular palette used across the industry is called **Fusion** within ResearchIR. It is very similar to its close cousin **Ironbow**. The fusion palette is designed to represent some of the colors that one sees in a piece of material (like iron in a blacksmith's forge or a hot coal in a campfire) that has been heated to incandescence. Like the colors of the rainbow, people instinctively know what the colors in the fusion and ironbow palettes represent, since everyone has at one time or another looked into a fireplace or campfire and seen the colors within. The hottest color in the fusion palette is white, then yellow, orange, red, magenta, blue and black. The last two colors suggest coolness. Figure 13a shows the ship image with fusion applied.

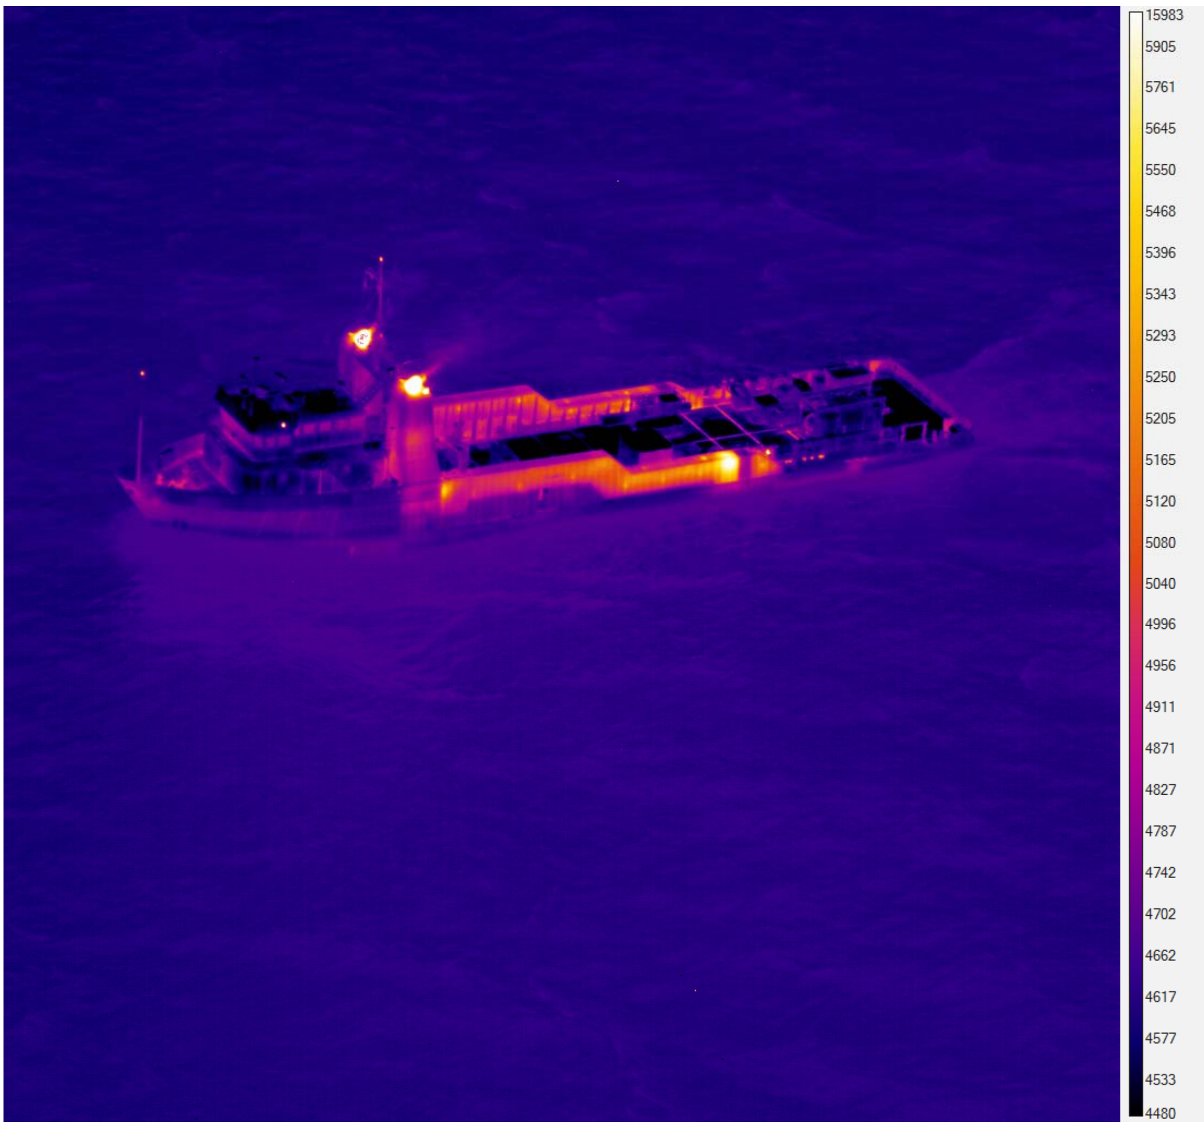


**Figure 13a. Fusion palette applied to the ship image**

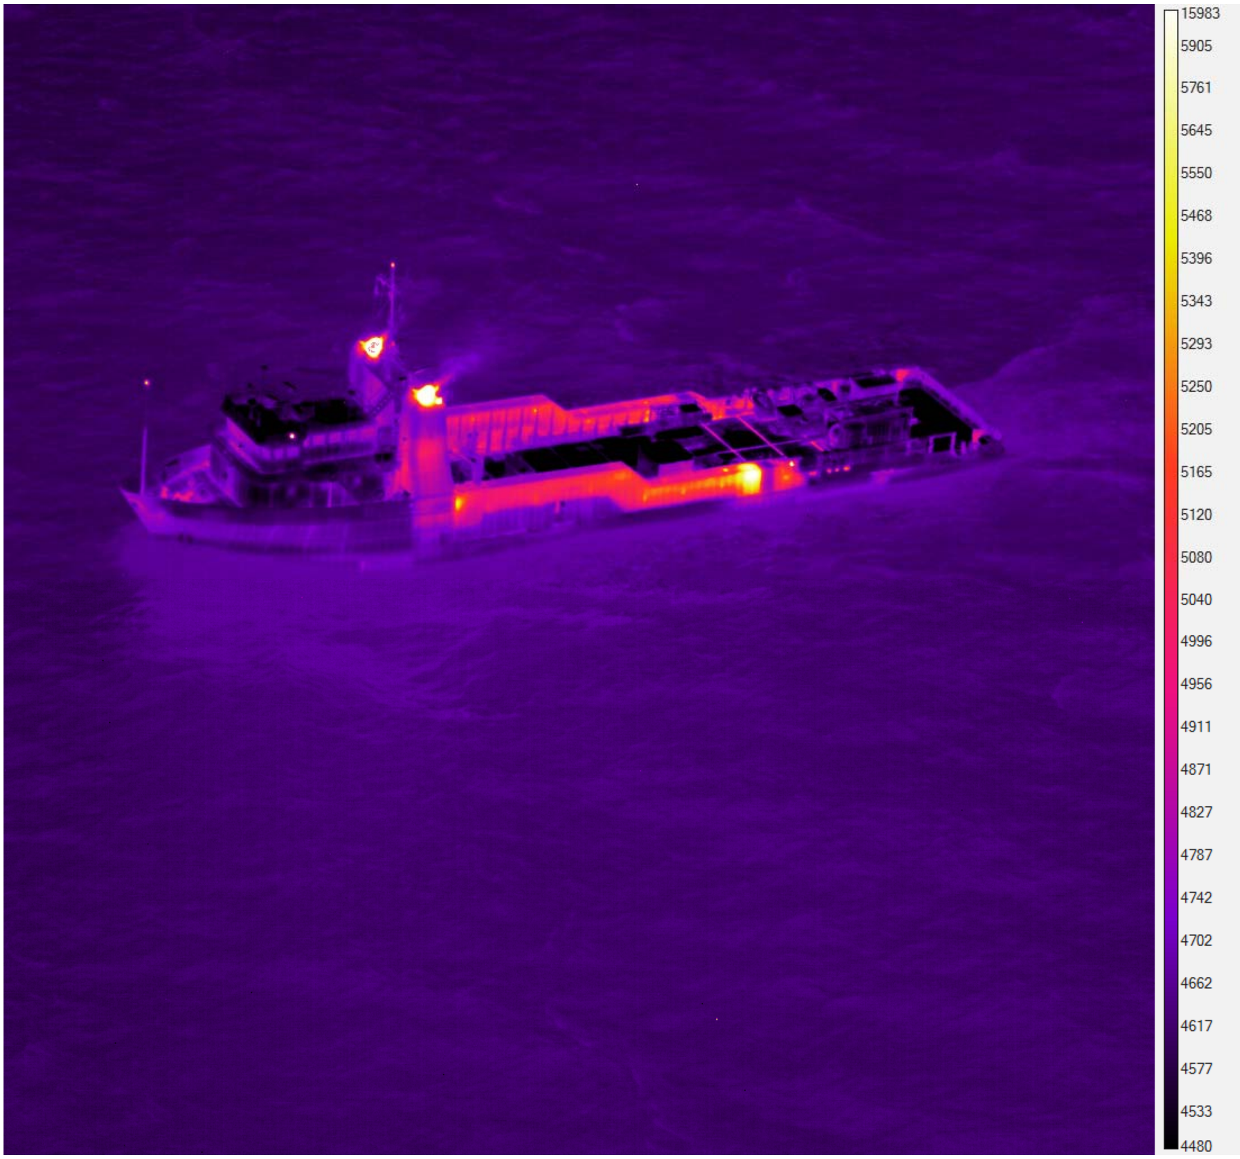


**Figure 13b. Ironbow palette applied to the ship image**

The ironbow palette is different only in that the blue near the bottom of the range is supplanted by purple fading into black, and a wider band of pink in the middle. Again, most casual viewers of a thermal image colored with this palette will "get it" without having to be shown the accompanying color bar.

The palette known as "1234" is popular with many users of ResearchIR software. The scheme of this palette is that there are two classes of colors: warm and cool colors. The warm colors are white, yellow, orange red and deep red. The cool colors are green, cyan, dark blue, purple and black. Like fusion and ironbow, there is an instinctive understanding of the order of colors in this palette, as shown in Figure 13c below:

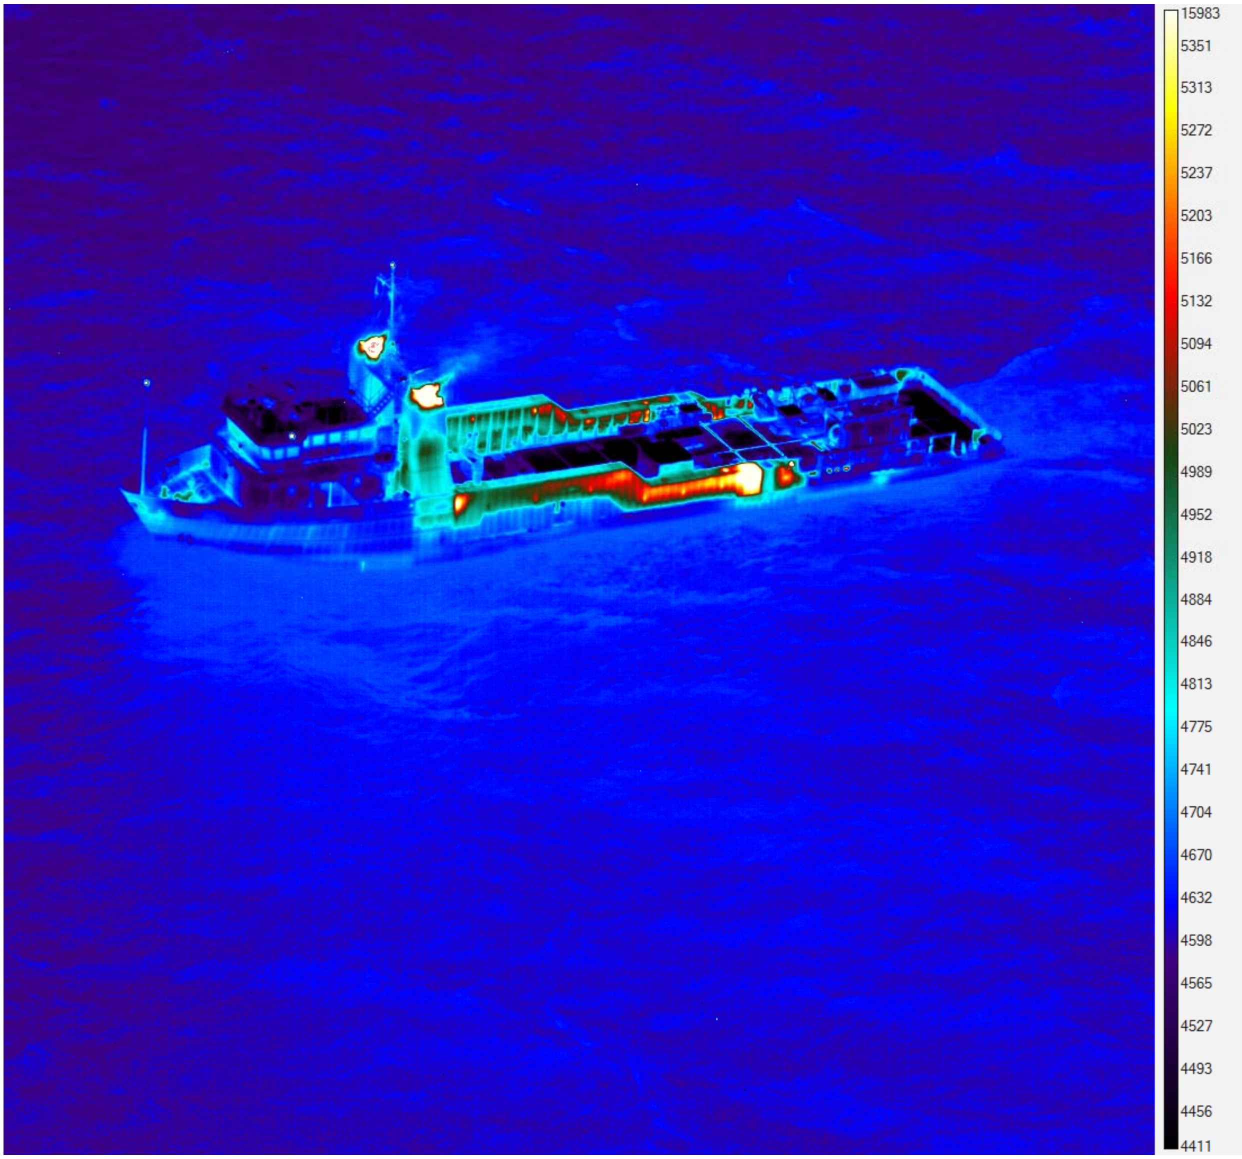


**Figure 13c. 1234 palette applied to the ship image**

NOTE: I have adjusted the upper and lower limits of the 14-bit histogram carefully to get the most pleasing images to my eye. The histogram limits are quite different from the limits set before I exported this image in rainbow palette. The AGC applied by ResearchIR does not always give pleasing results without operator intervention. For example, in Figure 14a, the fusion palette has been applied to the raw image with automatic AGC settings. The histogram limits are set from the image (From Image) and are quite wide limits, as shown in Figure 14b. The AGC type is PE Scale, which stands for plateau equalization, a non-linear intensity transform table generated from the image histogram which distributes the display levels in a non-linear fashion.

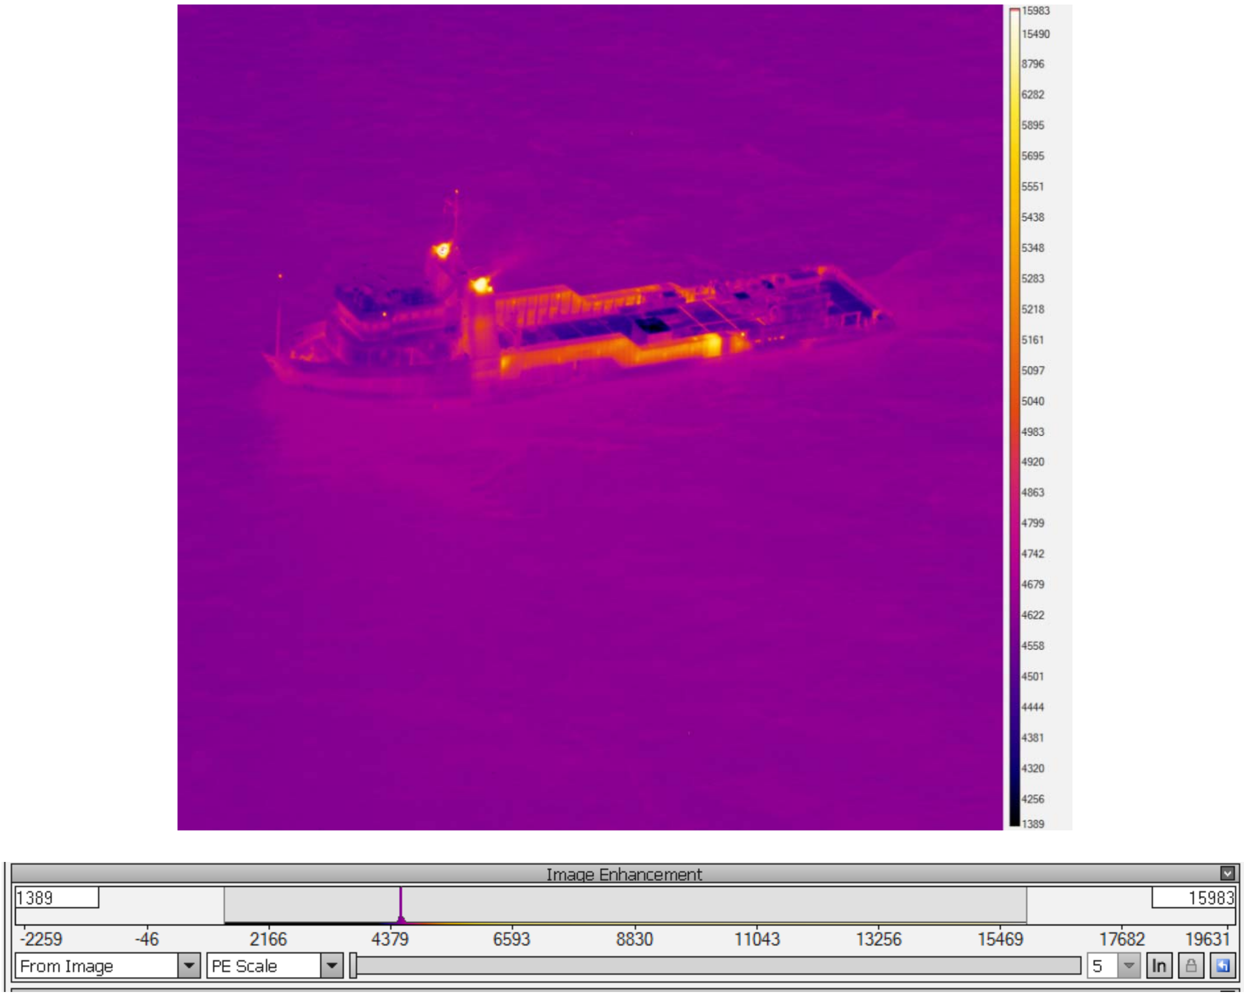


**Figure 14a-b. Automatic AGC applied to the ship image with the fusion palette. The histogram limits are quite wide, mostly because the hot exhaust stack on the ship is up at 15,983 counts. The camera operator can greatly improve the appearance of this image by simply manually adjusting the histogram limits, as shown in the next figure.**

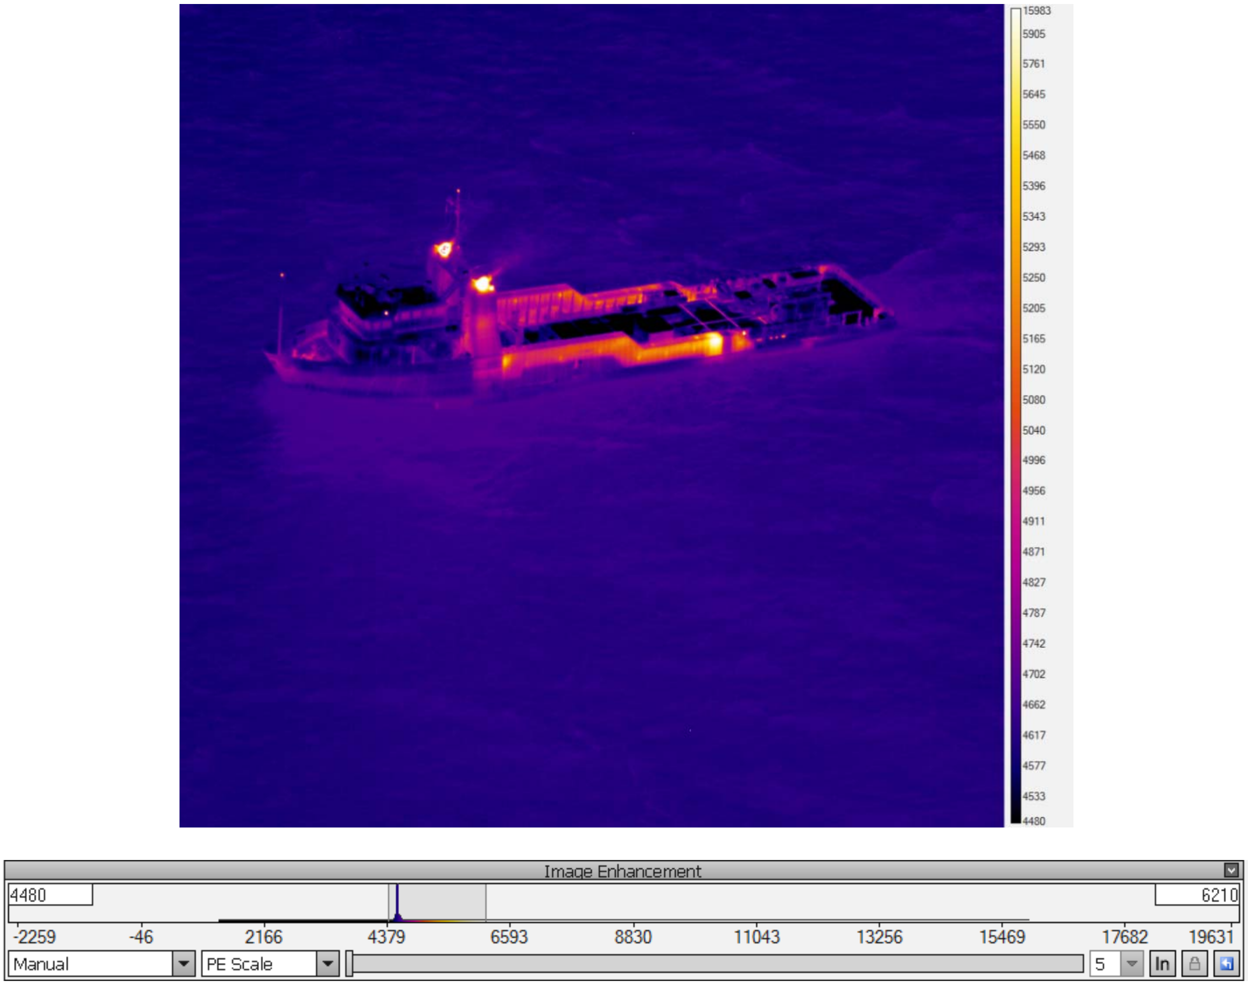


**Figure 14c-d. The display limits for the histogram of this image have now been adjusted to tighter limits, so that the bottom of the palette (black) starts at 4480 counts. Anything below that count value is mapped to black. The upper limit is set to 6210 counts.**

## Hybrid Palettes with Both Color and Grayscale

A very useful type of palette combines color and grayscale into one palette. The most well-known palette of this type is called **Ice_Fire**. The premise of this palette is that everything is represented by the white-hot palette, but then some range of the hottest and another range of the coldest pixels are colored solid red and solid blue respectively. Hence the Ice_Fire name. Figure 15a shows this palette applied to the ship image with 4329 and 5454 counts being the grayscale range. Any pixels below 4329 and any above 5454 counts are colored pure blue or pure red appropriately.

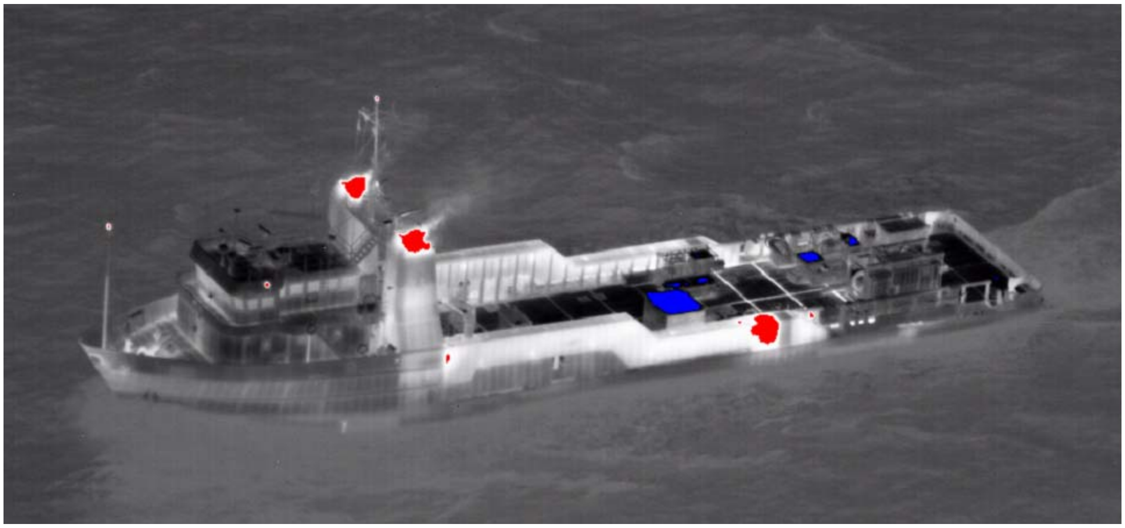

**Figure 15a. Ice_Fire palette applied to the ship image with histogram limits set to 4329 and 5454 counts.**

Applying the Segmentation tool in ResearchIR to the white-hot image with the same count limits gives essentially the same results:

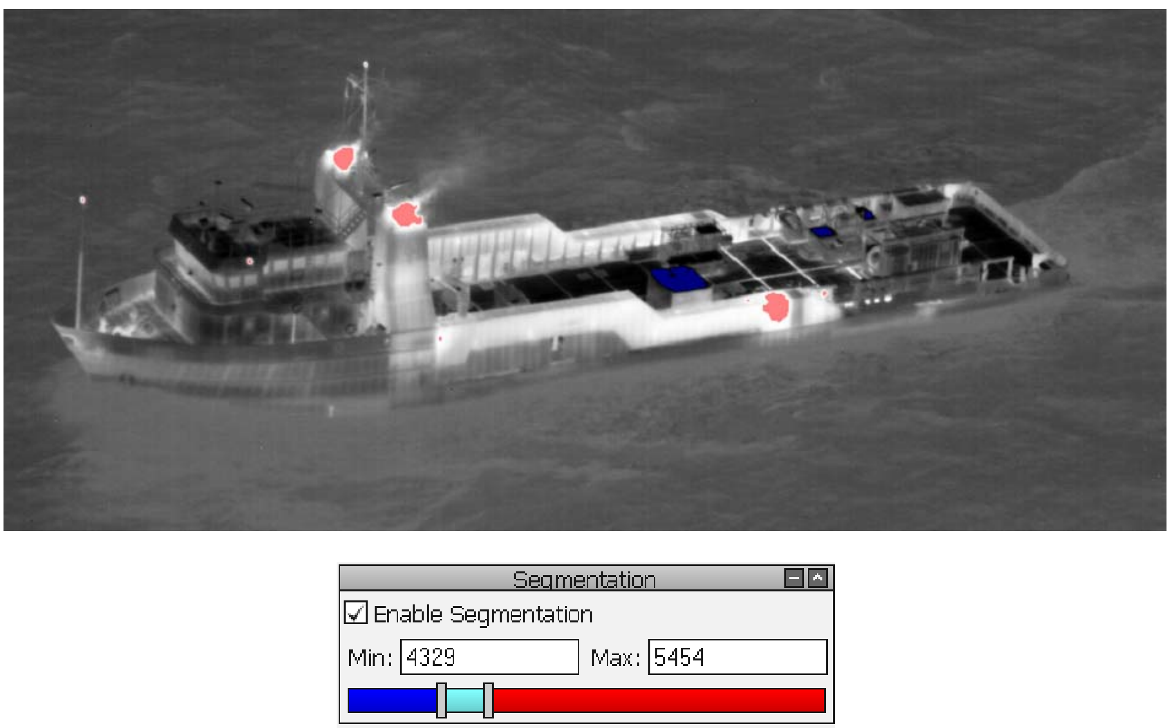


**Figure 15b-c. The ResearchIR Segmentation tool applied to the ship image with limits set to 4329 and 5454 counts.**

The problem with this palette is that it wipes out any details in the image in those "ice" and "fire" regions. I invented a set of new palettes that I call **graded fire and ice** which have the advantage of putting color gradients in the hottest and coldest parts of the image. Here is a classic example of an image that benefits from these palettes. The following image in Figure 16a was taken from a helicopter above the Oxnard California airport about 1 hour after sunset. The image is presented with the standard Ice_Fire palette, and then variants of the Graded Fire and Ice palette.

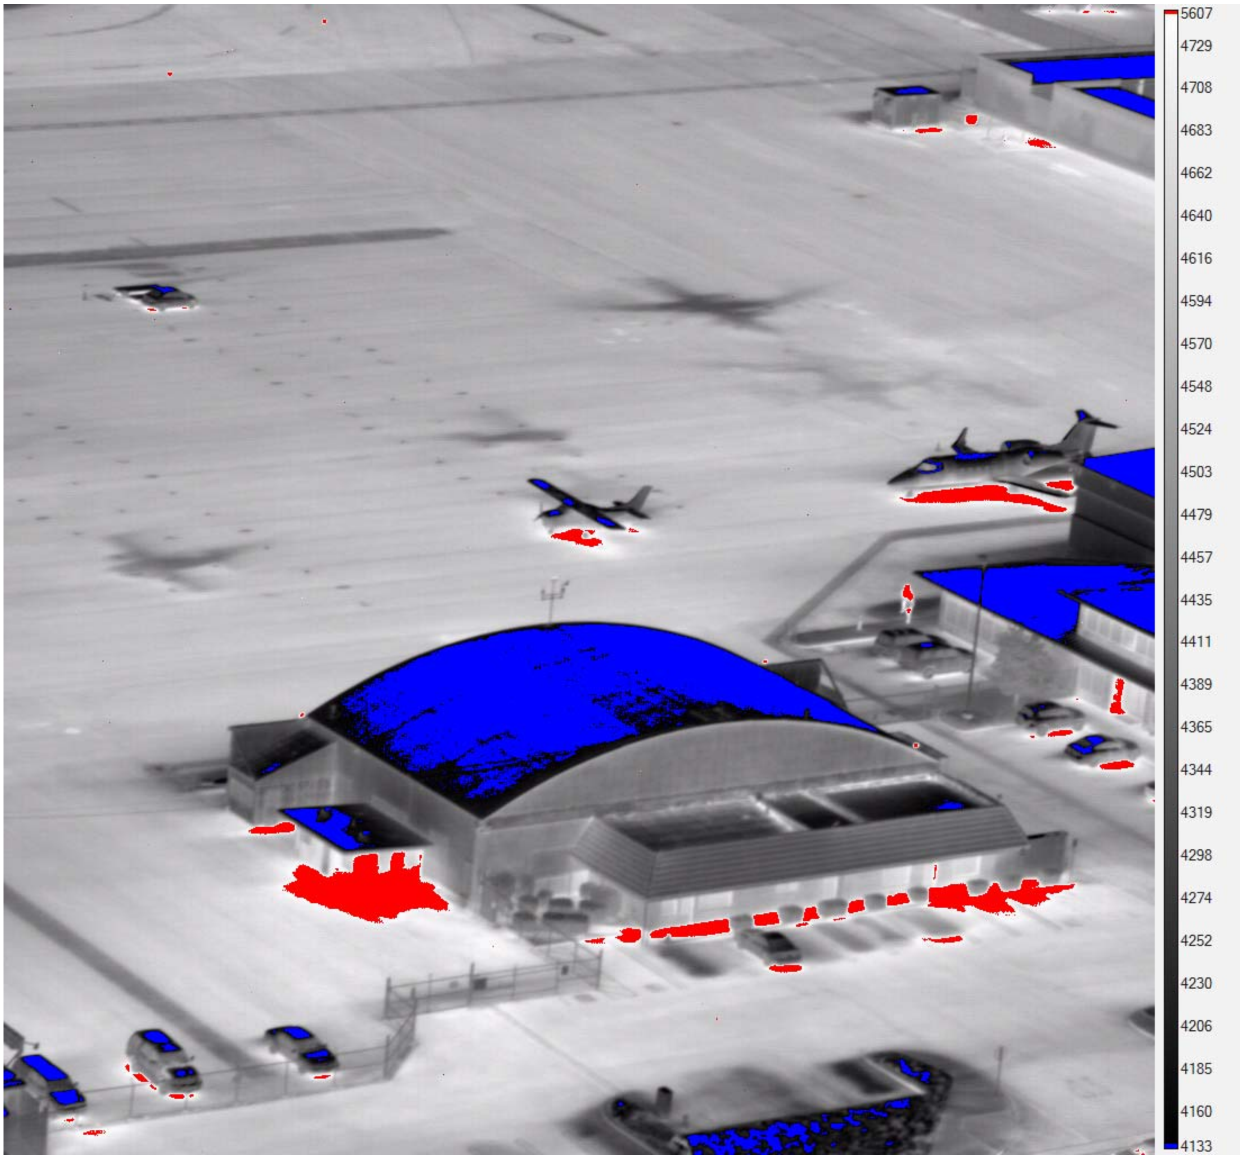


**Figure 16a. Ice_Fire palette applied to an image of the Oxnard Airport. The coloration is useful to quickly cue the eye to the coldest and hottest parts of the image, but all spatial detail is lost in the areas with the color, due to the fact that they are solid colors.**

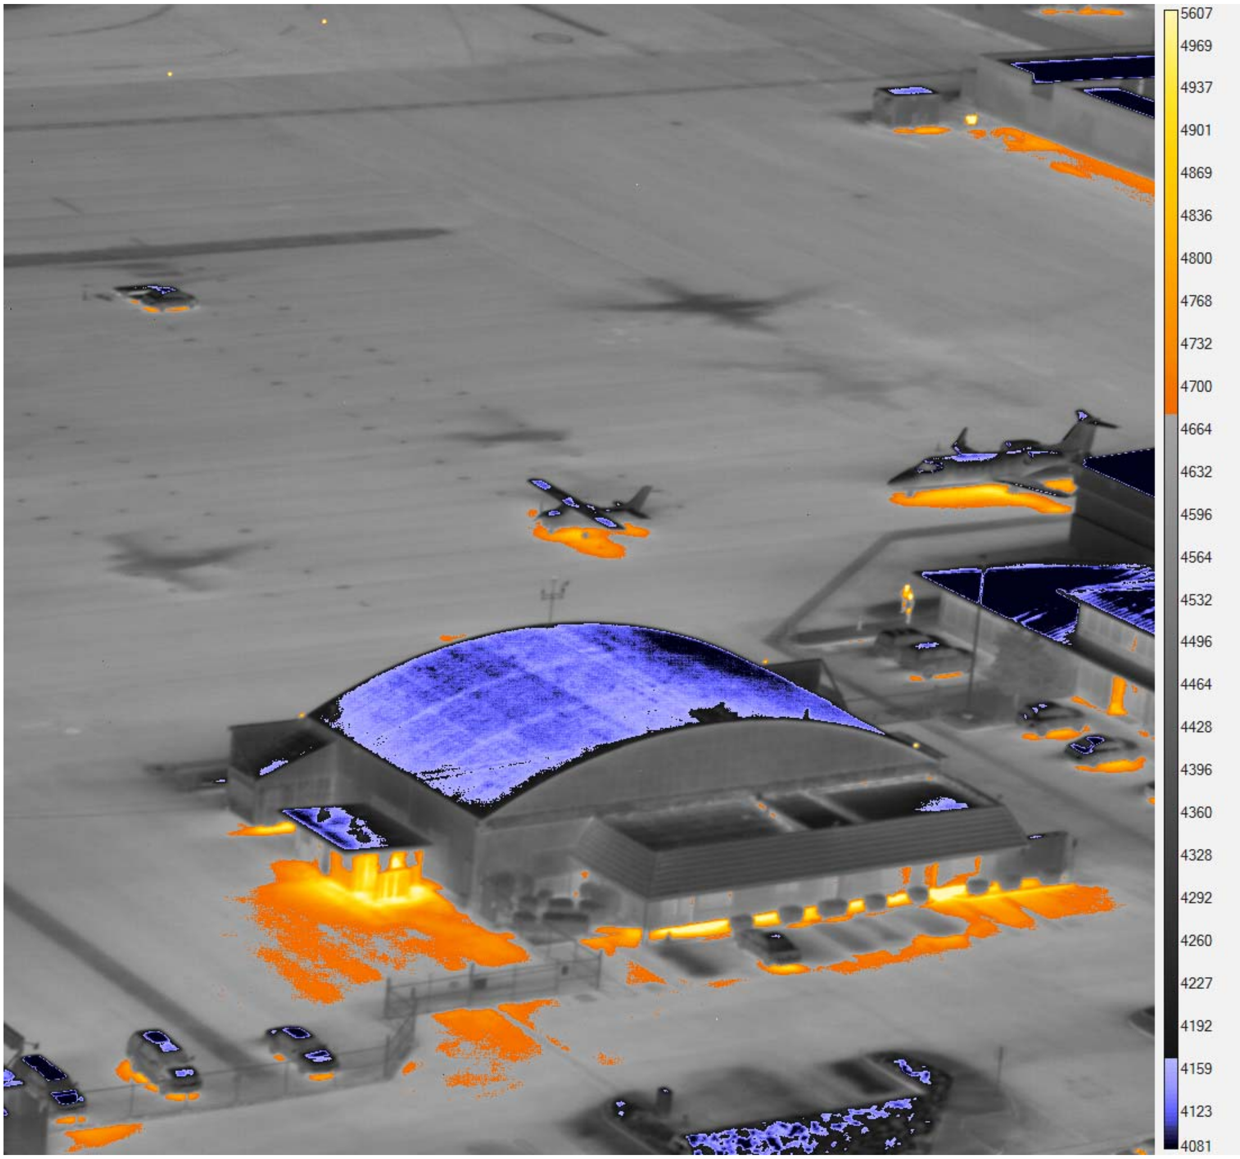


**Figure 16b. The same image colored with a palette called GFIV7, which stands for graded fire and ice version 7. The histogram limits are different from the limits in Figure 16a. Now one can still see the structural details in the hangar roof, and the entranceway to the hangar where a lot of heat is coming out.**

The trick with palettes like this is that they often require manual adjustment of the histogram levels to properly accentuate the details that one wants to see. This class of palette is more useful for post-processing of single images than it is for a video stream from a dynamic scene. One way to use this class of palette on a dynamic video stream is to put the camera into a radiometric mode, then set the limits of the histogram to create colors in a desired apparent temperature range or radiance range. Figure 17a shows four men walking in the woods at night with the GF White Hot palette from ResearchIR applied. The color starts for apparent temperatures greater than 22.1 °C and goes up from there. This palette will enable the user to identify the objects very quickly in an image that are in a certain apparent temperature range. This is sometimes called a “people-finder” palette.


**Figure 17a. Midwave image of four men walking in the woods at night. The color palette is GF White Hot. The AGC mode is linear with the histogram limits manually set to 6.3 °C and 30.2 °C. Anything in the image with an apparent temperature below 6.3 °C will be mapped to pure black. Anything above 30.2 °C is colored pure white.**

ResearchIR offers the Isotherm tool which can be set to color in an interval of temperature, radiance or digital counts. Figure 17b shows the same image with the color palette set to White Hot and an isotherm interval of 22 to 28 °C. Every pixel in the image in this temperature range is colored a solid green color. A possibly more useful isotherm variant puts a color palette across the isotherm interval instead of just a solid color. This is not available in ResearchIR at this time.


**Figure 17b. Midwave image of four men walking in the woods at night with an isotherm interval applied between 22.5 °C and 28 °C**

## High Contrast Color Palettes:

Many thermal images have relatively low thermal contrast in areas of interest, but high contrast overall. A case in point is a powered circuit board, with integrated circuit packages that show heat transfer and diffusion from the chips within. The chip cases themselves can be relatively low contrast, but the overall variations in the image brightness can be quite large. Standard color palettes may not be very effective in showing contrast. Here is a circuit board imaged with some standard color palettes applied – the color bars have been omitted for clarity. In each case, there is little thermal contrast to be seen on the main integrated circuit. But the information is there – we just need to either enhance it with a region of interest AGC, or else we need to use a "stronger" color palette to enhance the displayed contrast.

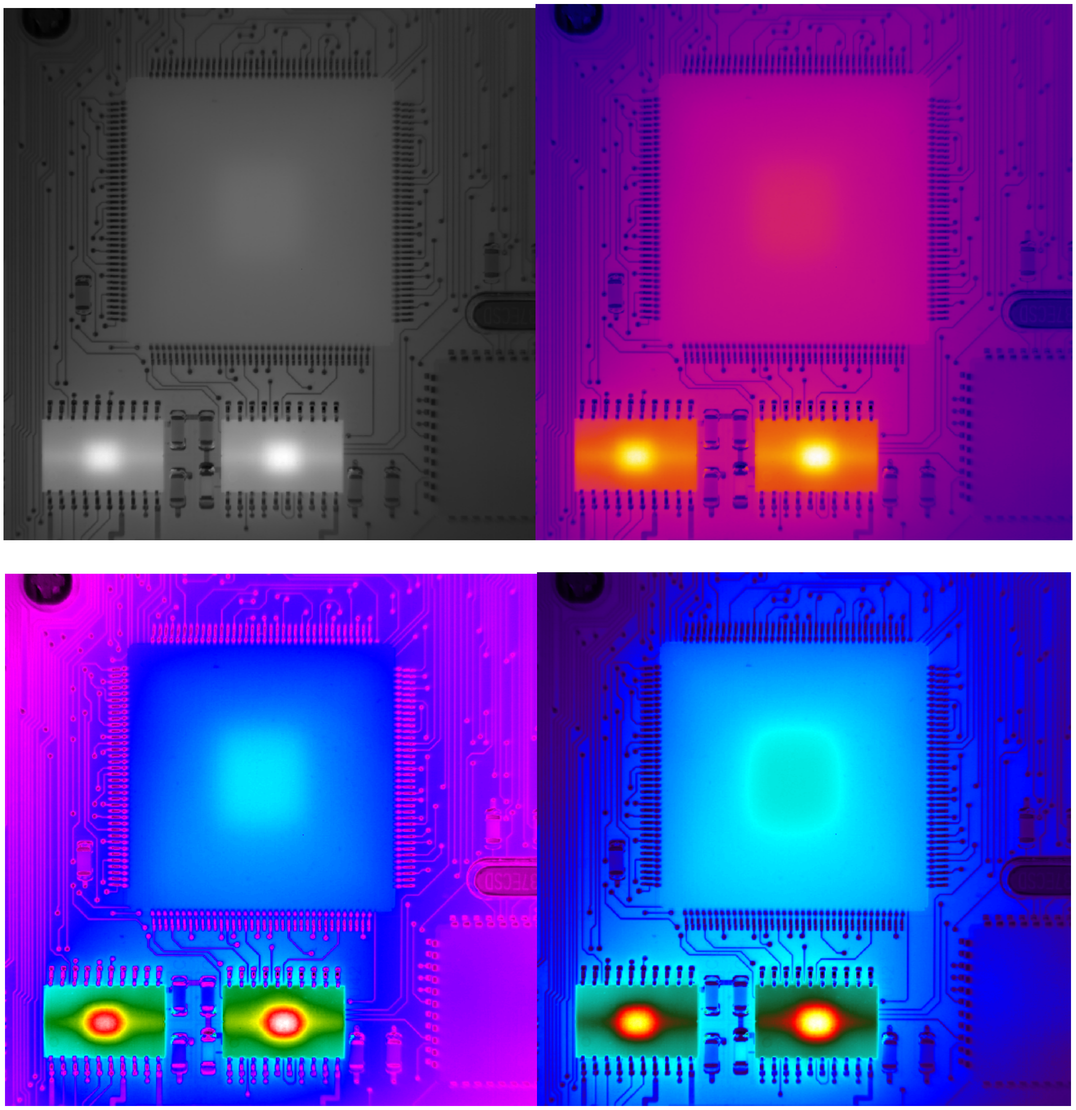

**Figure 18a-d. Four views of a circuit board in, from top left clockwise, greyscale or white-hot palette, fusion, 1234 and "RainbowGraded". The display contrast tends to increase with the higher number of base colors in the palette.**

A region of interest can be applied to the AGC to make the central IC "pop" in contrast. This enhances the central IC, but now the two hotter ICs in the lower part of the image are "blown out".

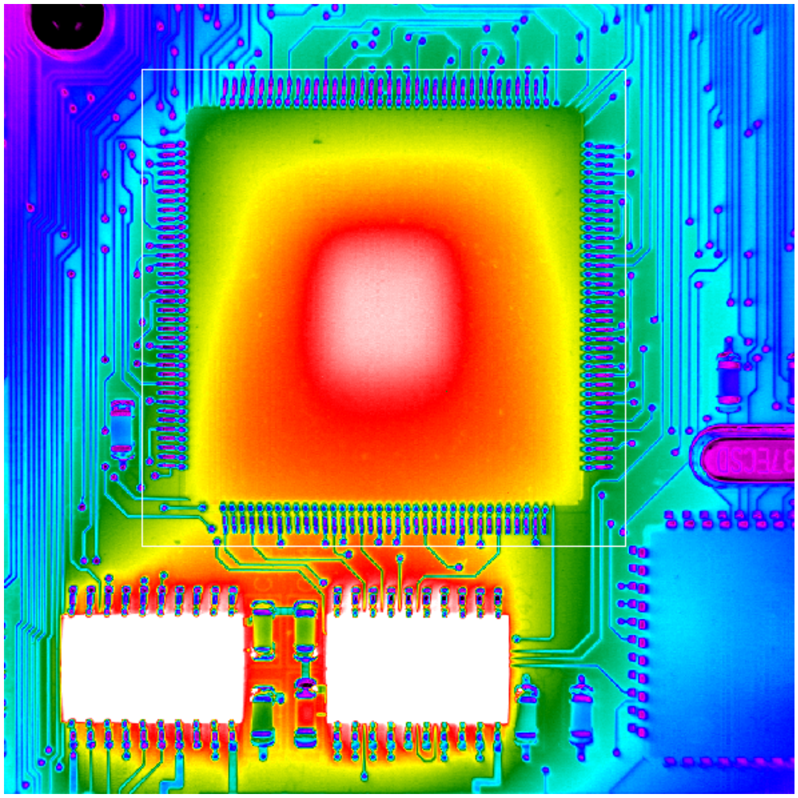

**Figure 19. Circuit board image with rainbow graded palette and AGC limits set by the pixels in the white region of interest.**

Some novel palettes were developed by me to enhance contrast in an image like this. These special high-contrast palettes use color theory, putting complementary colors in adjacent color pairs to enhance perceived color contrast. These three images show this phenomenon. Magenta and green, cyan and red, and yellow and blue are complementary colors, as shown in this image from Wikipedia:

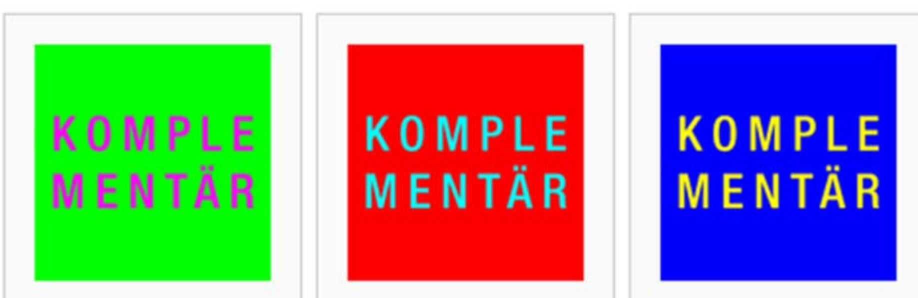


**Figure 20. Complementary colors used to make new palettes.**

Three palettes developed from these base colors and mixtures of them are shown below applied to the circuit image:

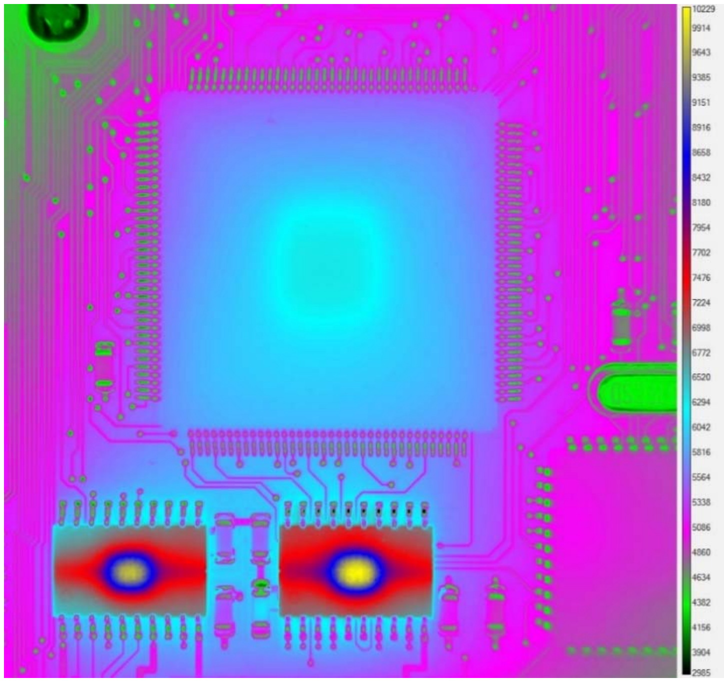


**Figure 21a. Color Wheel 6 palette. There are 6 base colors, made up of either primary colors or mixtures of two primary colors.**

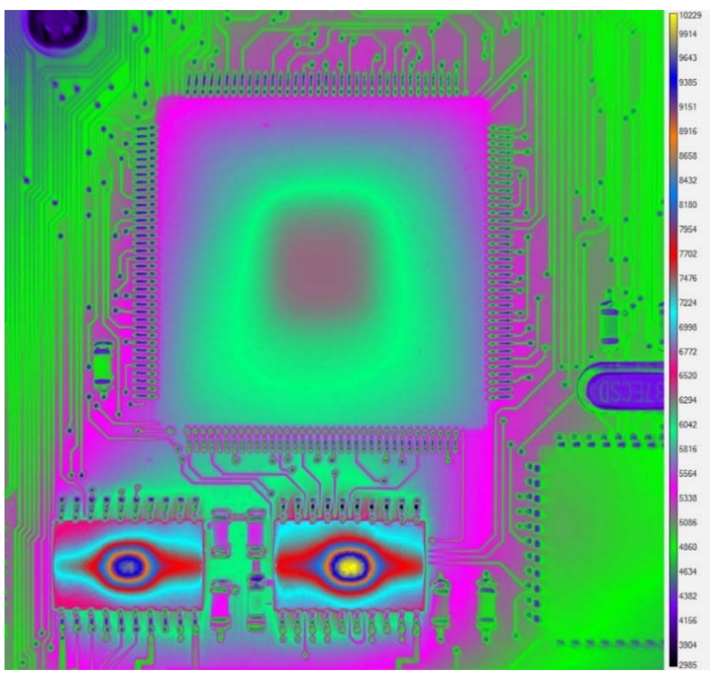


**Figure 21b. Color Wheel 12 palette. There are 12 base colors, which include primary colors, mixtures of 2 and mixtures of all three primary colors. You need a good monitor to distinguish every color from similar ones, but all colors are unique.**

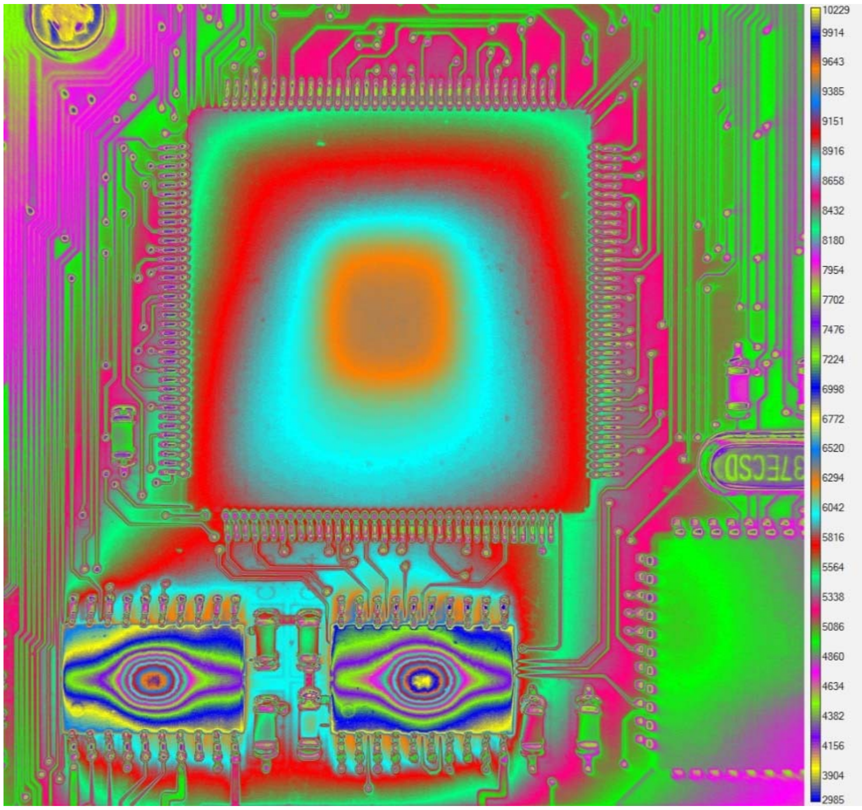


**Figure 21c. Color Wheel 26 palette. There are 26 base colors, which include the colors of color wheel 12 repeated again.**

The first palette in the images above is Color Wheel 6. It has 6 colors with black placed as a bottom color. The color order is counterintuitive – dark blue is never a hot color in other palettes, but with a little practice, it starts to become familiar. To really show the temperature gradients in the ICs, more base colors are needed. Color Wheel 12 uses 12 opposing pairs colors from the color wheel, with black as a bottom base color. It is harder to look at this image and determine what is hot and what is cold, because the colors start to look more alike. But the big advantage of this palette over Color Wheel 6 is the enhanced contrast.

Color Wheel 26 abandons any attempt for 1 to 1 mapping of colors to temperatures – it is basically color wheel 12 repeated twice, so there is ambiguity about what a color means. But it greatly increases contrast everywhere in the image all at once. Note the "banding" effect mentioned earlier in this paper on the chip in the lower right of the image. This is caused by the fact that there are only 256 discrete colors in the palette – a relic of the days when computers only displayed 256 colors. Higher bit depth color palettes reduce this effect, but it has not as yet been implemented in ResearchIR software.

**Figure 22. Banding seen in a closeup detail from Figure 21c.**

Another interesting feature of color palettes with many closely spaced based colors is their ability to impart "motion" to scenes that are changing temperature. Here are four images of a circuit board taken at 10 second intervals as the components warm up. The colored bands move across the IC and animate the scene in a way that simpler color palettes do not.

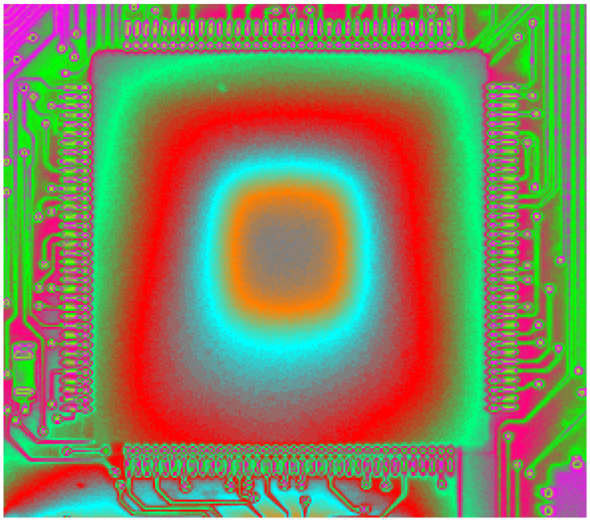

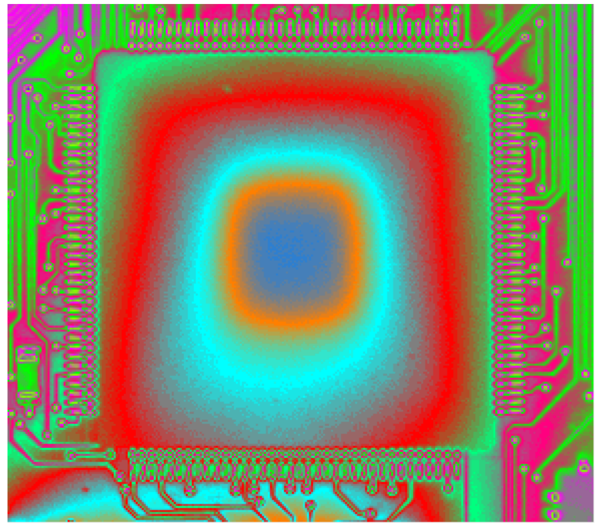

**T = 60sec** **T = 70sec**

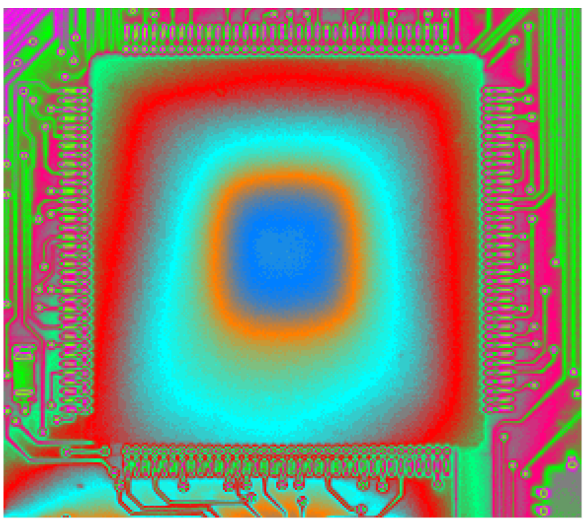

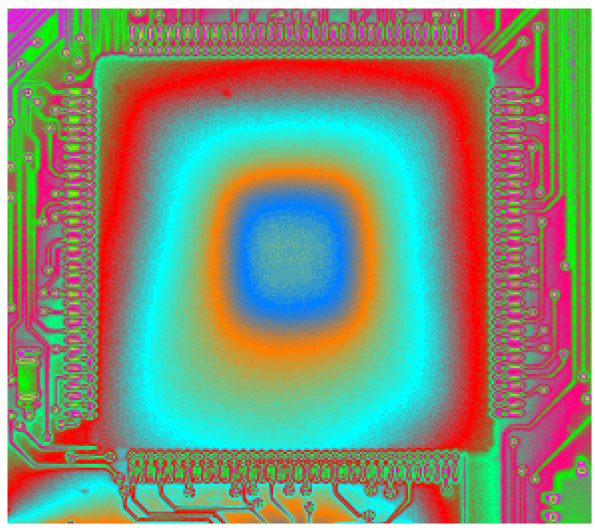

**T = 80sec** **T = 90sec**

**Figure 23a-d. A time series of midwave IR images of a heating integrated circuit package on a circuit board. The color wheel 26 palette causes the viewer of this time lapse movie to perceive motion caused by small temperature changes.**

## Summary

AGC or automatic gain control is the FLIR term for the conversion of 14-bit raw images to a form that is suited for 8-bit image displays. The conversion is sometimes done by the use of an intensity transform table or ITT. The ITT is generated continuously and applied to the displayed images to map 14-bit data to 8-bit display data, though there are many other ways to use an ITT. Color palettes are a means to increase perceived contrast between a target and the scene, and they also can make IR images look better and "sexier" for presentations. For most scientific and especially military applications, color palettes should be used with caution to avoid confusing the eye of the viewer.  Color can also make an image with high spatial frequencies very confusing to look at.

# Chapter 9 : User Calibration Step-By-Step Process

User Calibration : The Setup ........................................................................................................ 9—359

## User Calibration : The Setup

Here is a step by step account of a user calibration. Before I started, I looked to make sure that ResearchIR would not record ATS files. I want to record SFMOV files, because they have separate metadata files (*.inc and *.pod), which makes it far easier to change out the include file to apply a calibration.  Here is where that is done – you just uncheck the checkbox shown below in the red rectangle and hit OK:

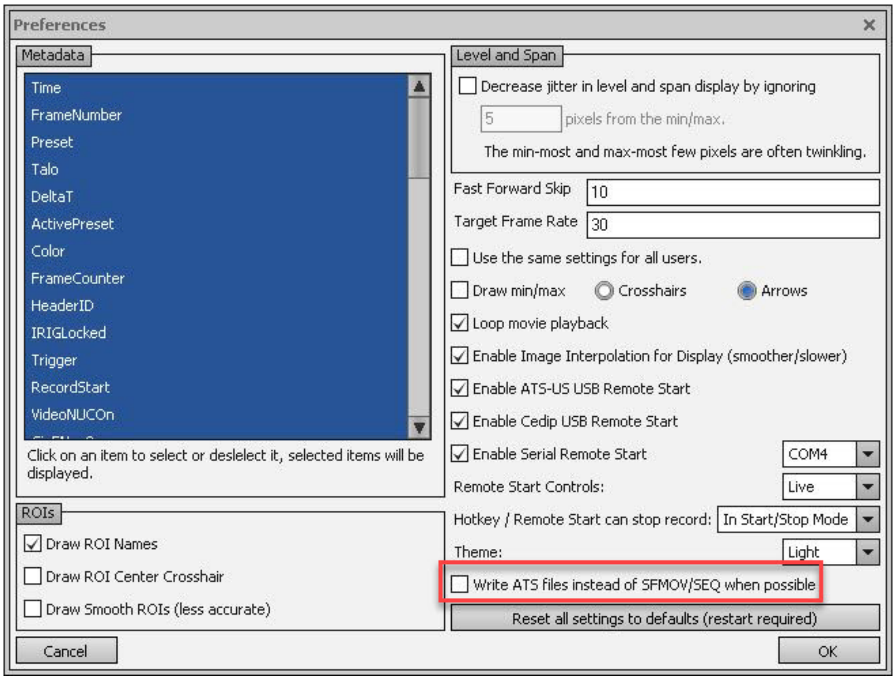


**Figure 1. Setting File Type preferences**

I changed the hotkey (using the Tools/Hotkeys menu item) that gets me to the Record Settings menu to the F1 key, so I can get there with a single keystroke.  I do this since I will be going back there again and again:

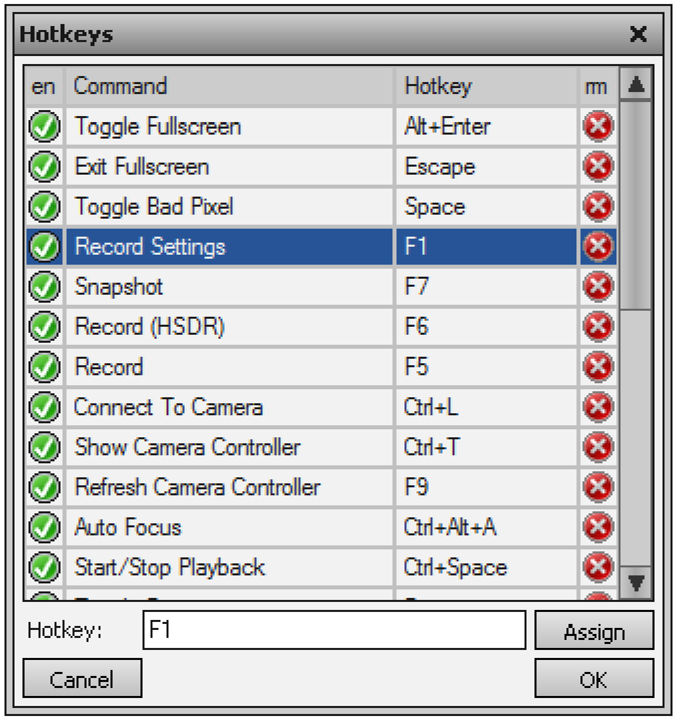


**Figure 2. Record Settings in Hotkey menu**

1) Now I turned off the factory calibration by selecting No Factory Calibration from the pulldown menu in Preset 0. I did not have to do this to create a user calibration, and in fact it had other consequences which I realized later. A better strategy might be to unload all the factory calibrations in the three presets that had them, thereby eliminating any trace of factory calibration from data that I collected later with the user calibration loaded. When the factory calibration is unloaded, the Setup window will show just the 0.095msec integration time:

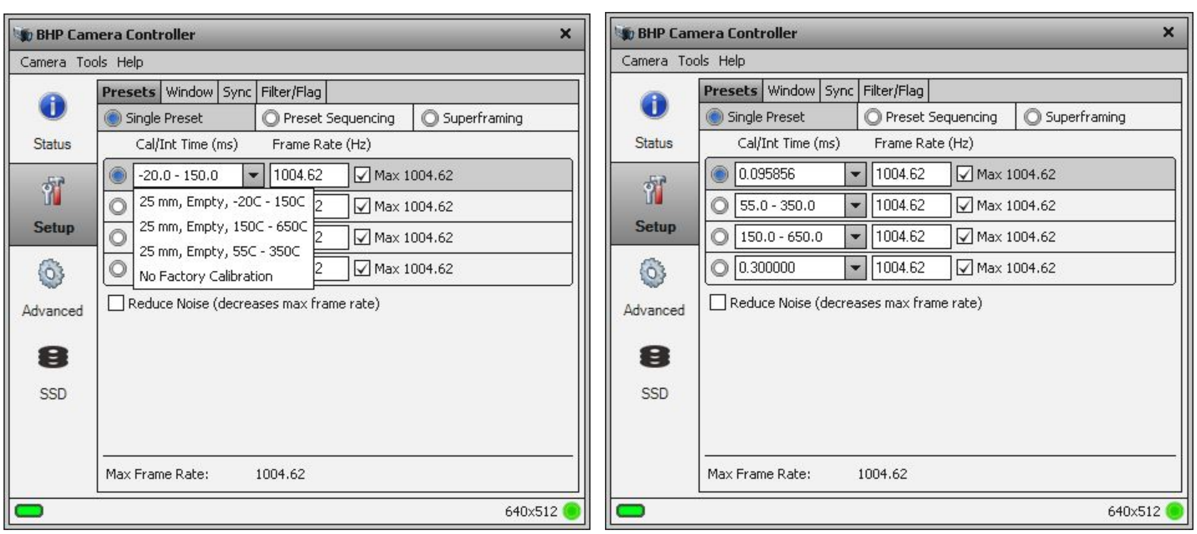


**Figure 3a-b. Unloading a factory calibration**

2) Unloading the factory calibration on the active preset immediately "breaks" the temperature and radiance display options, and the image goes black when either of those two options are selected. The temperature is reading at absolute zero, which is -273.15 °C, and the radiance is reading at the radiance offset coefficient C0, which is 1.1408e-2 W/sq. cm/sr:

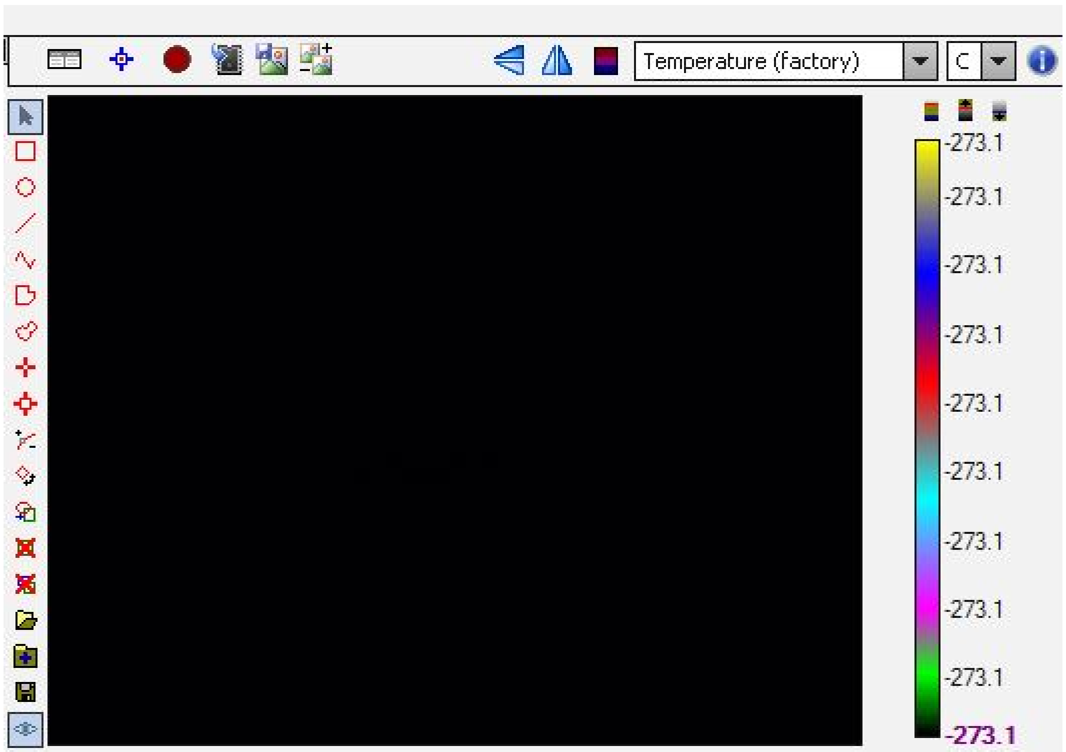


**Figure 4a. Black image at absolute zero when factory calibration is unloaded, and the units are set to Temperature**

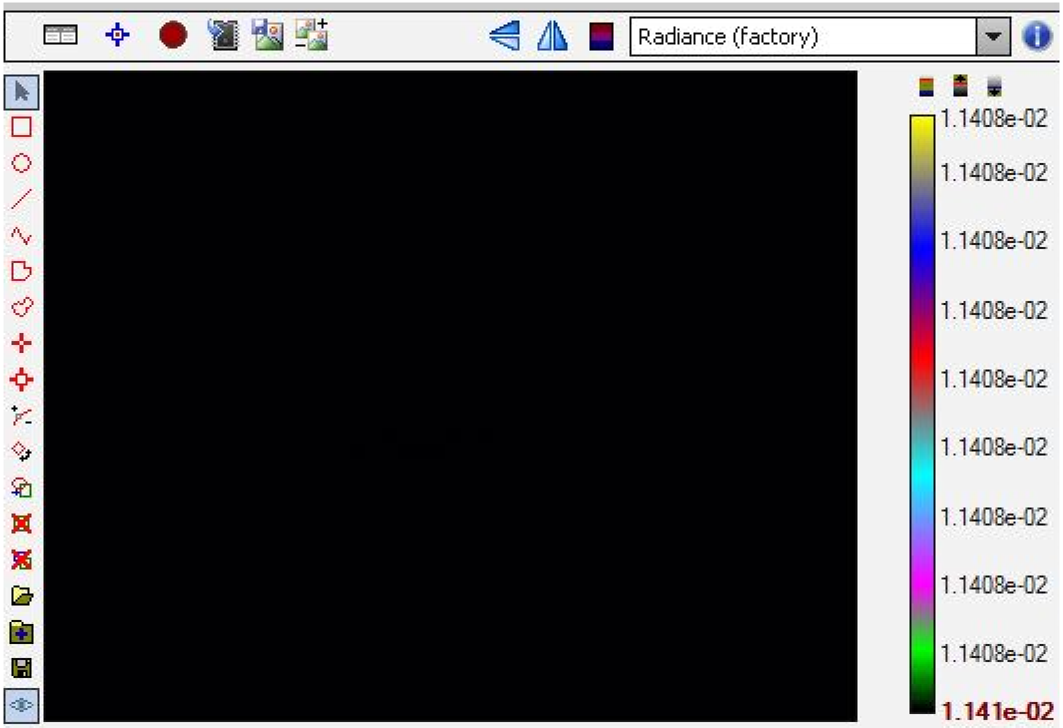


**Figure 4b. Black image at calibration radiance offset C0 when factory calibration is unloaded, and the units are set to Radiance**

3) But displaying images in Counts units still works, as it should, since it does not depend on a calibration, indeed, it is pre-calibration data:

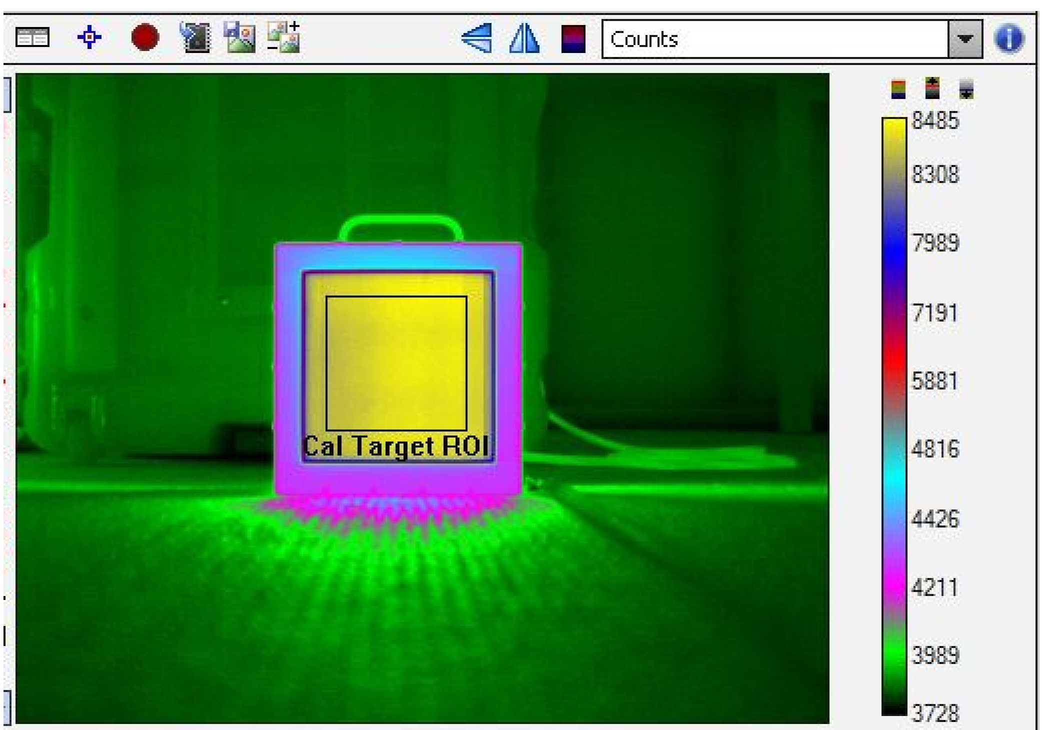

**Figure 4c. Image with factory calibration is unloaded, and the units are set to Counts**

4) At this point, the Object Parameters in Tools are mostly disabled by the act of turning off the Factory calibration, so I closed that part of the ResearchIR window by clicking on the little arrow marked in red in the figure below:

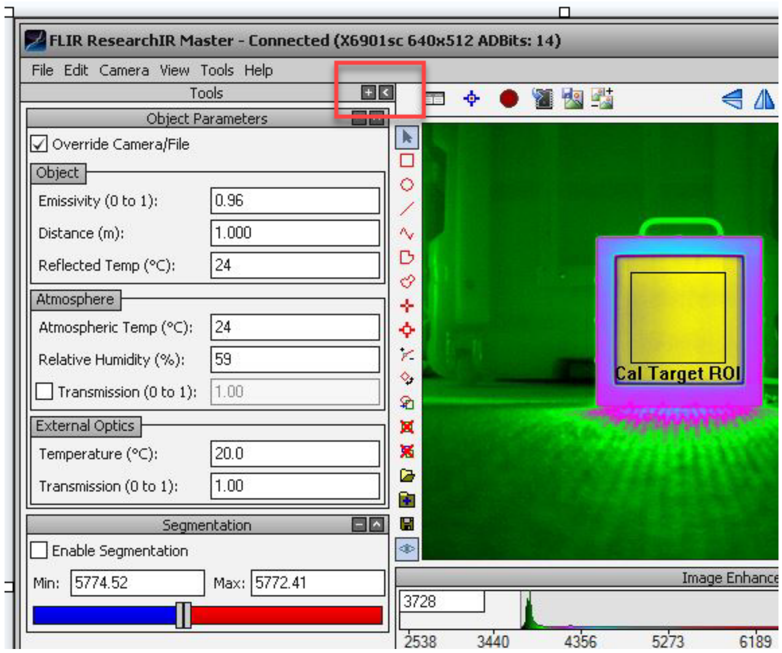


**Figure 5a**

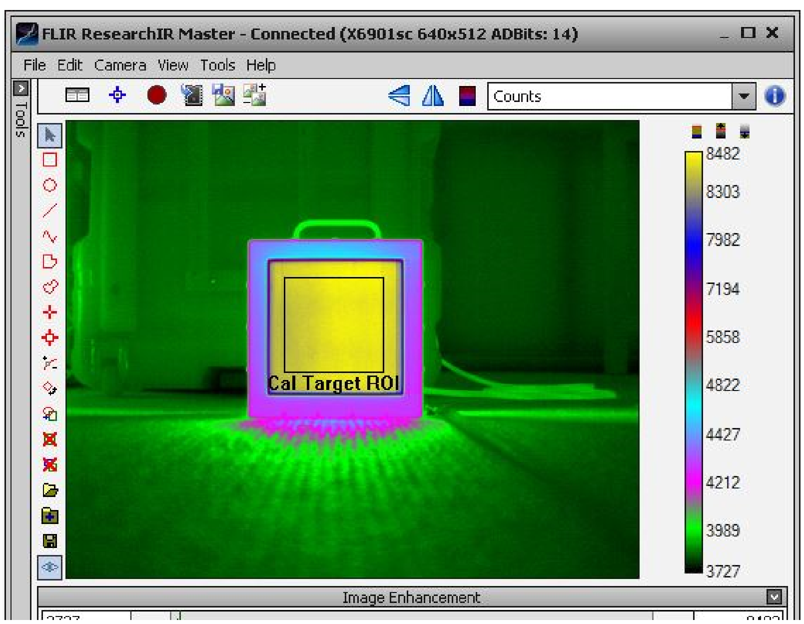


**Figure 5b**

5) I am going to record 16 frames of the blackbody at each calibration source temperature, so I set up the Record Settings menu appropriately. I choose 16 frames to record to memory, I choose a folder for those 16 frame movies to go into, and I create a file name template for the movie files. I set the Count to 1, and I leave this alone during the data acquisition. Using the counter adds a file number suffix to each new file, and it prevents one from accidently overwriting earlier data. Accidently overwriting a 16-frame movie made on the blackbody when it was at a different temperature than what it is now is very annoying and time wasting, since you have to set the blackbody controller back to an earlier setpoint and wait for it to stabilize again.

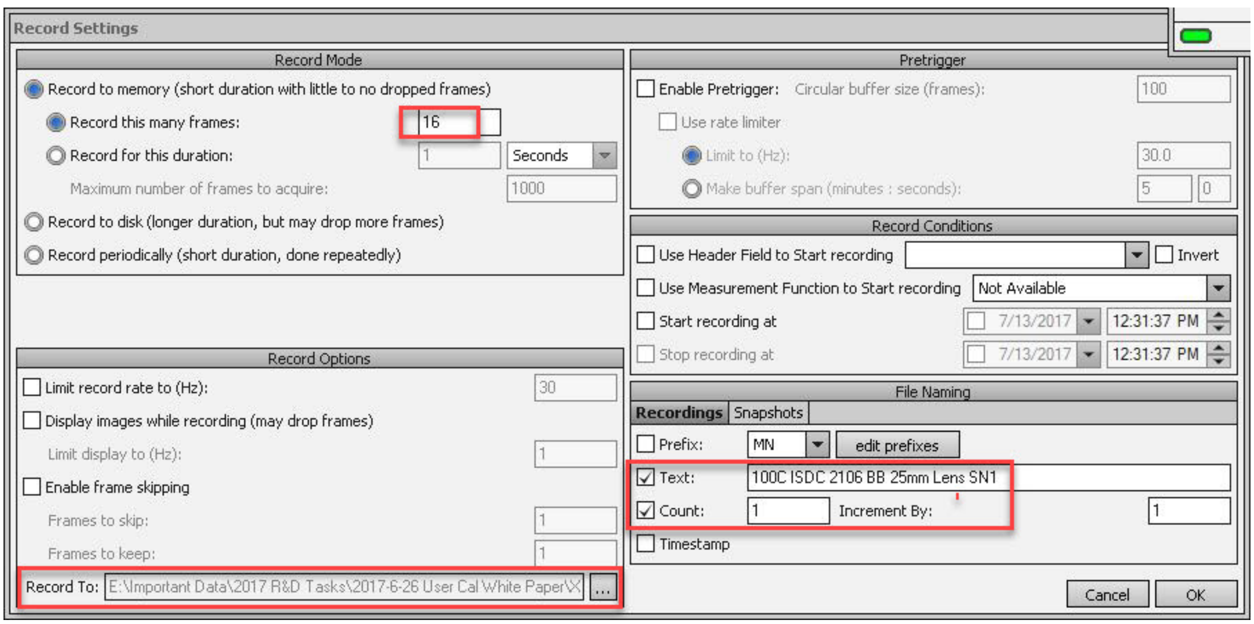


**Figure 6. Setting recording file path and filename text**

6) The blackbody was at 100 °C already, so that will be my first movie. Here are the three files I get when I record a movie, which I do with the Hotkey F5:

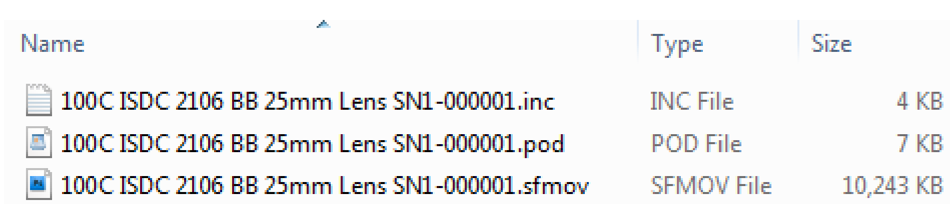


**Figure 7. Files recorded**

7) Now I change the blackbody to 55 °C. You can see the 3-by-3 matrix of thermoelectric coolers in the blackbody cooling down as the temperature ramps down:

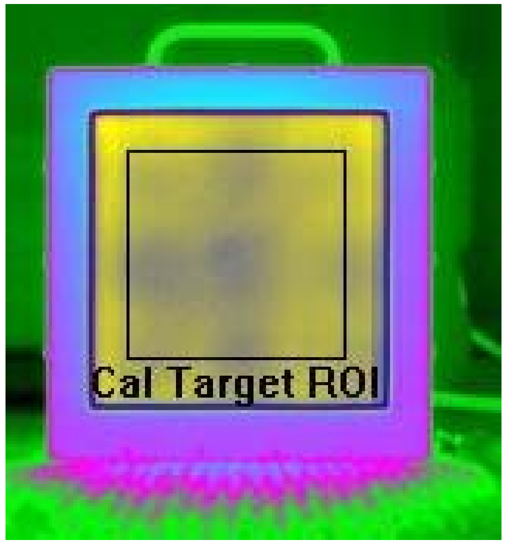


**Figure 8. Temperature nonuniformity during blackbody cooling.**

8) I waited until the blackbody controller indicated that the blackbody was stable, then I took a 16-frame data movie on the blackbody at 55 °C, with the settings shown below in Record Settings. The only change is the Text says 55 °C instead of 100 °C, and the counter for the file suffix has now incremented to 2.

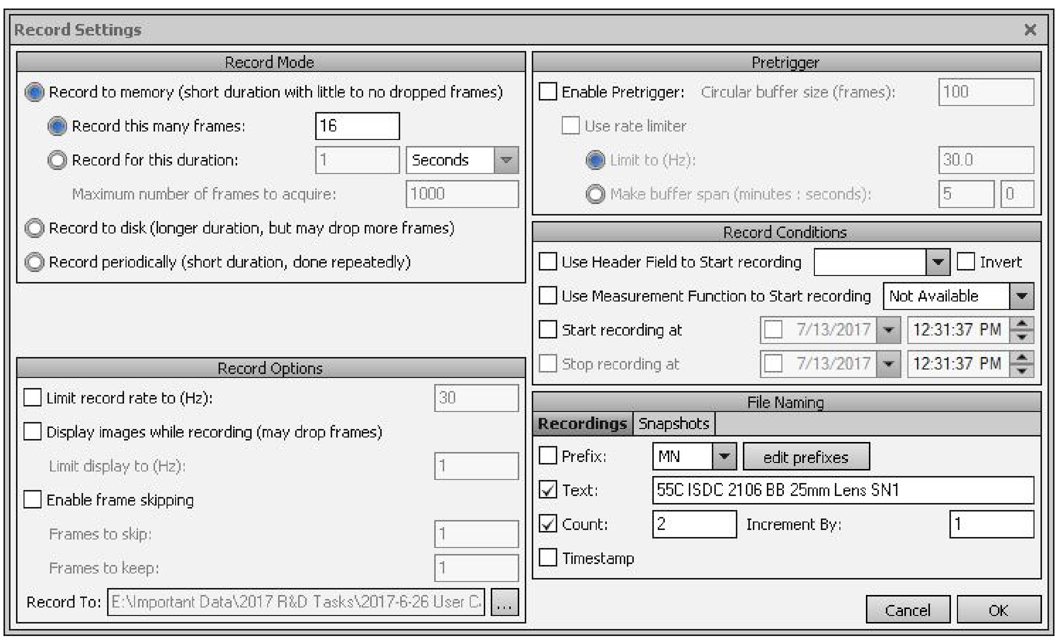


**Figure 9. Record settings after a movie has been acquired.**

9) Then I cranked up the temperature to 150 °C and let the blackbody controller tell me when it was stable. You can see the 3 by 3 matrix of thermoelectric coolers heating up in the blackbody during the ramp-up of temperature:

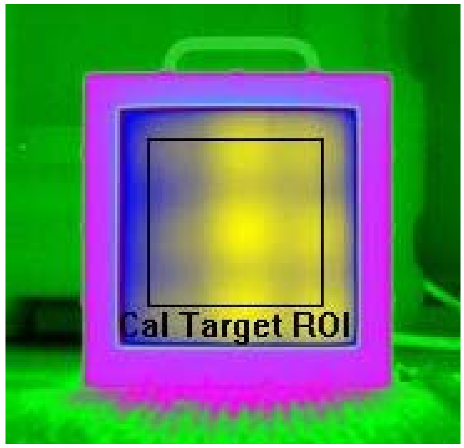


**Figure 10. Temperature nonuniformity during blackbody heating.**

10) This 150 °C temperature is pretty close to where the camera FPA will go into saturation, so it is a good idea to turn off the NUC, and then monitor the counts in the ROI on the blackbody to see what the maximum count value is. The maximum value in the ROI is 14,250 counts, so this is not too hot for the camera with these settings, but I won't push it higher.

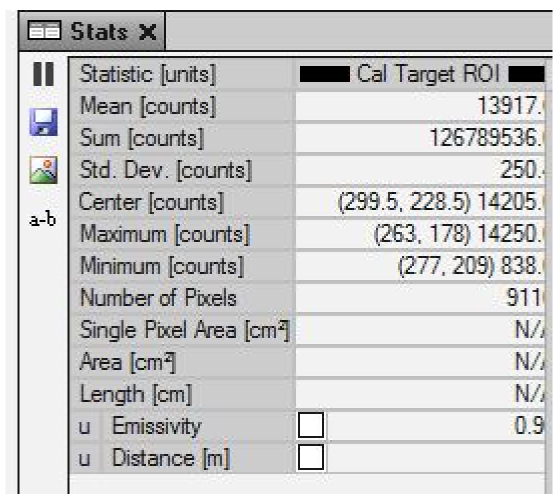


| Statistic [units] | | Cal Target ROI |
|---|---|---|
| Mean [counts] | | 13917. |
| Sum [counts] | | 126789536. |
| Std. Dev. [counts] | | 250. |
| Center [counts] | | (299.5, 228.5) 14205. |
| Maximum [counts] | | (263, 178) 14250. |
| Minimum [counts] | | (277, 209) 838. |
| Number of Pixels | | 911 |
| Single Pixel Area [cm²] | | N/ |
| Area [cm²] | | N/ |
| Length [cm] | | N/ |
| u Emissivity | ☐ | 0.9 |
| u Distance [m] | ☐ | |

**Figure 11. Statistics on ROI during calibration.**

11) Finally, I took a 5 °C data movie. I now have four points which is enough for a user calibration. I don't leave the blackbody at 5 °C, because that is below the dew point and it will get covered with condensation pretty fast.

12) It is now time to create the two calibration files in the CalibratIR utility. This can be found in here:

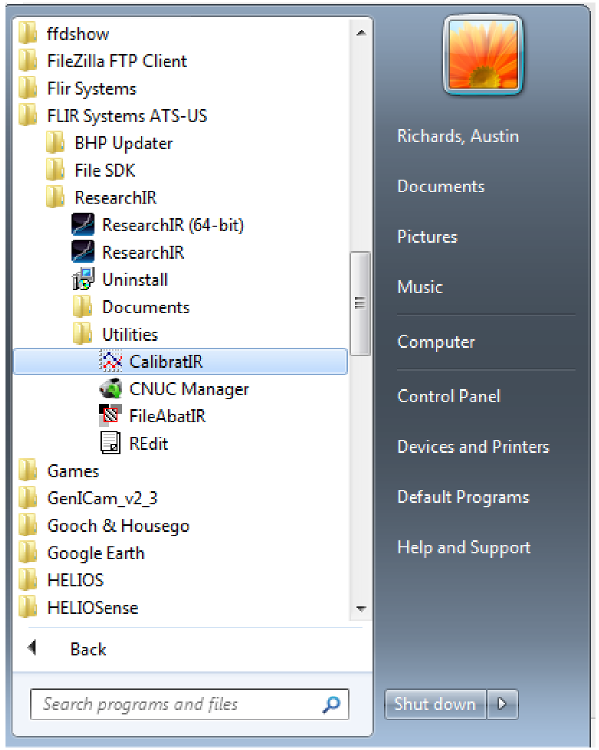


**Figure 12. Start menu with CalibratIR**

13) First, I check again to make sure there is no user calibration loaded into ResearchIR. The easiest way to do that is to look to see if I can edit an existing calibration. Edit is greyed out, so there is no calibration in ResearchIR:

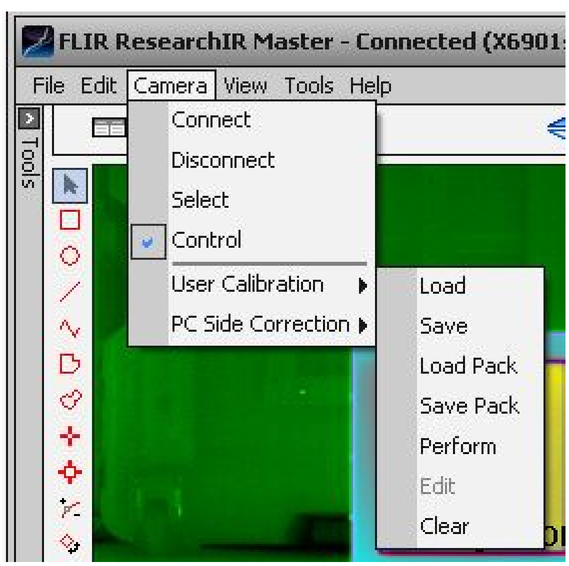


**Figure 13. Checking that no user calibration is loaded.**

14) Rather than perform a user calibration in the live mode, I am going to start up CalibratIR and use it in concert with ResearchIR, which will be playing back the 16-frame movies I

already took. Doing the user calibrations in post processing mode is a much better way to go, because if you try to make the user calibration in live mode, and you have to wait many minutes for the blackbody to settle at the various temperatures, you run the risk of something going wrong (power failure or ResearchIR crash) and losing the data in the Calibration Wizard, which might take hours to collect!

I choose the expert mode in CalibratIR:

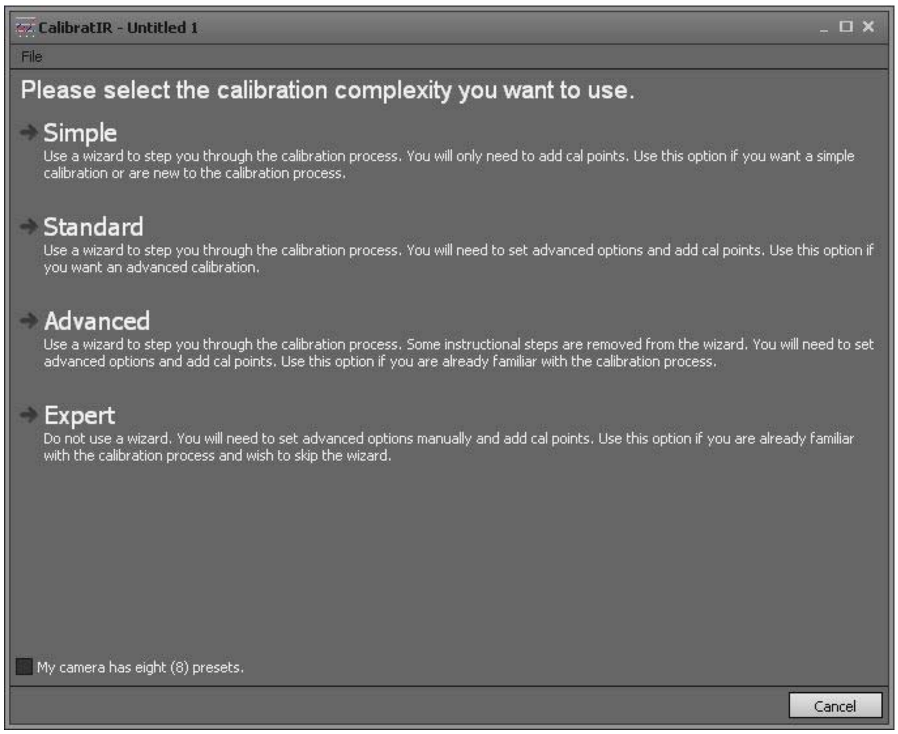


**Figure 14.**

15) The first screen you see in Expert mode is the response screen with a Tophat response and 3 to 5 microns spectral band limits. These are the default values intended for a midwave InSb camera, and they won't work well at all for the SLS camera used for this user calibration, in fact you will get a poor linear fit to the radiance versus counts curve. For the puposes of this calibration, I will use the Tophat spectral response for the camera with limits set for the Longwave IR SLS detectors at 7.5 microns and 10.5 microns. Later on, I will reprocess the calibration with a measured spectral response file (*.prn file) to get more accurate in-band radiance measurement results.

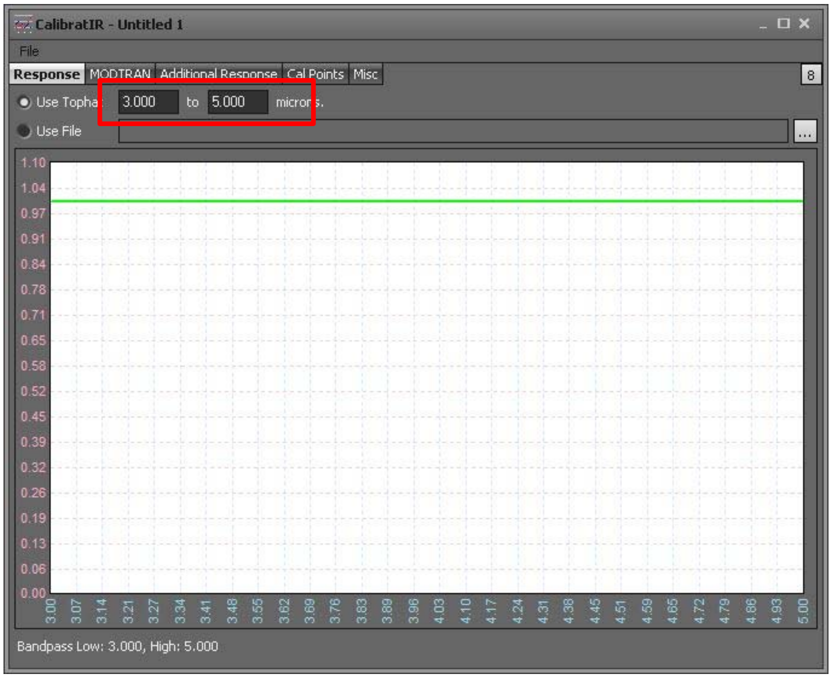


**Figure 15.  Default 3-5μm spectral limits. Don't use these for an LWIR SLS camera.**

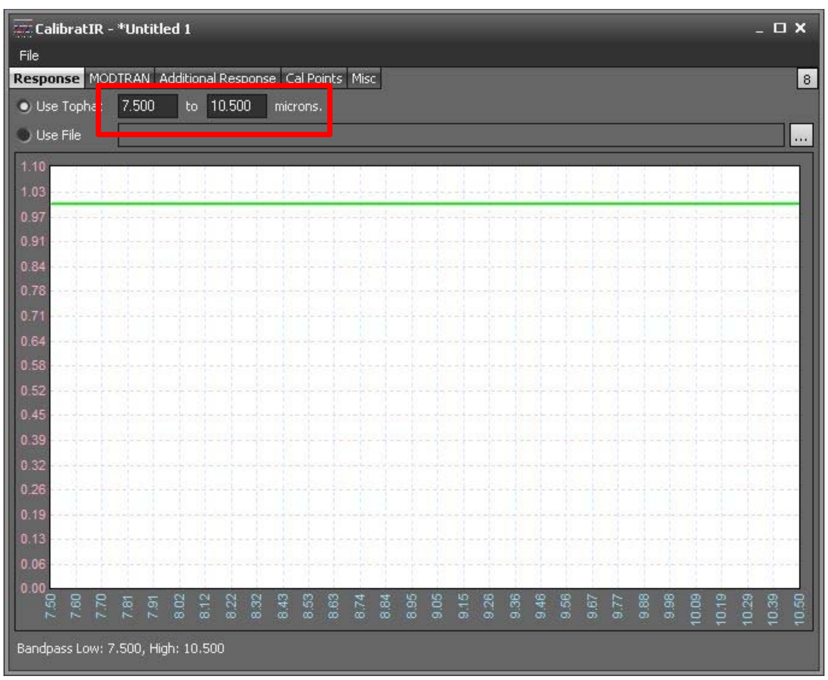


**Figure 16. Selected 7.5-10.5μm spectral limits for LWIR SLS camera**

16) The next screen is MODTRAN.  This is a way for the camera user to account for the air path transmission losses between the camera and the calibration source. It is only applicable to calibration – it does NOT correct for atmospheric transmission losses when

the camera is subsequently used to collect data in the field!  I will leave the tab at the default setting of 100% transmission over the band of the camera, since the camera and the blackbody were only 1 meter apart, which is an essentially perfectly transmitting air path in the LWIR band.

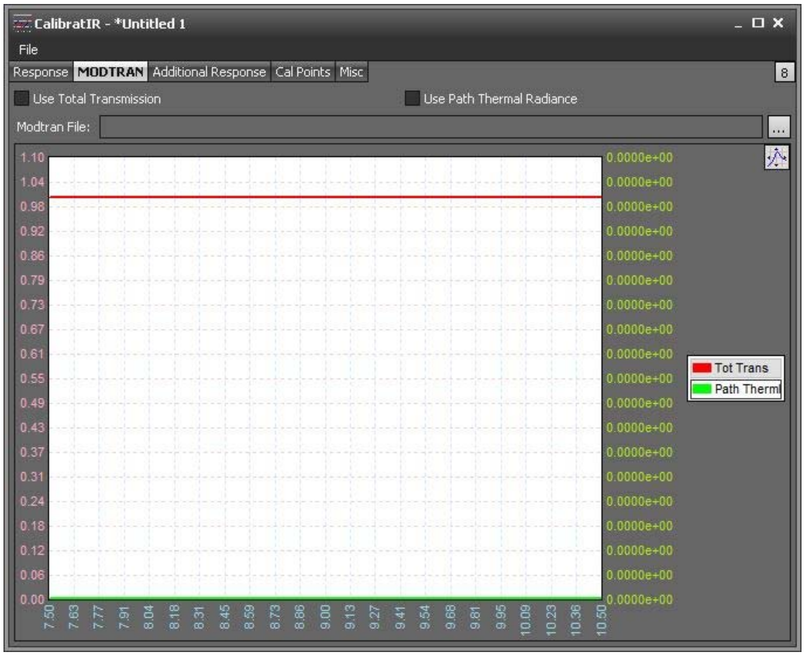


**Figure 17.  MODTRAN tab in CalibratIR.**

17) The next screen is called Additional Responses. I am going to leave that at the defaults as well, which is no added responses. This screen allows the user to account for other optics that have some spectral shape to them, like an enclosure window or a fold mirror.

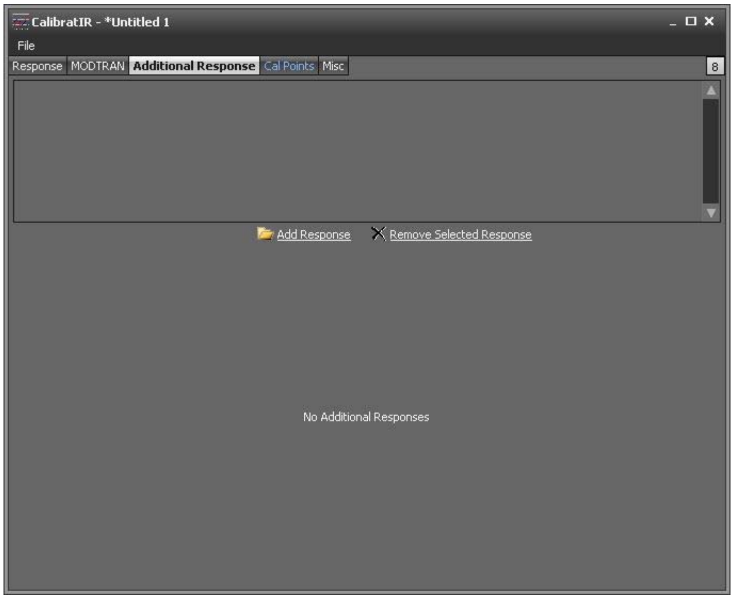

**Figure 18. Additional Response tab in CalibratIR**

18) Finally, we are at the Cal Point screen. This is where I will populate the measurement grid with the four data points for counts at the four different blackbody temperatures: 5 °C, 55 °C, 100 °C and 150 °C. The first thing I do is click on Add Calibration Point:

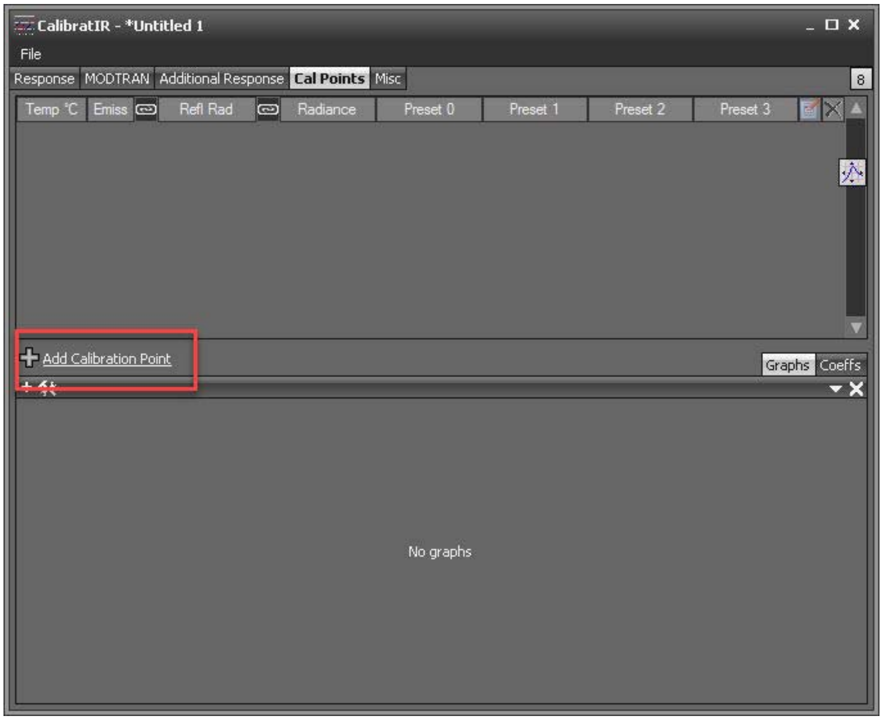


**Figure 19. Adding a point in the Cal Points tab in CalibratIR**

I get this view with a 0.98 default emissivity value, which I will change to 0.96:

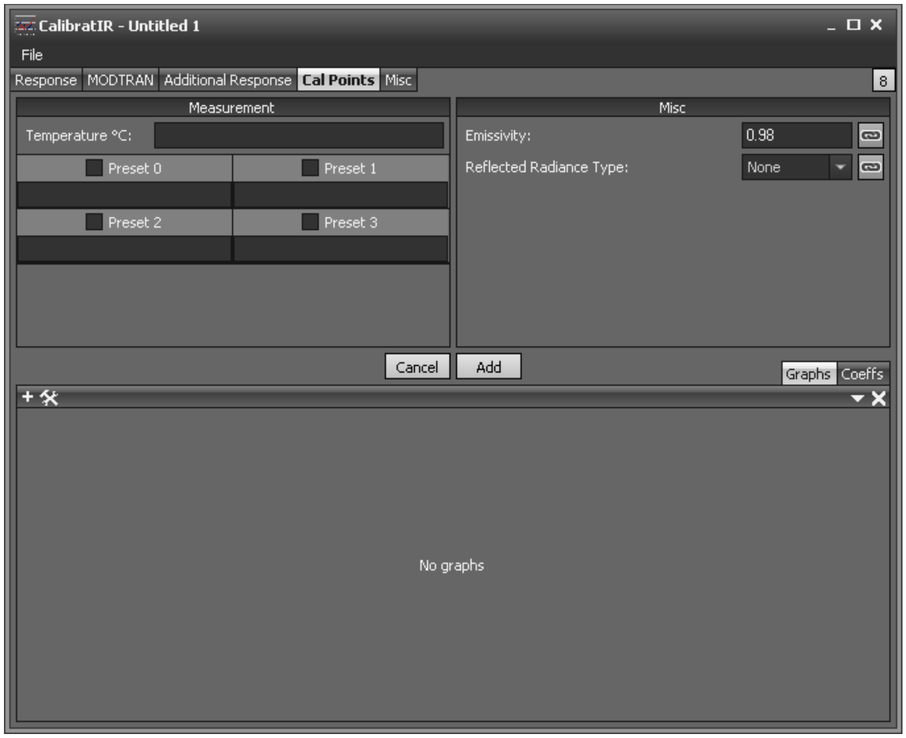


**Figure 20. Boxes to enter temperature, emissivity and count values. Emissivity default is 0.98.**

19) I will now add the 55 °C blackbody temperature data. I will open up the 55 °C movie called “55C ISDC 2106 BB 25mm Lens SN1-000002.sfmov” in ResearchIR, run a 16-frame average on it using a Tools filter called Frame Average, and take the mean of the ROI on the blackbody source. Note that if I open an SFMOV that was made when there was no calibration loaded in the active preset, but there were factory calibrations loaded in other presets, then ResearchIR will automatically default to the temperature (factory) units display view. Since there is not any calibration on the preset, the screen will look black, as shown in Figure 21. I then have to select counts as the display units, and then I will get a proper counts image.

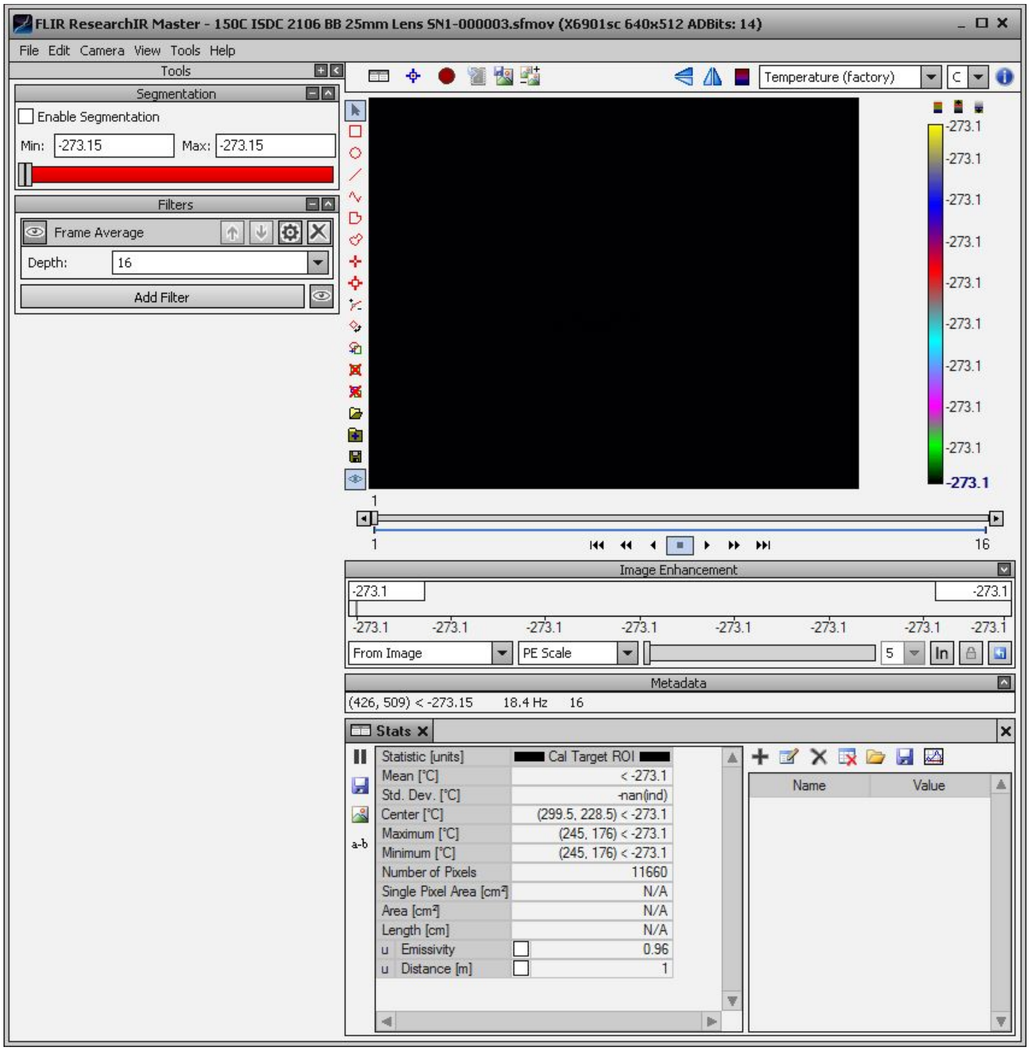


**Figure 21. Frame Average filter with 16-frame depth active. The display units need to be set to counts.**

The Filters are in in the Tools menu as shown below. I let the movie play through once and I get a 16-frame average value in the ROI called Cal Target ROI. I also save this ROI so I can open it later and apply it to the data I have taken. This is yet another advantage of the post-processing calibration: it is like the camera is live and you can change settings and move around the ROI if you need to. I get a value of 5328.8 counts, which I will enter into CalibratIR.

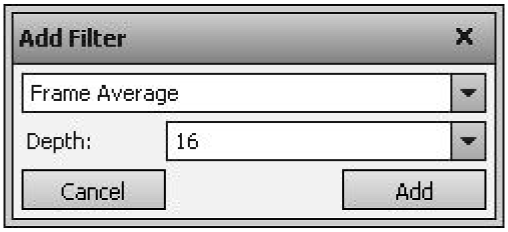


**Figure 22. Add frame averaging filter with 16-frame depth**

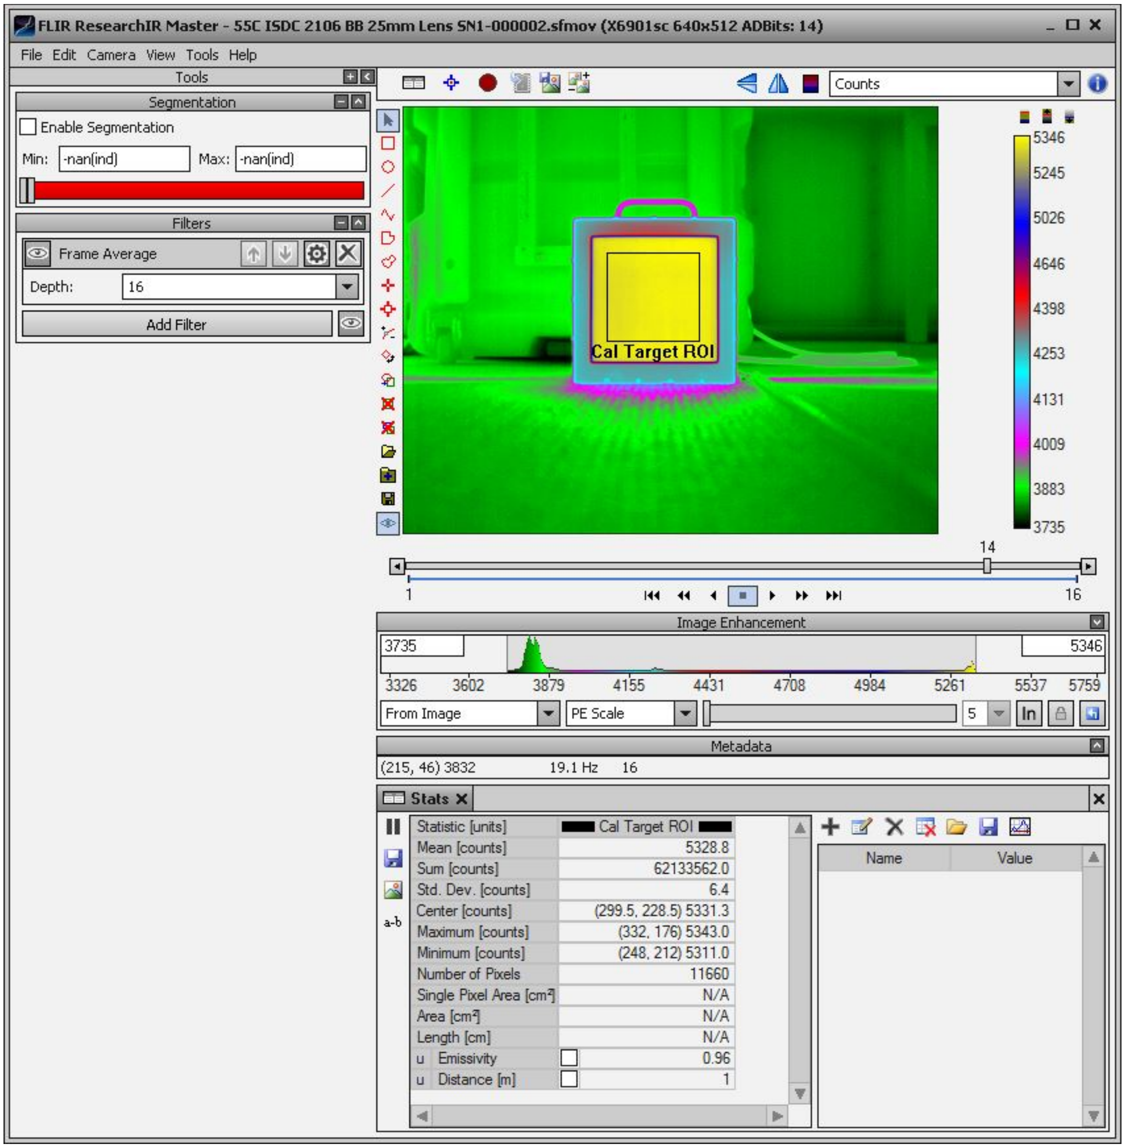


**Figure 23. Run movie through once to make the ROI mean value settle.**

20) Figure 24 shows the entry process into the measurement grid of CalibratIR. First, I change the emissivity to 0.96, which is the manufacturer specified value for the ISDC Model 2106 blackbody emitter surface, then I check the Preset 0 checkbox which is unchecked by default, I type in the temperature and mean counts, then finally I hit the Add button in the middle of the screen. The little chain icon next to the emissivity means that the next calibration point I put in will use the same emissivity value, instead of the 0.98 default. You can uncheck that and enter the emissivity for each point individually, in case you are using various different blackbodies for the user calibration.

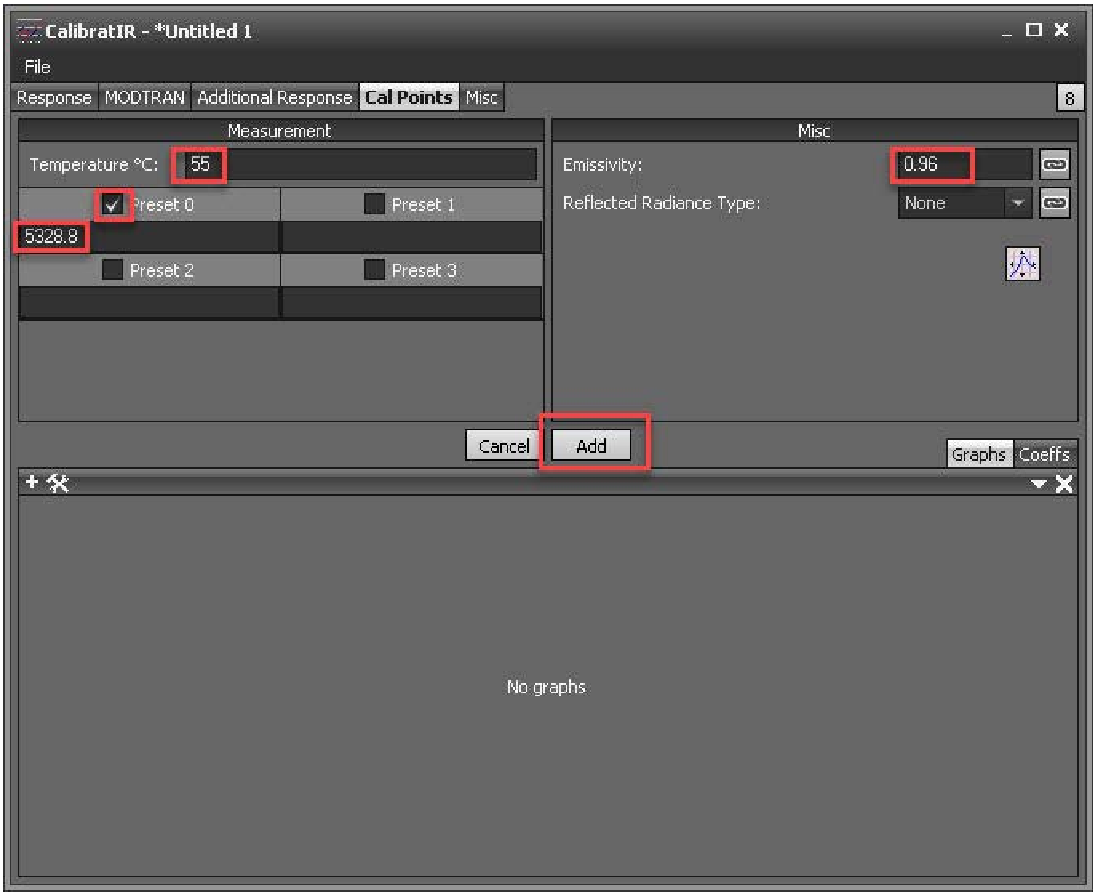


**Figure 24. Entering a 55 °C calibration point into CalibratIR**

21) I now see this screen which is called the Measurement Grid. The software is computing the in-band radiance to be 4.3867e-3 watts/sq. cm/sr. There is an option to change the reflected radiance, marked in red in Figure 25. This reflected radiance control will add in the small amount of reflected energy off the surface of the blackbody from the room I worked in. It is only a 4% reflection of a 24 °C room, so it is not a big correction, but a good one to do, especially for LWIR cameras where the thermal derivative is smaller than in the MWIR band and so room temperature radiance is not all that small compared to emitted radiance levels in the 30-60 °C range.

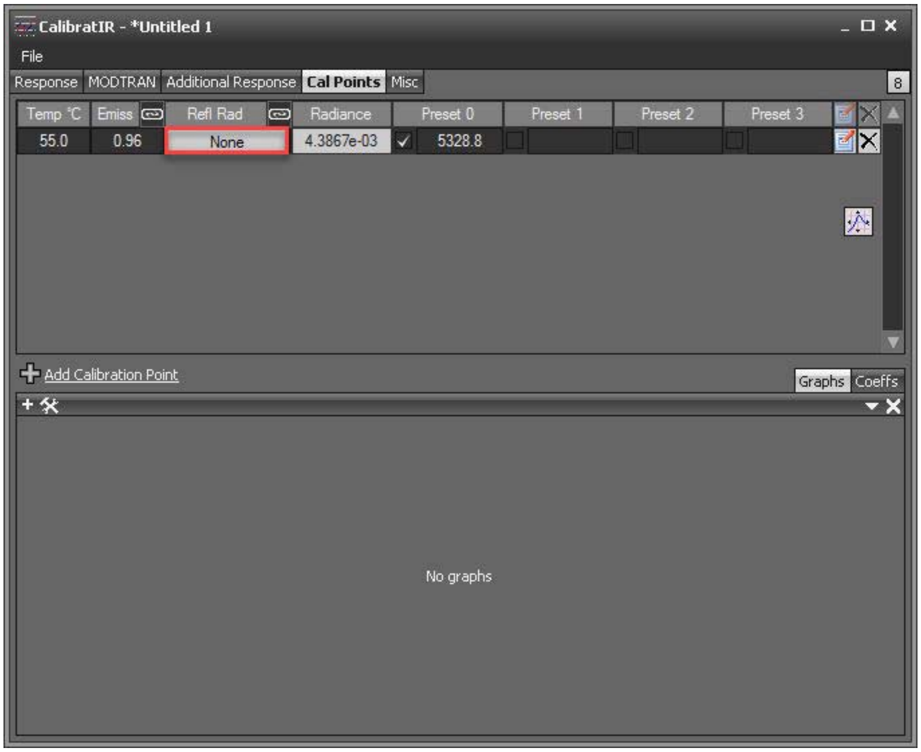


**Figure 25. Measurement grid with button to change Reflected Radiance**

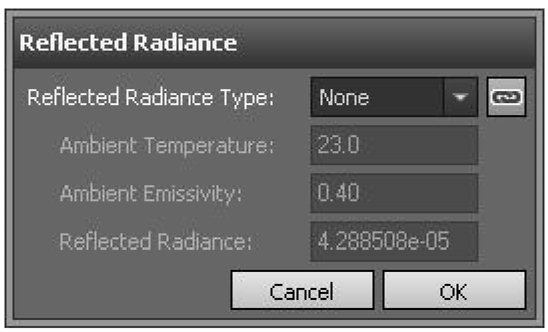


**Figure 26. Reflected Radiance dialog**

The above screen is the default settings. I will use Calculated radiance with an ambient temperature of 24 and an ambient emissivity of 1. The radiance is about 40X smaller than the 55 °C emitted radiance, so it is a ~3% correction, roughly.

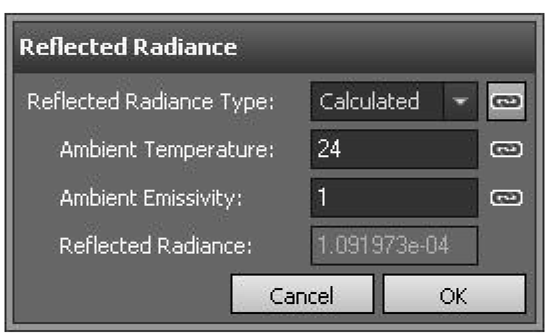

**Figure 27. Changing reflected radiance to a calculated value**

Now the reflected radiance is applied. This means that the software is accounting for both the emitted and the reflected radiance contributions to the total radiance detected by the camera.

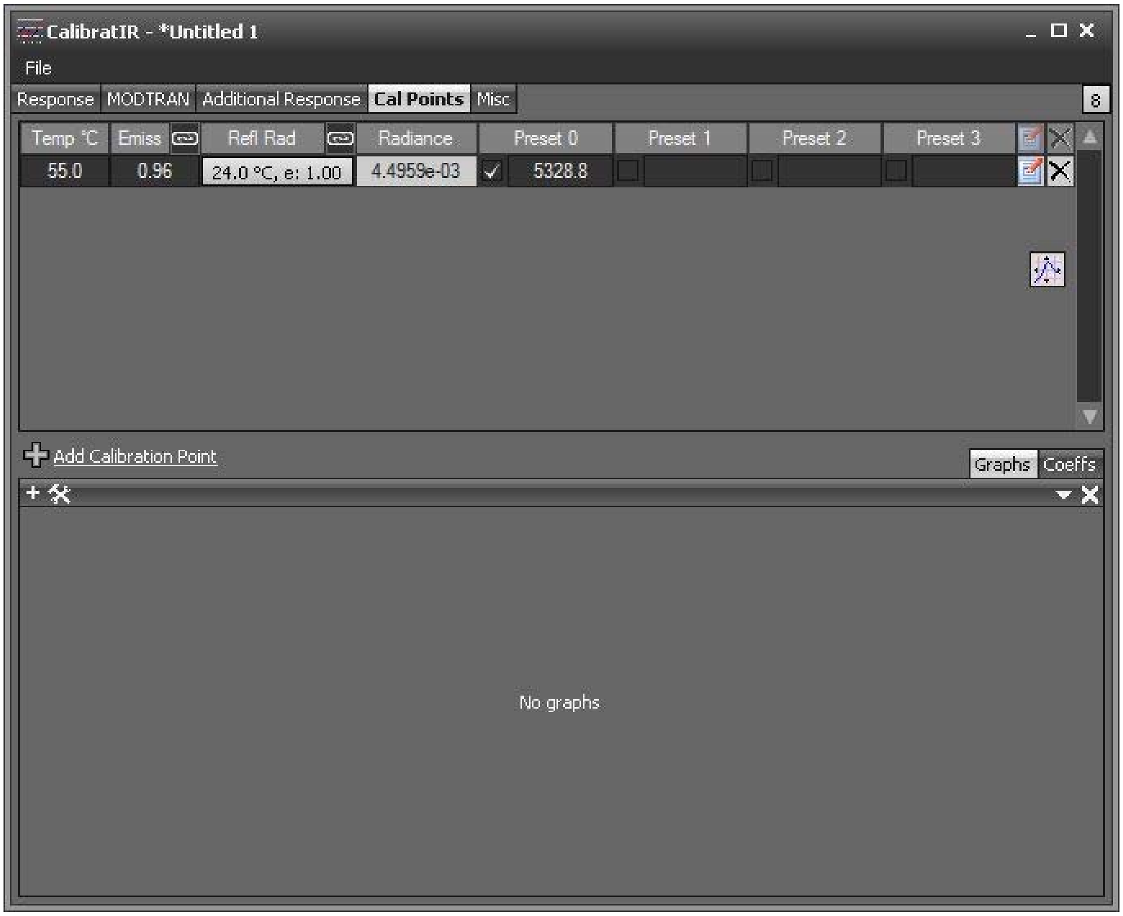


**Figure 28. New reflected radiance value**

22) I now repeat the steps of opening the other three data files, computing the 16-frame average and entering the ROI mean into the measurement grid one point at a time, until I have this set of points and a graph with a curve fit. The curve fit is linear by default and should be left linear always, even though there are provisions for higher order polynomial fits.

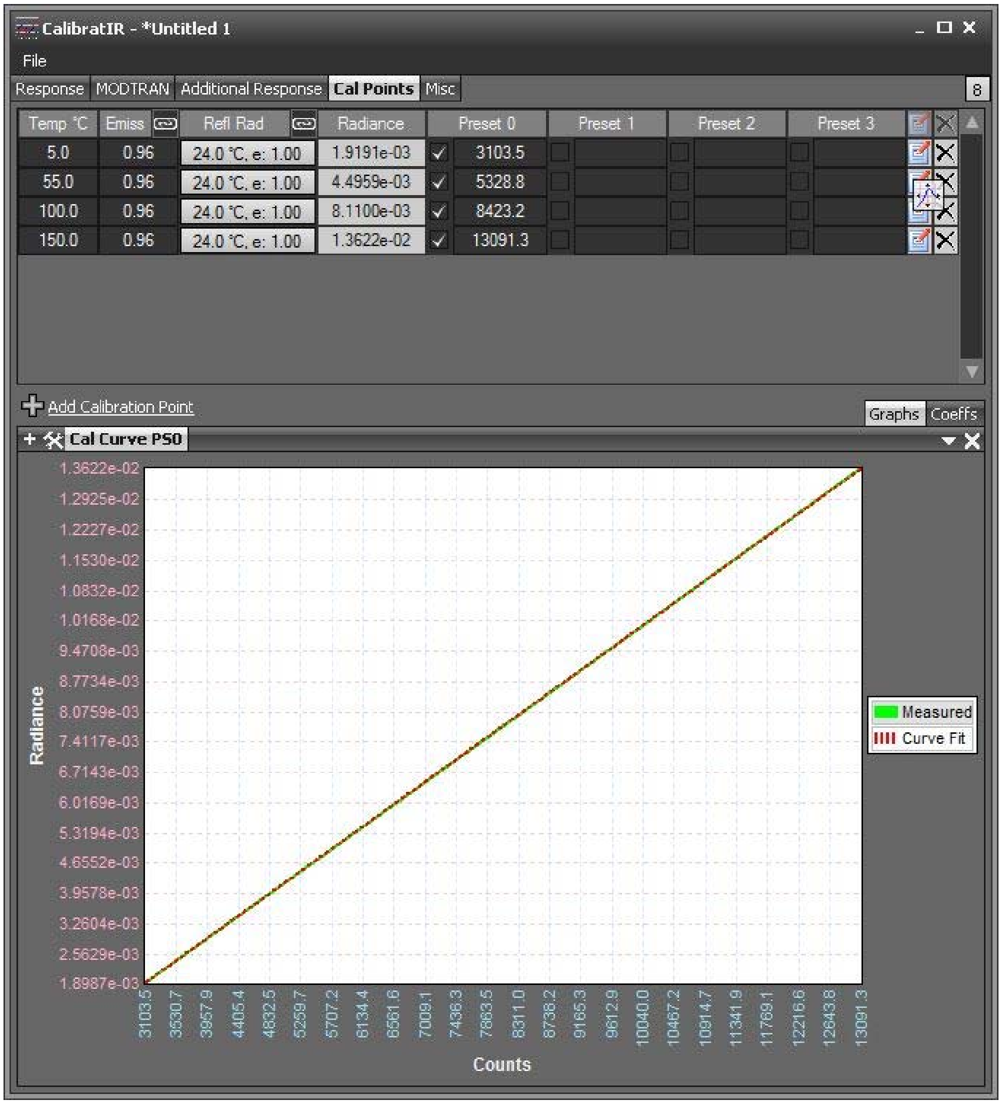


**Figure 29. Measurement grid with four values and radiometric transfer function**

23) We check the radiometric transfer function coefficients by clicking on the Coeffs tab and then the Preset 0 tab:

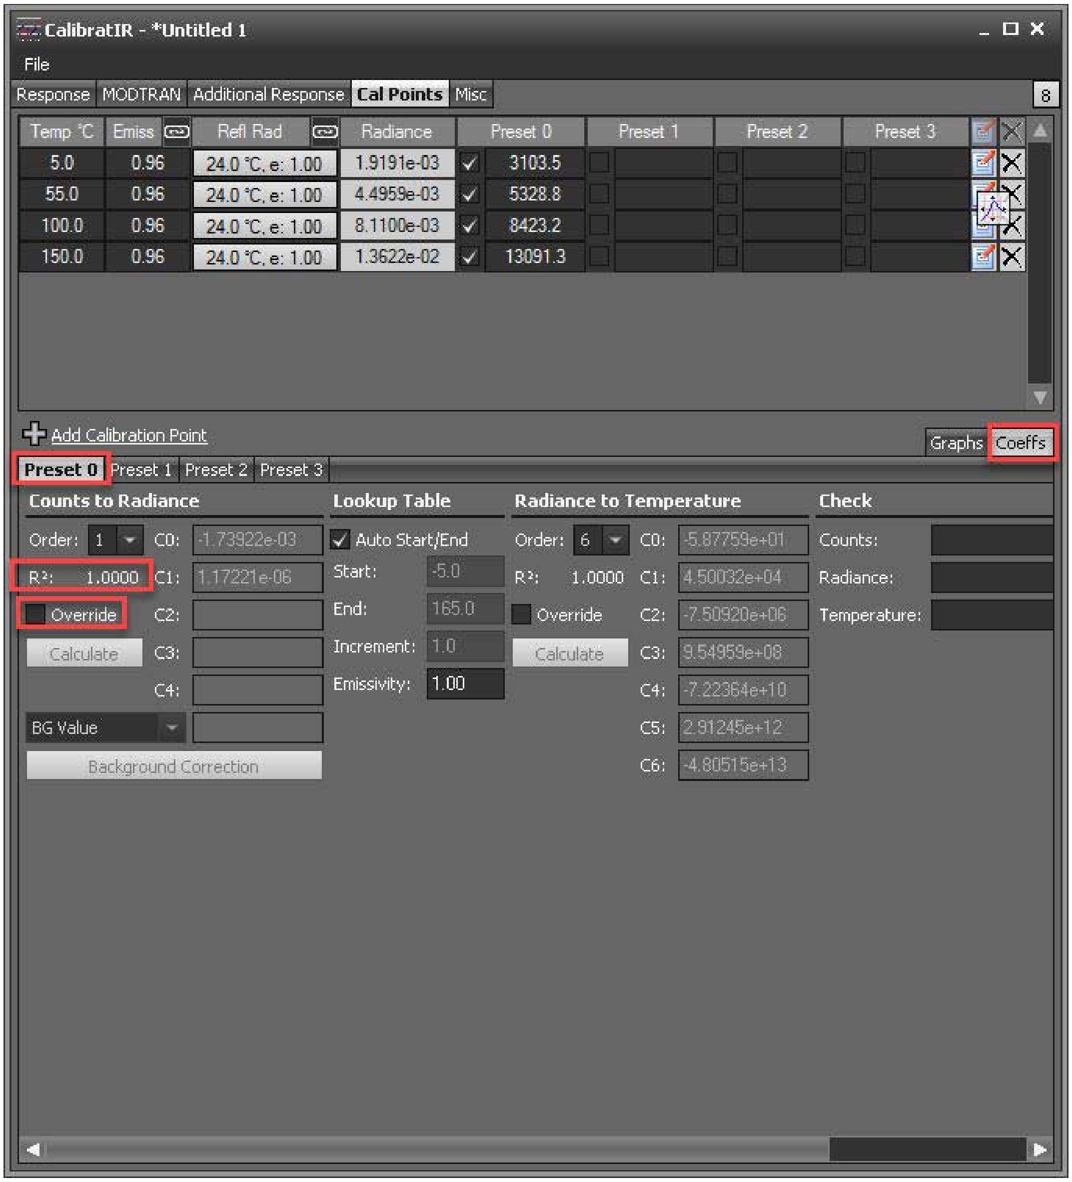


**Figure 30. Coefficients tab in Cal Points window with important controls marked in red**

The linear (order 1) polynomial fit is excellent with an $R^2$ value of 1.0000. The C0 and C1 coefficients are shown slightly greyed out, which means they are being calculated. If I check the Override checkbox, I could adjust their values manually, but I won't. The Auto Start/End is checked by default, and CalibratIR thinks you can get a good calibration from -5 °C to 165 °C. You can leave that alone, or you can be more conservative and uncheck the box and manually enter values like 5 °C and 150 °C, just to be true to the range of temperatures of the calibration sources you used in the calibration.

24) Lastly, I will go to the Misc tab where I will set the IFOV tags for the camera and lens. The values are 1000 microradians for the horizontal and vertical angular size of the 25µm pixels and the 25mm lens. I can also add a slant range of 1 meter, just because that was

the range of the blackbody during the acquisition of the movies. When I look at data that I collect on the blackbody later, the Stats distance parameter will be 1 meter by default, which is convenient.

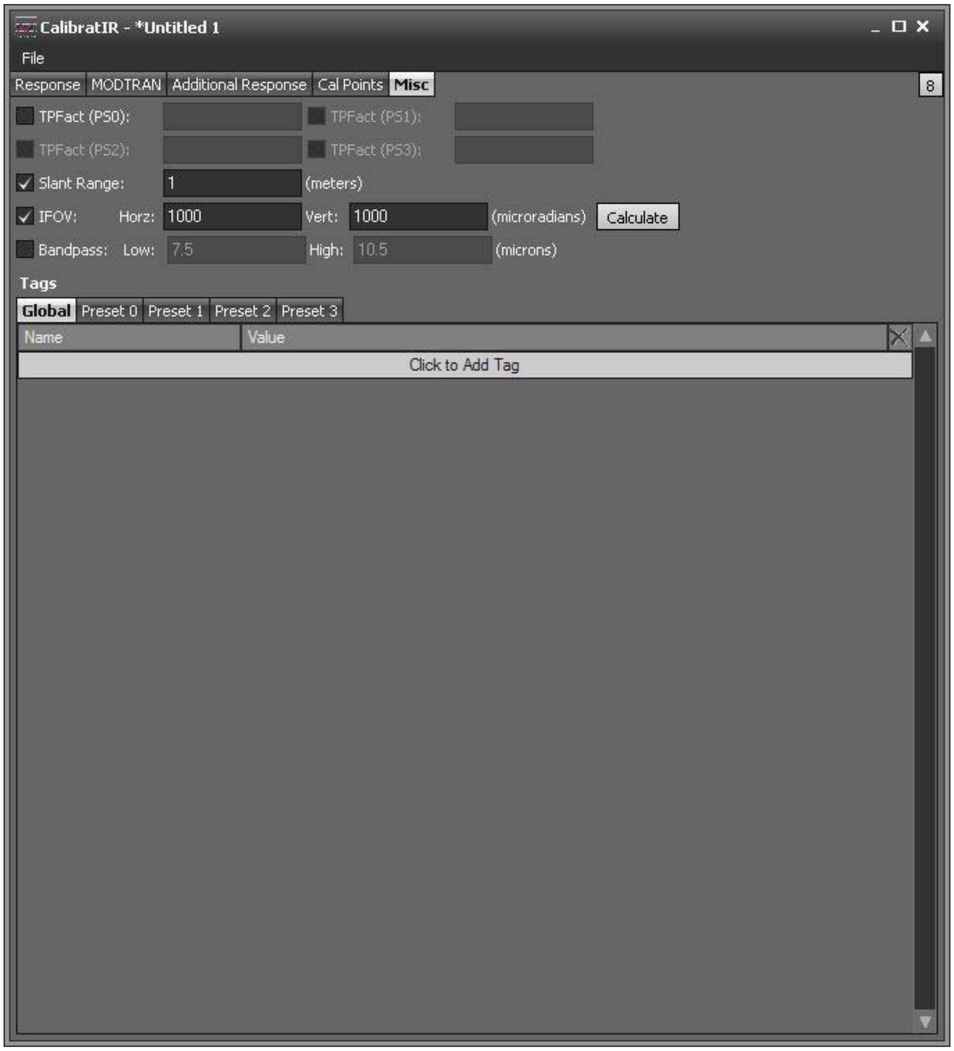


**Figure 31. Miscellaneous tab with extra tags that can be edited by the user**

25) I can now save the calibration file. Notice how there is no name for the file yet along the top of the screen. The name I gave the calibration files (*.cal and *.inc) contains enough info to unambiguously describe what I had set up in terms of the camera and the optics for the cal. This is the filename I chose:

**"-20C to 150C 95pt8559usec 25mm SN1 Lens Calibration"**

I need to generate an include file at the same time, so I need to use the Save As option and save as an include file. I will automatically get a *.cal file at the same time with the same name. That does NOT happen if I save the *.cal file only by using the Save option.

In that case, I have to go back and do a Save As and choose include file.  Just a little foible!

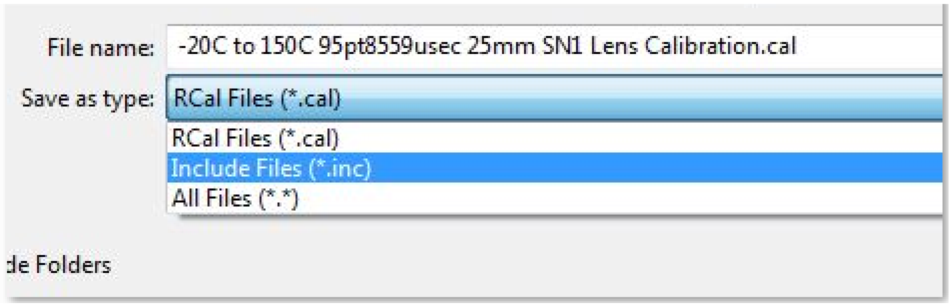


**Figure 32.  Two types of files that can be saved out of CalibratIR**

26) Finally, I can load the calibration into ResearchIR using Camera/User Calibration/Load:

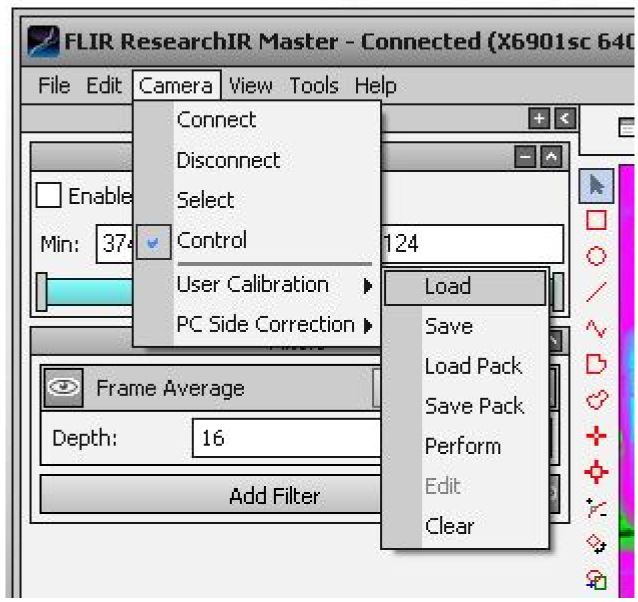


**Figure 33.  Loading the new calibration into ResearchIR**

You can load either the *.cal or the *.inc file in ResearchIR and the other file comes along automatically:

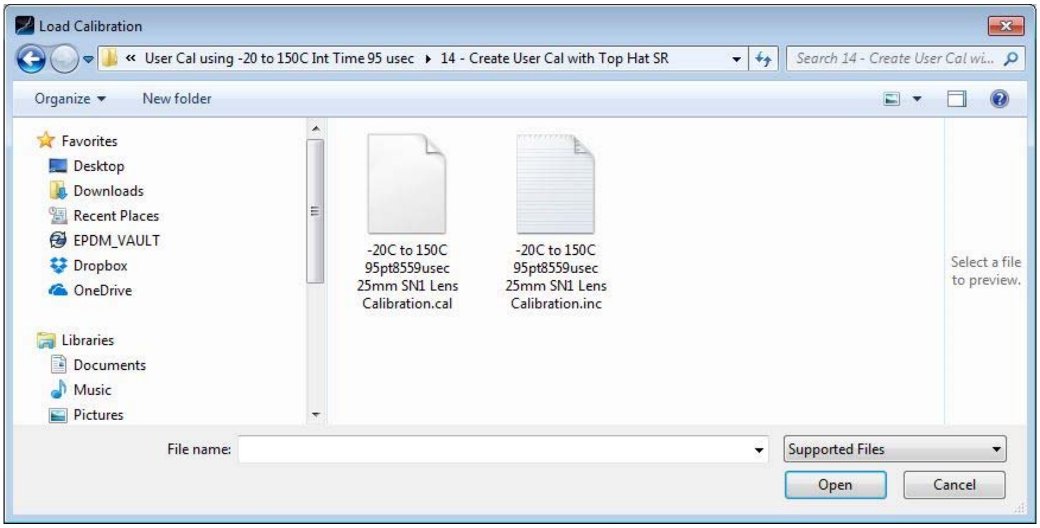


**Figure 34. Load Calibration window**

27) It is always a really good idea to examine the user calibration using the Calibration Wizard screen after loading it in to ResearchIR with a live camera connected. There is no indication on the main screen of ResearchIR that a user calibration is loaded and active unless you either scroll though the display units options and see Temperature (user) or Radiance (user), or you look in this pull-down menu and see that Edit is an option that is not greyed out. Edit it and it will open in a screen that looks like CalibratIR:

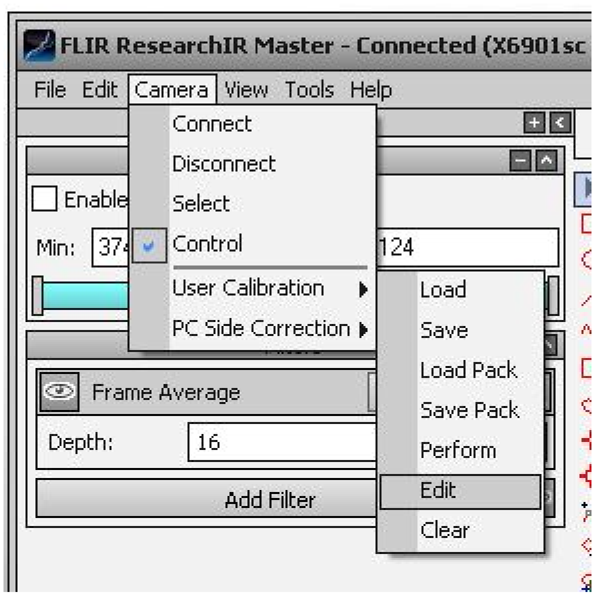


**Figure 35. Using the Camera/User Calibration/Edit function to check the loaded user calibration within ResearchIR**

I take a look at the Cal Points tab and I see the four calibration points in the measurement grid, so it looks correct, and the linear fit is very good looking.

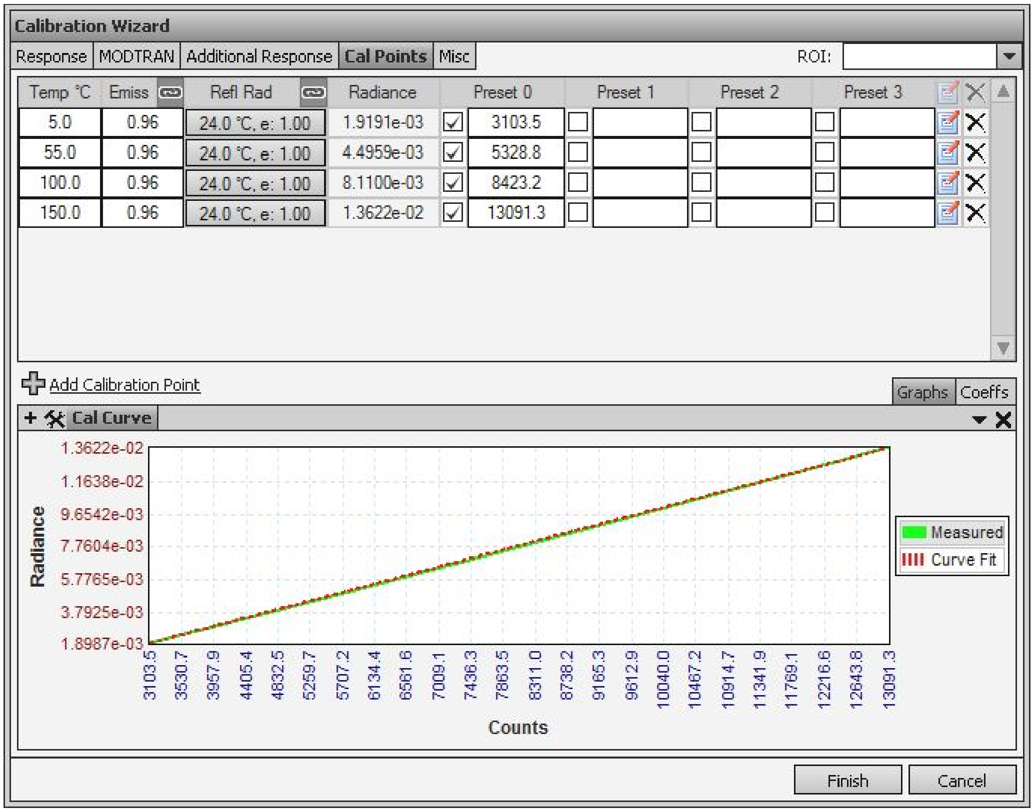


**Figure 36. Displaying the loaded calibration graph**

28) Here is another little foible of ResearchIR that you should be aware of. If you look at the Spatial Calibration window, it will not show you the two IFOV tags, though they ARE loaded and active if they are in the include file. You get to the Spatial Calibration window by going to Edit/Spatial Calibration in the top menu ribbon in ResearchIR:

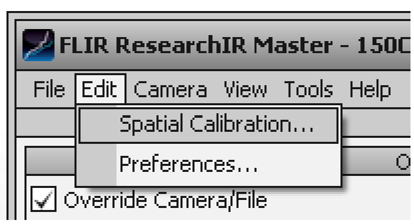


**Figure 37. Edit/Spatial Calibration command**

The two tags at the bottom are blank:

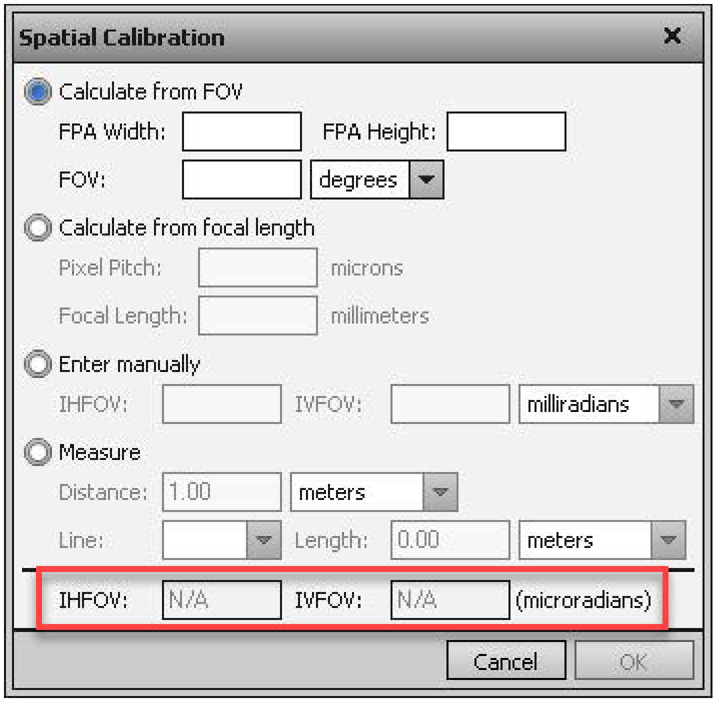


**Figure 38. View of spatial calibration window.**

But now if you look at the Misc tab of the Calibration Wizard, the tags are there:

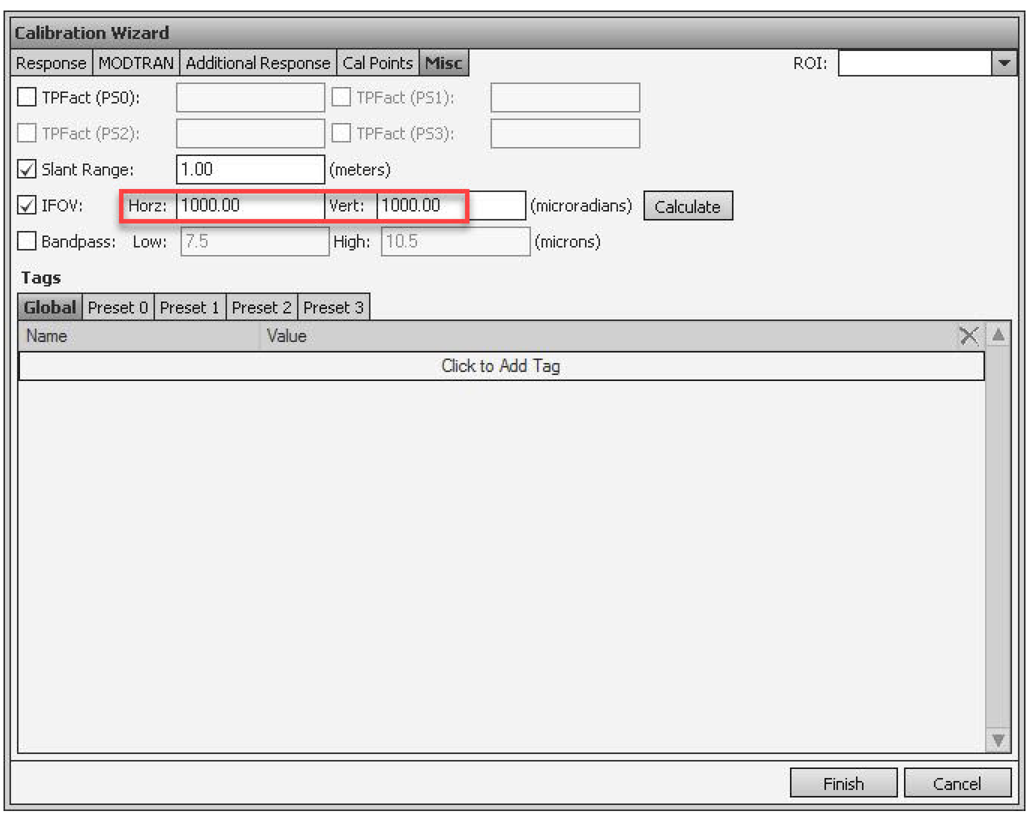


**Figure 39. Active IFOV tags from Include file shown in Miscellaneous tab**

The IFOV tags are also in the Include file, which because it is a simple text file, you can also look at in the Windows utilities Notepad or WordPad.

```
HdSize auto
Group Misc Tags
BgFile_0
BgType_0 none
BgValu_0 0
CAFile -20C to 150C 95pt8559usec 25mm SN1 Lens Calibration.cal
Coeff0 -1.739221e-03
Coeff0_0 -1.739221e-03
Coeff1 1.172209e-06
Coeff1_0 1.172209e-06
DaUnit_0 W/(sr-cm^2)
EuRaw_0 Raw
IHFOV 1.000000e+03
IVFOV 1.000000e+03
PolyOrder_0 1
SBPLo 7.500000e+00
SBPUp 1.050000e+01
SltRng 1.000000e+00
StdUnt_0 17
TempCoeff0 -5.877592e+01
TempCoeff0_0 -5.877592e+01
TempCoeff1 4.500317e+04
TempCoeff1_0 4.500317e+04
TempCoeff2 -7.509197e+06
TempCoeff2_0 -7.509197e+06
TempCoeff3 9.549586e+08
TempCoeff3_0 9.549586e+08
TempCoeff4 -7.223645e+10
TempCoeff4_0 -7.223645e+10
TempCoeff5 2.912450e+12
TempCoeff5_0 2.912450e+12
TempCoeff6 -4.805150e+13
TempCoeff6_0 -4.805150e+13
TempPolyOrder_0 6
TPFact 1.0
TPFact_0 1.0
DATA
```

**Figure 40. A screen shot of the IFOV tags in an include file. The values are 1000 microradians.**

29) Now we have two new options for display units: Temperature (user) and Radiance (user). You have to scroll down a little to see all five options. I found out after the fact that you generally do not want to have both user and factory calibrations loaded simultaneously. It is easy for ResearchIR to get confused. We recommend that if you are creating a user cal, you should unload the factory calibrations from all the presets, not just the one you are using for the user cal, which is usually Preset 0.

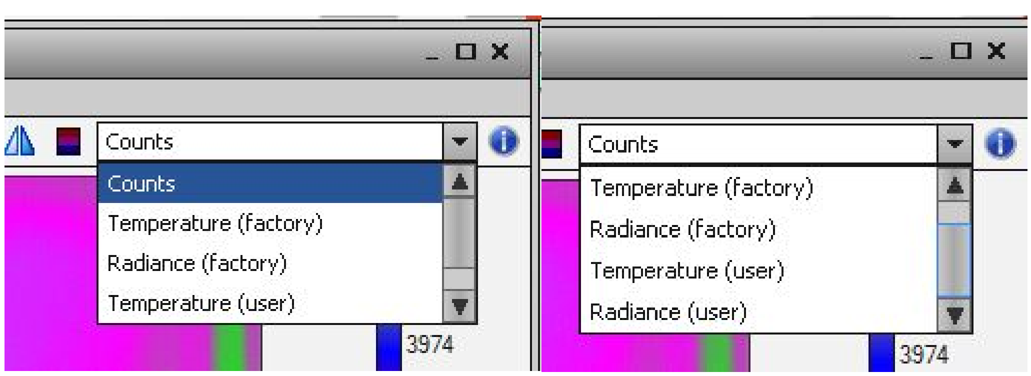


**Figure 41. All the options for displayed units for this camera. You should always unload all the factory calibrations first to avoid confusion!**

30) I tested the user calibration by looking at a known temperature and emissivity target and recording 16 frames to see how well the calibration works. The target is 29 °C and the emissivity is 0.96. The results are very decent. The apparent temperature of a 29 °C blackbody is 29.9 °C with the emissivity correction. That is much better than the rated spec of 2 °C or 2% of the temperature in °C, whichever is greater. But in a sense, this is a little bit of a circular argument, since I used the same blackbody that I used to create the calibration as I used for a target. So if my 0.96 choice emissivity setting was a little off of the "true" value, that value was "burned in" and we get very accurate temperature measurements, without having to know what the actual emissivity was. It might be a good idea to use a different blackbody to test the calibration accuracy, particularly a cavity blackbody which has an emissivity that is typically 0.99 or greater, or to get one's area blackbody in-band emissivity measured by a third party like Surface Optics in San Diego.

I took a 16-frame movie of the target, and I got four files saved into the directory that was specified in the Record Settings. They are the .cal, the .inc, the .pod and the .sfmov:

| | | |
|---|---|---|
| Top Hat User Cal 29C ISDC 2106 BB 25mm Lens SN1-000005.cal | CAL File | 51 KB |
| Top Hat User Cal 29C ISDC 2106 BB 25mm Lens SN1-000005.inc | INC File | 3 KB |
| Top Hat User Cal 29C ISDC 2106 BB 25mm Lens SN1-000005.pod | POD File | 7 KB |
| Top Hat User Cal 29C ISDC 2106 BB 25mm Lens SN1-000005.sfmov | SFMOV File | 10,243 KB |

**Figure 42. Files saved by ResearchIR with a user calibration loaded**

With a Factory Calibration loaded and no user calibration loaded, I would only get the last three files, NOT the *.cal file. Any additional information needed for the Factory calibration lives inside the camera in an XML file that we do not give the customer access to.

***Note that the Object Parameters fields are not greyed out even though my camera system is operating with an active user calibration and not an active factory calibration. The object parameters do not work in this case – none of them do anything except the emissivity control for temperature. Radiance (user) is completely unaffected by the Object Parameters settings being changed. I like to close that Tool panel so I don't start being tempted to enter values in there hoping that it will correct my measurements.***

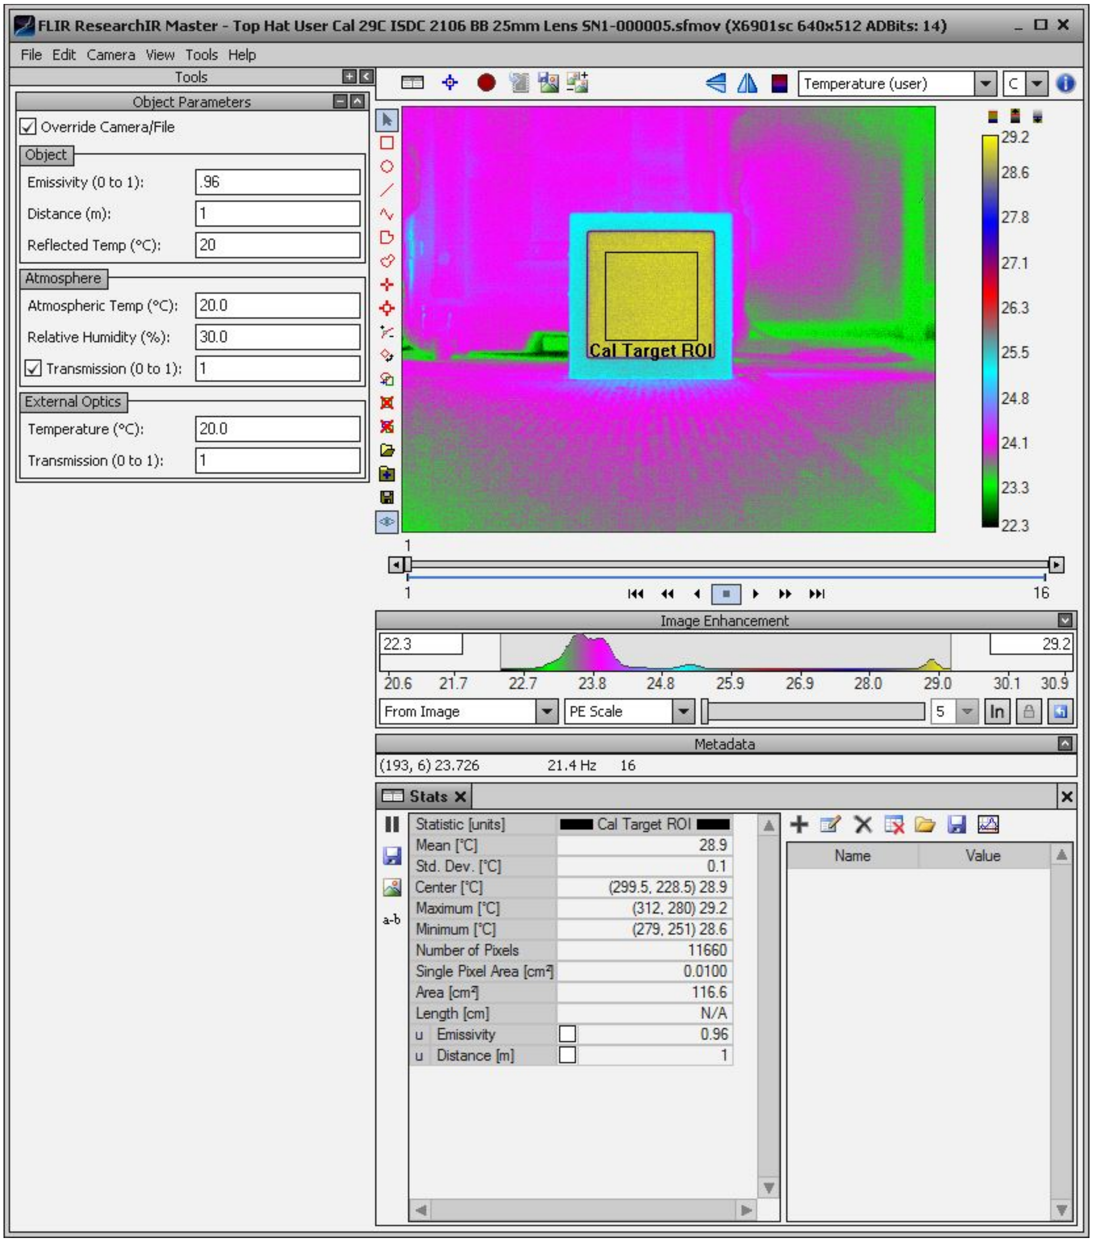


**Figure 43. Don’t’ try to use the Object Parameters when a user calibration is loaded.**

31) Note: It is generally a good idea to unload the factory calibrations from EVERY preset before embarking on a user calibration measurement. If you don’t unload all of them (not just the active preset you are using for the user calibration), then the *.include files you get will have both types of tags in them: factory calibration tags that start with the string “nfc” (short for Niceville Factory Calibration, followed by a user calibration tag). Failure to unload the calibrations may cause problems in ResearchIR and should be avoided, because otherwise you may find that the temperature and radiance measurements don’t work properly. Similarly, when using factory calibrations, it is a good idea to make sure you have no user calibration loaded. Otherwise, the *.include files that are generated

when you take a movie will have user calibration tags in them that may confuse ResearchIR. User calibrations require that the user pay close attention to what they are doing. For example, if I have a user calibration matched to the integration time for Preset 0 loaded up in ResearchIR, and then I change presets to one with a different integration time, I will get something like this Stats result, where temperature values are obviously not correct:

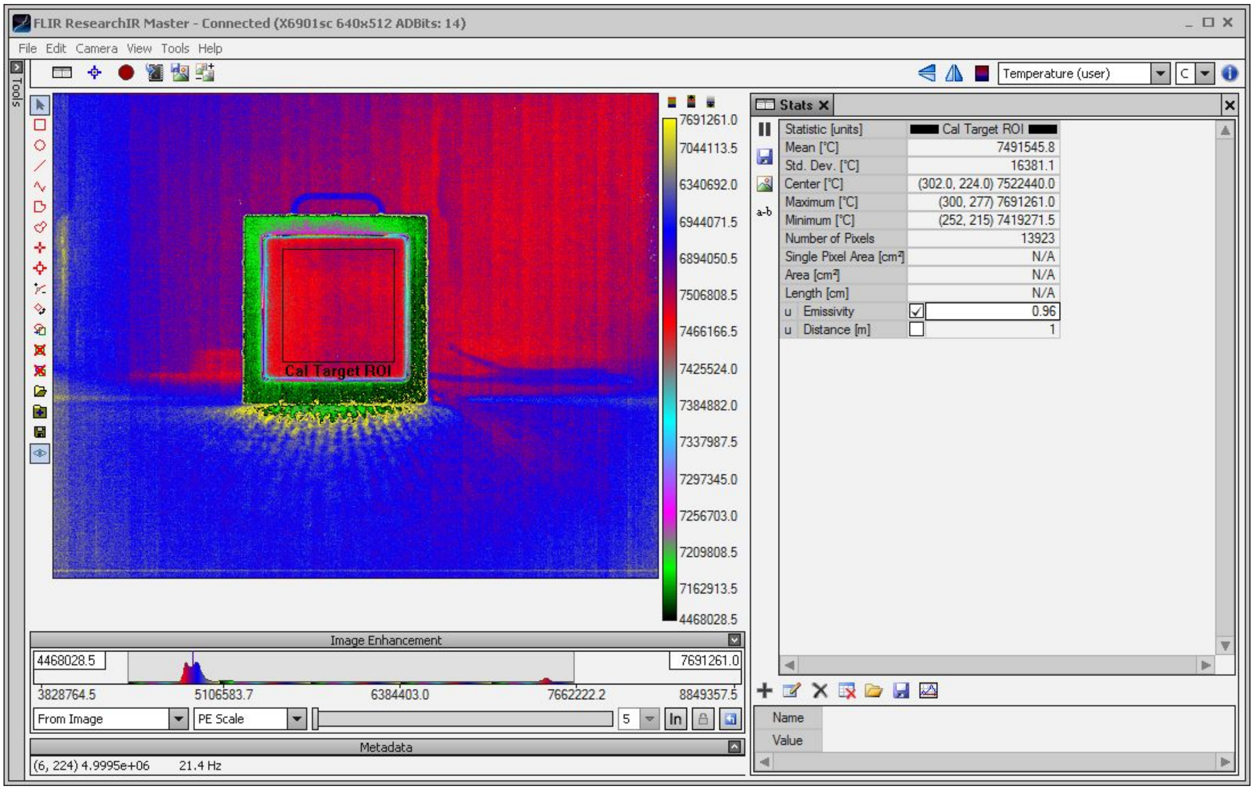


**Figure 44. That mean temperature is as hot as the interior of the sun! This cannot be physical.**

32) After creating the pair of calibration files with the Tophat spectral response, I opened the Calibration Wizard in ResearchIR, and changed the spectral response to use a spectral response file (*.prn file) to get more accurate radiance measurements.

Here are the steps:

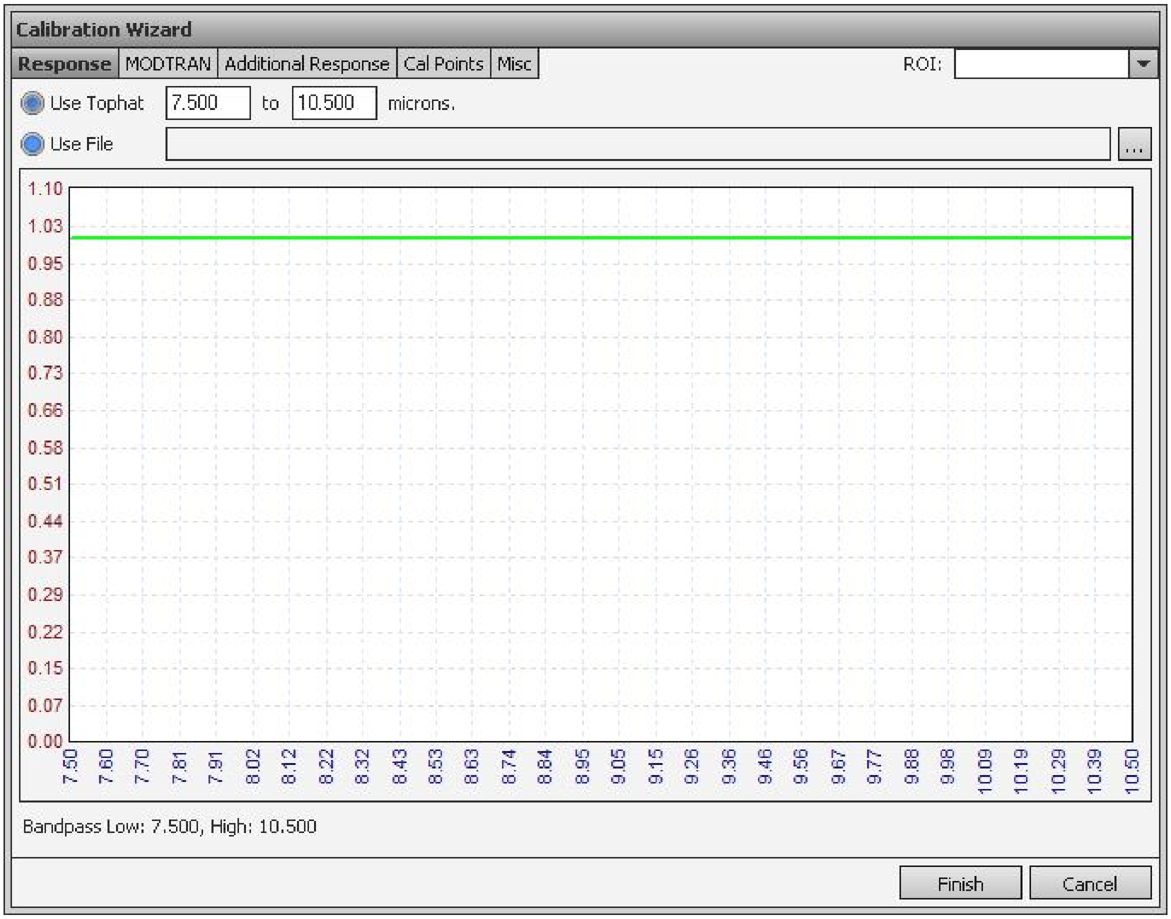


**Figure 45. Response tab of the Calibration Wizard. Hit radio button to change to a File for the response.**

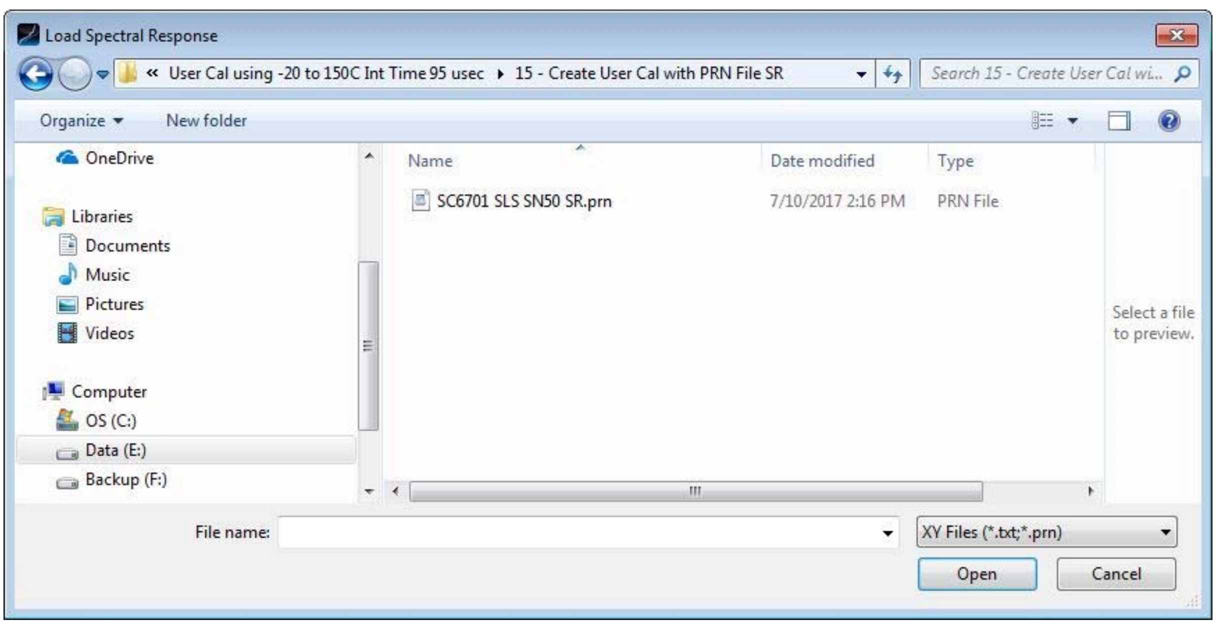


**Figure 46. Loading a spectral response file for this camera**

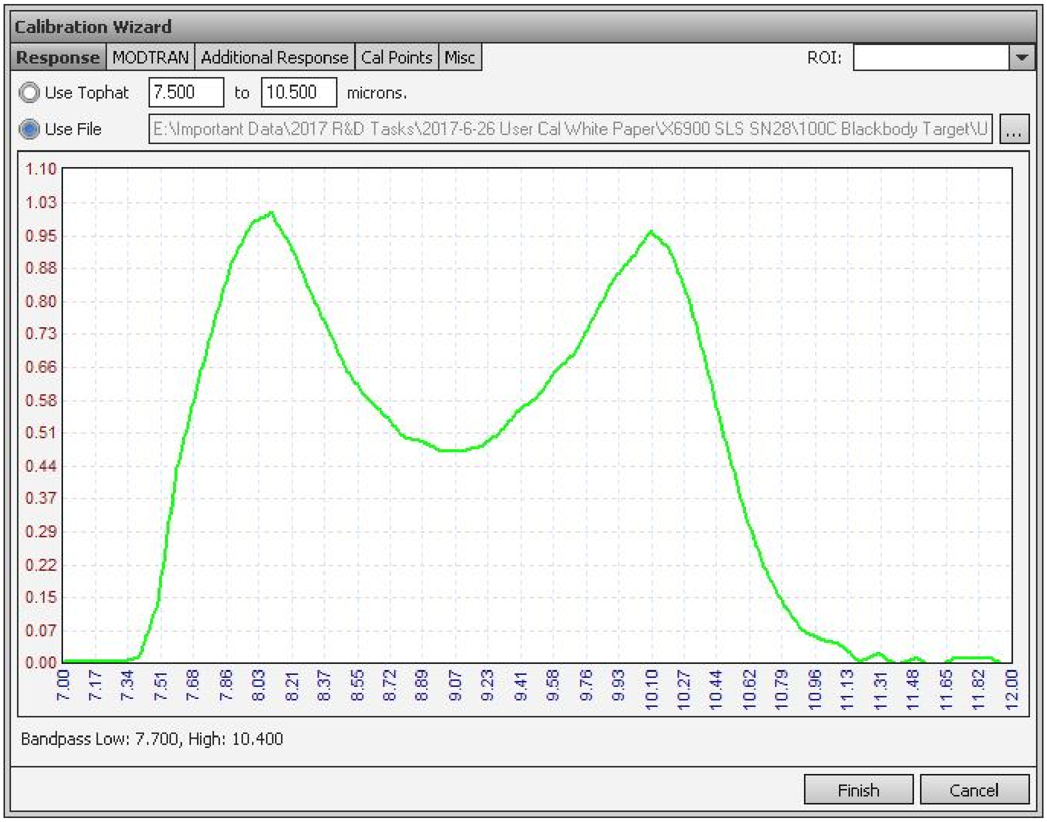


**Figure 47. The spectral response curve that I loaded.**

I then saved the calibration files with a name that indicated that it was different from the earlier tophat calibration.

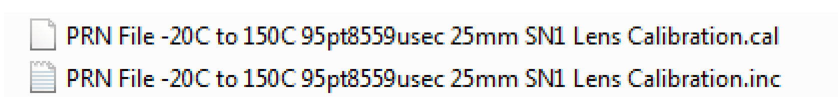


**Figure 48. The PRN File prefix shows that the calibration has a spectral response file loaded, not a tophat**

The two calibration files are different in terms of the radiance values they will yield, but they will both give very similar temperature measurements.  It makes sense that they would – the ROI mean digital counts in the measurement grid for each calibration point temperature are identical. Only the mapping to *radiance* is different because of the different spectral response choices.

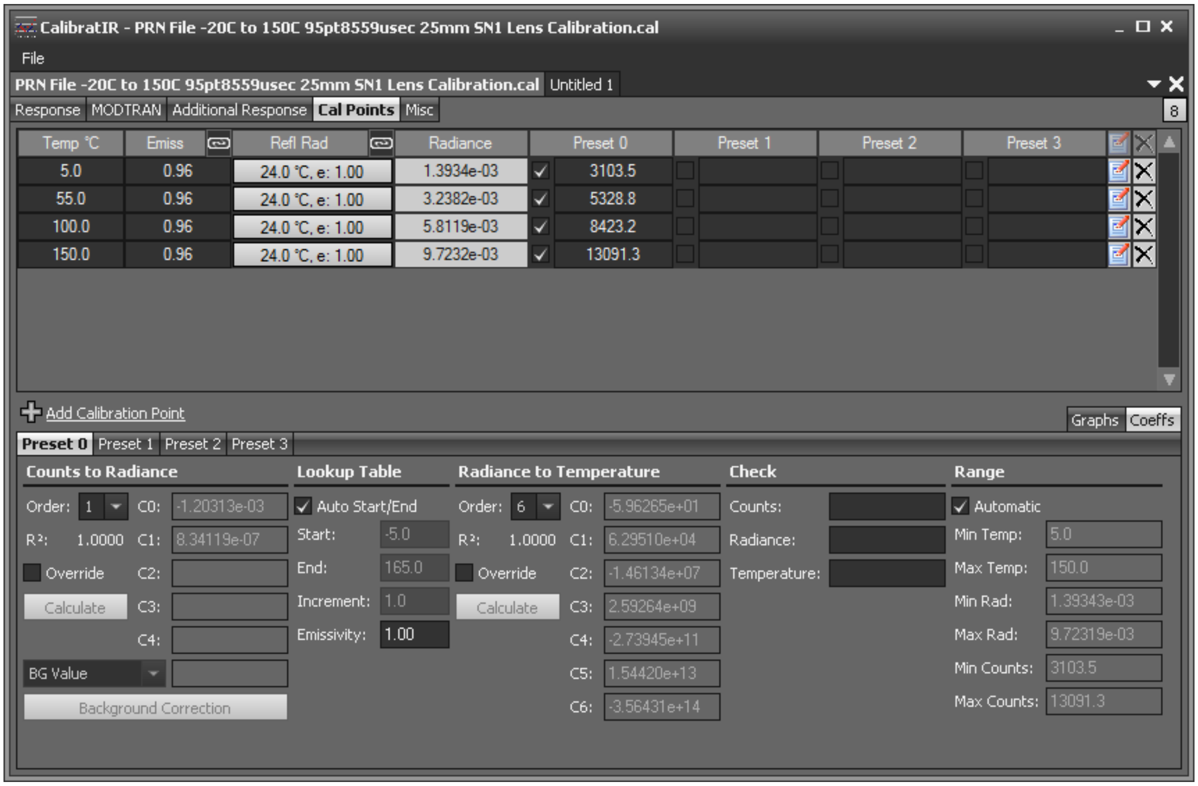


**Figure 49a.  Note the 5 °C radiance value of 1.3934e-3 for the PRN file spectral response**

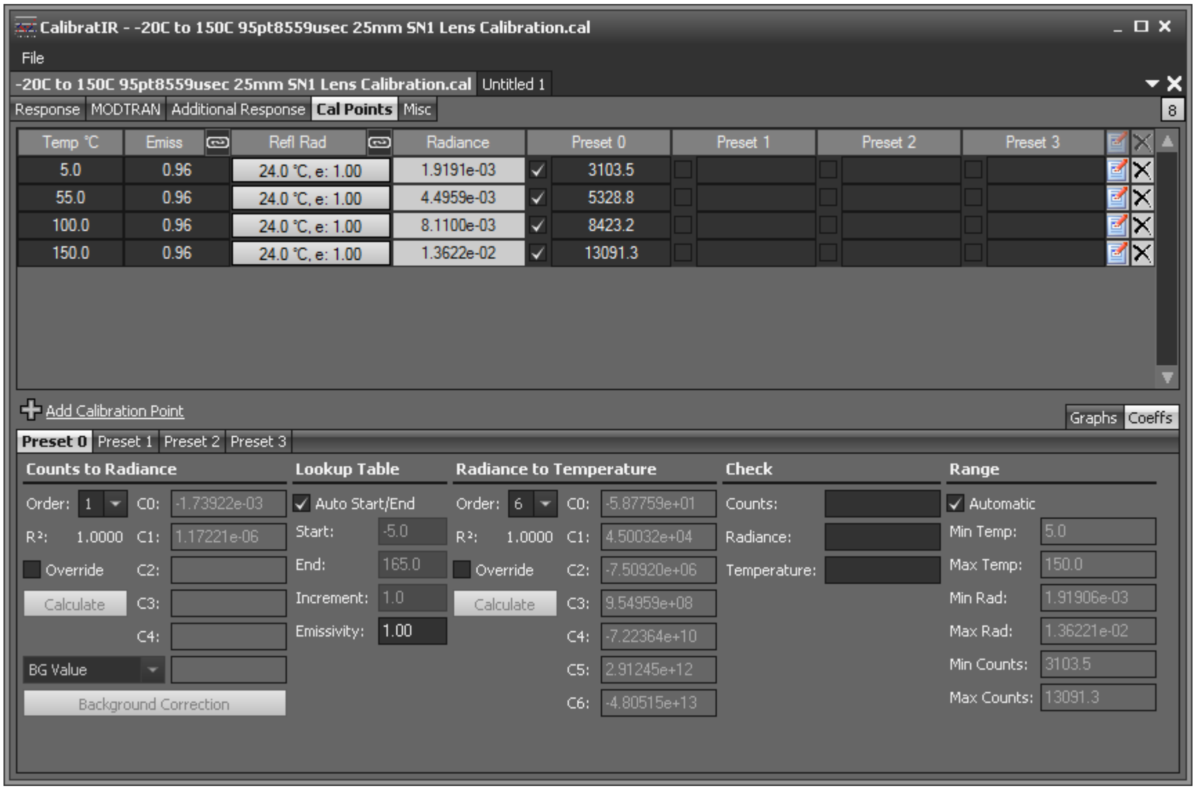


**Figure 49b. The same 5 °C point gives a different radiance value of 1.9191e-3 when I use the 7.5-11µm tophat limits.**

Figure 50a shows what happens to the tophat calibration graph of radiance versus counts if I set the spectral band limits to 3 and 5 microns. The fit is much worse (0.9499 $R^2$ value) and the curve is no longer highly linear. This is a warning that something is not right. Note also the * symbol in the ribbon at the top in front of the file name. That tells you that the file is no longer matching the last saved version, another clue that your co-worker has just monkeyed with the file while you were at lunch!

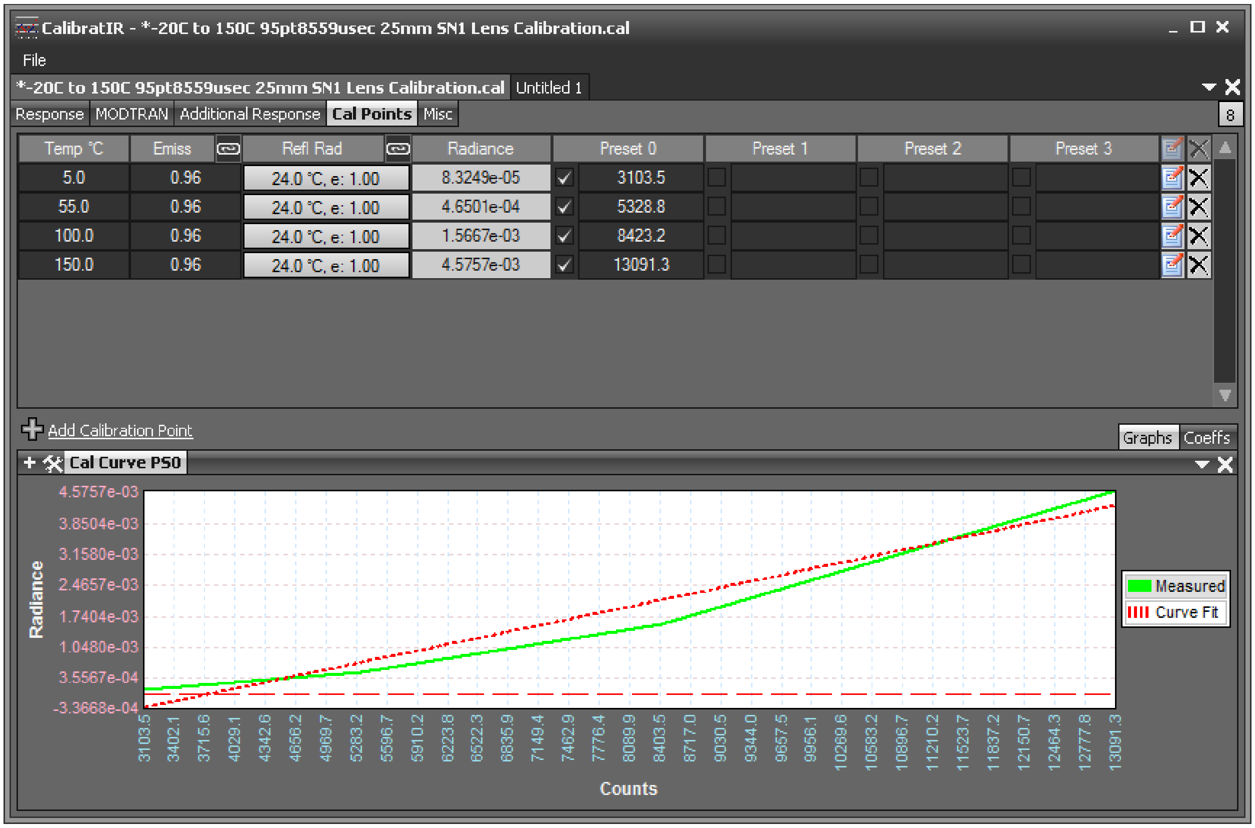


**Figure 50a. Bad fit of measured radiometric transfer function to a straight line caused by incorrect band limits set in Response tab**

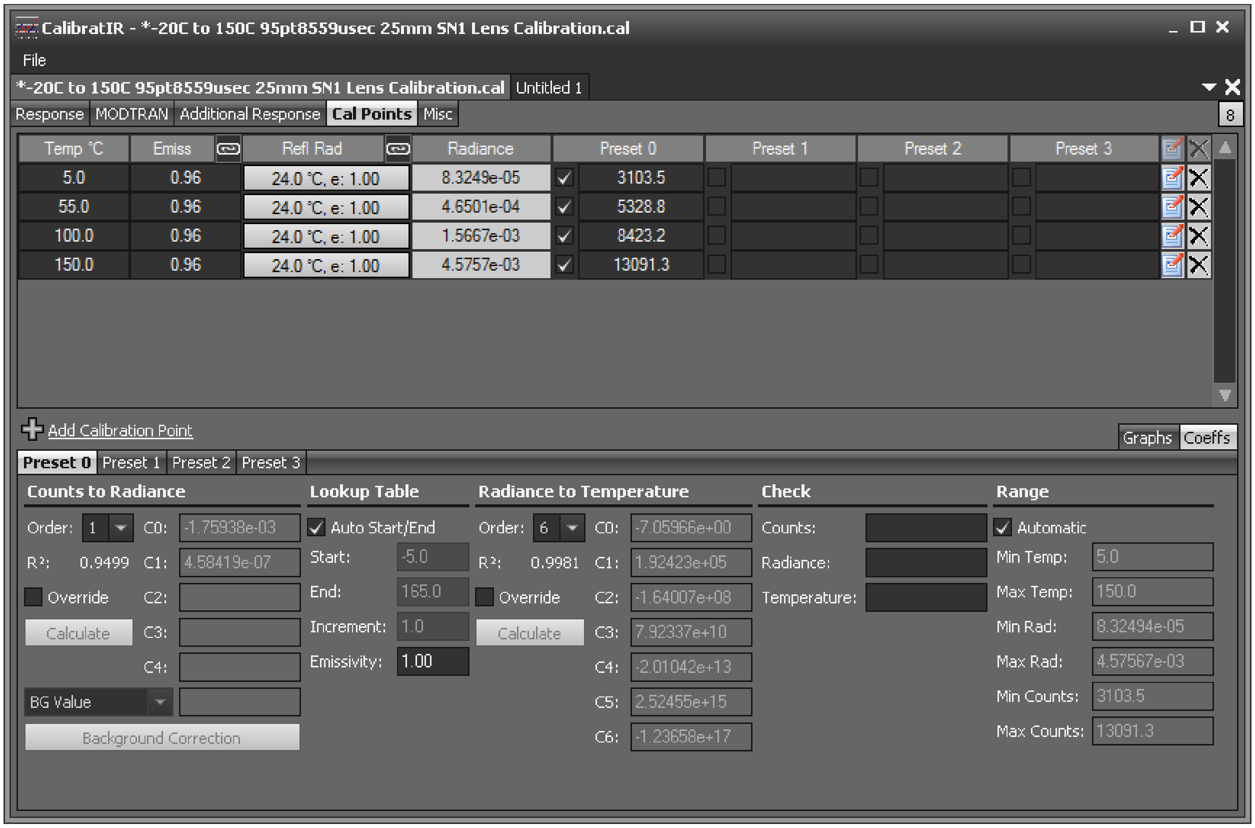


**Figure 50b. Lousy fit shown in $R^2$ display box on the left.**

Using 7.5 to 10.5 microns for the response band limits works well and gives $R^2 = 0.9999$. I also tried 7 to 11 microns and I got $R^2 = 0.9999$ as well.

# Chapter 10 : User and Factory Calibration Comparison

# Contents

Appendix 2 : User and Factory Calibration Comparison ........ 10—395
Introduction: ........ 10—397
Standard Factory Calibrations ........ 10—397
$T_{Drift}$ Correction ........ 10—401
User Calibrations ........ 10—402
Factory Calibration Accuracy Evaluation ........ 10—403
Temperature Image ........ 10—406
Radiance Image ........ 10—409
Radiance measurements are sensitive to spectral response ........ 10—411
Using Two Different Ranges of Factory Calibrations to Measure the Same 100C Target ........ 10—412
Comparison of Factory Calibration to User Calibration ........ 10—415

## Introduction:

For FLIR science cameras, both factory and user calibrations are essential in transforming the raw data captured by the camera into units the engineer or scientist can utilize for quantitative measurements, as well as intuitively understand. Of the two types of calibrations, factory calibrations provide several advantages by accounting for systematic effects on measurements that are often overlooked during the use of user calibrations, or which are not easily dealt with when using the user calibration. To this end, most of the FLIR science cameras constructed in Niceville, Florida can be factory calibrated to provide radiometric measurements of scenes and targets. The exceptions are the RS6700, RS8300 and other long-range cameras, which are not factory calibrated as a standard practice because of the difficulties associated with corrections for atmospheric transmission over long air paths, and because factory calibration on long-range cameras should be done with both a collimator and a thermal chamber for the camera, and these assets are not in physical alignment in the Niceville facility.

A camera that is factory calibrated or user calibrated can be used to measure both apparent temperature and in-band radiance. In-band radiance is the radiance of the scene measured in watts/sq.cm/steradian in the spectrally-weighted bandpass of the camera system. Apparent temperature is the object temperature that the camera “thinks” it is seeing while capturing data. This varies from the kinetic surface temperature, which is the actual temperature of object surfaces in the scene – the temperature you would measure with a contact temperature probe. Achieving a good correlation between this apparent temperature and the kinetic surface temperature can be complex for scenes and objects that do not emit continuum radiation like a theoretical Planck blackbody or graybody.[21] Both the accuracy of the information on the physical parameters of the setup during calibration, and the accuracy of the physical conditions during data acquisition determine how accurate both the in-band radiance and apparent temperature measurements are going to be. That is why it is important to record air temperature and humidity measurements during user calibration, and also during experimental data collection, especially if high accuracy is a goal. It is also vital to know the emissivity of the calibration blackbodies used to create a user cal. Erroneous emissivity values will “bake in” errors in the user calibration, and limit measurement accuracy.

## Standard Factory Calibrations

The Niceville factory calibrations are done in the plant in Niceville, Florida using a “calibration arc” or a series of blackbodies arranged in a curved arc. A camera is placed on a swing arm that moves the camera in front of this series of calibrated laboratory blackbodies, which are arranged in a progression of temperatures. The response of the camera in digital counts is then measured for each blackbody. This process can be done over a very wide range of blackbody temperatures, from -20 °C to 3000 °C, depending on user needs. The standard calibration for a MWIR camera is -20 °C to 350 °C. The calibration file set consists of five different factory

[21] A graybody is a blackbody with an emissivity that is less than 1 and spectrally flat.

calibration sub-ranges to cover overlapping temperature ranges (-20-55 °C, 10-90 °C, 35-150 °C, 80-200 °C, 150-350 °C). The high temperature calibration options for a standard 3-5µm InSb cameras use a neutral density filter (ND2) to reduce the incoming radiance, and these calibrations have temperature spans from 150 °C up to 2000 °C, in three overlapping ranges (250-600 °C, 500 °C to 1200 °C and 850-2000 °C). The blackbodies used in the calibration arc are calibrated periodically and have known emissivity values of 0.99, since they are cavity-type units whose emissivity values are driven more by geometry than surface coatings.

These factory calibrations are done in concert with factory NUCs (non-uniformity corrections) that are tied to the radiometric calibrations in the camera software. It should also be noted that these calibrations are done for each lens that is used with the camera, since different lenses have different transmission values. The factory calibration process for a camera that has a number of different lenses and desired temperature ranges can be lengthy and expensive.

Operating the camera with the factory calibrations is simple (much simpler than using any user calibrations a user creates). The camera user can select a pre-calibrated measurement temperature range from a pull-down menu in the Setup/Presets section of the ResearchIR (ResearchIR) camera controller GUI. The left image in Figure 29 shows an X6900 SLS LWIR camera GUI set so that preset 0 has a -20 to 150 °C calibration for a 25mm lens, while the right image shows the pull-down menu. One of the options is “No Factory Calibration,” which puts the camera into a mode where the integration time can be changed by the camera operator.

**Note: To avoid confusing the camera or ResearchIR, it is advisable to set ALL four of the presets to No Factory Calibration before creating or using a user calibration. The advantage of unloading all the factory calibrations from every preset is that the include file will be quite a bit simpler, as it will not contain any of the Niceville factory calibration tags and when a movie file is loaded into ResearchIR, the display units won’t default to Temperature (factory), which will result in a black screen if there is no factory calibration for the displayed preset.**

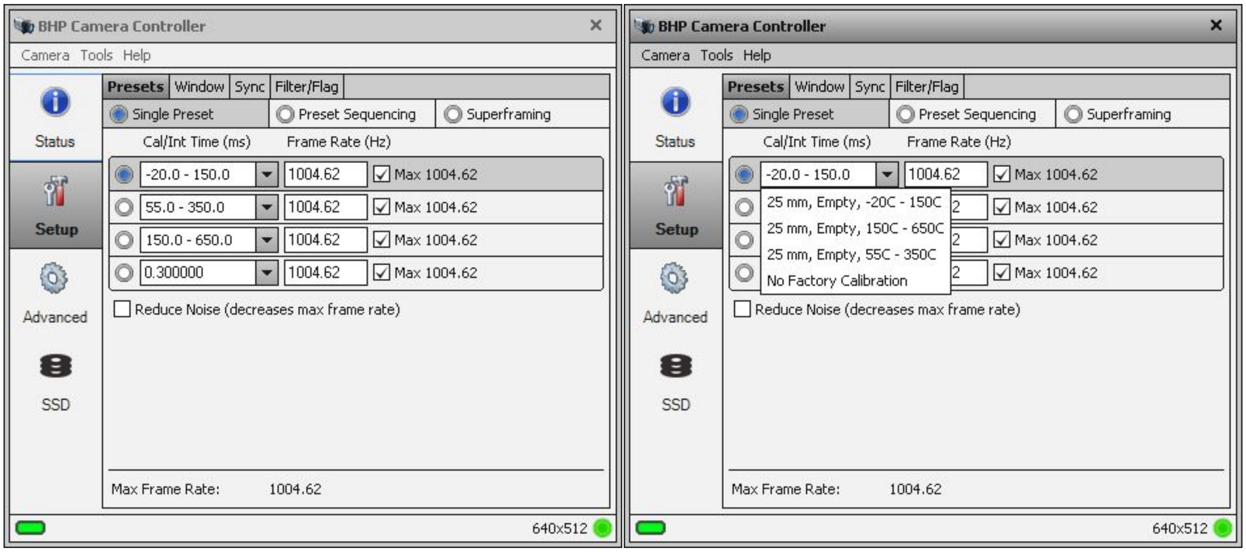


**Figure 29: (Left) An X6900 SLS LWIR camera controller GUI in ResearchIR. (Right) The pull-down menu.**

The factory calibration names are set at the factory and cannot be edited by the user. Selecting thoughtful names for user calibrations is important as these file names help the user make sure that the correct user calibration is loaded into ResearchIR. For example, the X6900 SLS camera shown in Figure 29 has only three factory calibrations. If there were any high temperature factory calibrations for this camera, they would not say "Empty" after the lens focal length – they would instead say "25mm, ND1, *temperature range*." The focal length of the lens is always specified in the longer version of the factory calibration name you see when you pull down the menu. The camera does not have a means to know what interchangeable lens is installed, so the user must select the correct calibration for that lens type, whether it is a factory calibration or a user calibration.[22] Typical lenses that are calibrated as a standard offering are 17mm, 25mm, 50mm and 100mm. Longer focal length lens factory calibrations need to be done using folding or collimating mirrors and are not standard product offerings. Instead, they require a special engineering request and increased labor cost and lead time.

It is possible to use a 25mm factory calibration with a 50mm lens without introducing significant measurement error, but, it is not recommended because the lenses have different transmission curves. The 25mm HDC lenses have a ~4% lower transmission because they have an additional distortion-correcting lens element compared to the 50mm lenses.

The inclusion of "ND1" or "ND2" in the long version of the factory calibration name refers to the inclusion of a neutral density filter with an optical density of 1 or 2 in the optical setup. Neutral density filters are used to decrease the amount of radiation reaching the camera uniformly over all wavelengths over which the camera is sensitive. An ND1 filter transmits $10^{-1}$

[22] We do not currently have any communication between the standard set of prime interchangeable lenses and the camera that would tell the camera what lens was installed or its focus position. The standard prime lenses are not motorized or equipped with electronics.

or 1/10 of the incident radiation. High and ultra-high temperature calibrations may require ND2 or ND3 filters, which pass 0.01 and 0.001 of the incident radiation, respectively.

Users must take care to ensure all of the different files required are changed when applying different calibrations. When the pulldown menu is used, the camera software does not simultaneously load the NUC associated with the chosen calibration range. **The camera operator should take care to make sure the correct NUC is loaded in the corresponding preset.** The camera comes from the factory with the factory calibrations and NUCs loaded in the correct presets. It is always a better idea to change the active preset if you want to access a different temperature range, rather than changing the loaded calibration and then having to also change the loaded NUC for that preset.

The factory calibrations and factory NUCs are both stored in the camera itself, rather than on the host PC. In contrast, “user calibrations” and “PC-side corrections” are stored on the host computer. Figure 30 shows that the controller GUI has four different preset options, which are all taken up by factory calibrations in this case. While it is possible to change the order of these presets, there is little need to change which factory calibration is loaded into a particular preset unless the camera has more than four factory calibrations, or if one is running in superframing mode. Figure 30 shows a screen shot of the camera controller window showing the factory NUCs loaded into the four presets:

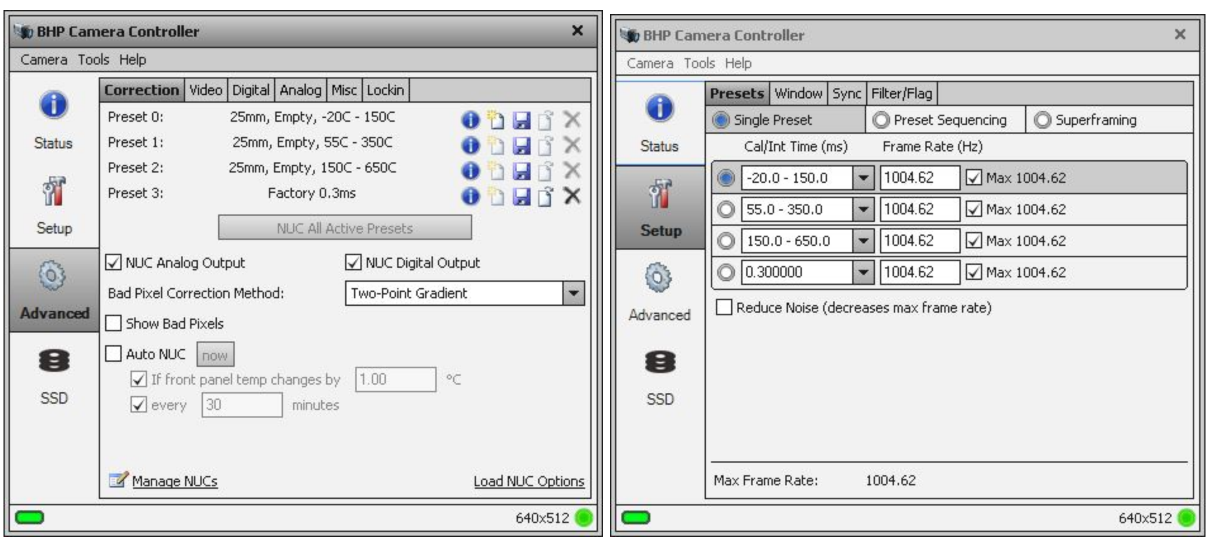


**Figure 30: A screen shot of the camera controller Advanced window showing the factory NUCs loaded into the four presets, with the Setup window next to it. Note that the factory calibrations are aligned with the correct factory NUCs, and one should generally leave the camera that way.**

The customer can complete a backup of the NUC files by going into the Manage NUCs section of the GUI and downloading them. Unfortunately, there is no easy way for a customer to back up the factory calibration files because they live in an XML file inside the camera. If those files get corrupted, the camera will generally have to be returned to FLIR for service.

The factory calibrations are extremely convenient for typical customers, because

1) The cost for the customer to buy a set of area and cavity blackbodies to create their own NUCs and user calibrations is quite considerable
2) The lead time on these devices can be quite lengthy.
3) They need periodic recalibration which is also expensive.
4) The management of the various settings and files for user calibrations is tricky and easy to get wrong, especially in field conditions with its many distractions and time pressures. A good example of this complexity is that the factory calibration in an X series camera makes sure that if a filter is needed for a calibration, the correct filter is deployed in front of the warm window. The filter holders have ID tags that the camera filter wheel can read.

## $T_{Drift}$ Correction

Another advantage of factory calibrations over user calibrations is the automatic $T_{Drift}$ correction feature, which accounts for the fact that a camera calibrated in the laboratory at ~23 °C ambient temperature will read the radiance and temperature of objects differently in very hot or very cold operational environments. This effect is caused by stray light generated inside the optics and the camera body itself. This self-radiation, as it is called, will produce a radiance offset on the FPA sensor. Put another way, the inside of the lens and optical path is never truly dark as it is in a visible-light camera. Figure 3 shows the lens mount and the inside of an SC6000 camera with the lens removed. There is a hollow cavity inside the body around the lens mount. Temperatures inside this cavity can easily get to 35 °C or higher, and that radiation reflects off the back of the lens and onto the FPA. The radiation is there when the camera is calibrated, and that level of self-radiation is “baked into” the calibration in the offset. If the temperature of the camera cavity changes, the offset changes from what it was during calibration. This offset change will happen if the camera is operated outside on a hot day in the sun, or on a cold night.

**Figure 3. The lens mount and inside of a SC6000 camera without a lens installed.**

To create the Tdrift calibration, additional factory measurements are done in a thermal chamber to correct the radiometric calibration for changes in ambient temperature of the camera/lens system. This Tdrift correction will then automatically adjust for systematic errors in the apparent temperature of targets that are caused by changes in the self-radiation of the optics and optical interfaces in the camera. It is possible for the user to do a self-radiation offset correction in ResearchIR, but it is not automatic like the Tdrift correction. The Tdrift feature is more important for measurements of ambient and below-ambient temperature targets where the optics self-radiation can be comparable to the radiance from the cold target itself. At higher target temperatures, shorter integration times are used to limit the exposure, and that makes the camera much less sensitive to operating temperature induced changes in optics self-radiation.

## User Calibrations

A user calibration is a process whereby a calibration is created inside the ResearchIR environment by the camera user. The user points the camera with appropriate lens and any filters they want at a calibration source which is either an area or cavity blackbody. The source is adjusted to a series of temperature setpoints, with digital image data taken at each temperature point. The data on the source temperature and emissivity and any reflected radiation is entered into a calibration wizard or standalone utility called CalibratIR, and a pair of files are generated. These files are loaded into the host computer running ResearchIR, and they are used to superimpose a calibration onto image data files that are subsequently collected. The user calibrations can also be added to existing uncalibrated datafiles that were previously captured, ideally with the exact same camera settings and with the same lens and optics.[23] This allows for users to go back in and edit the radiometric calibrations after the fact, something that is not

[23] Of course, one can apply a user calibration done with a 50mm lens to data captured with a 25mm lens, but there will be error introduced because of the different lens transmission.

possible with factory calibrations. See Appendix 3 of this document called Changing Include Files for more information about post-collection calibration.

User calibrations offer the customer the ability to fully customize the calibrations and do things like interpose a window in the optical path during calibration that then becomes “baked into” the calibration. Range customers are the most likely to create and utilize user calibrations as they almost always have access to a large pool of blackbody sources, and they often want to make their own custom radiometric calibrations with various combinations of interchangeable optics and filters in their often quite considerable equipment inventory.

As a rule, the factory in Niceville does NOT do user calibrations for customers. We have done them from time to time on an individual basis for customers that pay for SME support. When I do them, I use a small inventory of blackbodies that I have in Santa Barbara, or I use the customer’s blackbodies at their site.

## Factory Calibration Accuracy Evaluation

The following is a description of accuracy evaluations of existing factory calibrations in a Niceville science camera, as well as comparisons to two different user calibrations performed on the camera and running in parallel with the factory calibrations.

The tests were performed using an Infrared Systems Development Corp. Model 2106 blackbody placed 1 meter from the lens front of a FLIR X6901sc SLS camera with a 25mm lens on it, as shown in 4. It should be noted that it is highly recommended that the user capture photographs of the setup AND the nameplates on the major components for any test series or user calibration completed, as shown in Figure 4 and 5, so that he or she can go back later and figure out what was done, or exactly what equipment was used. This is especially important for Range customers that may have multiple examples of particular cameras and lenses!

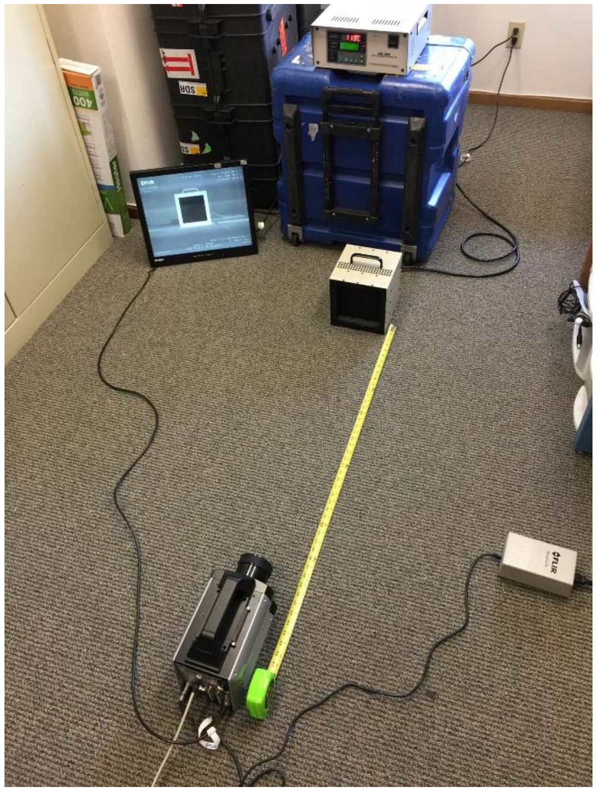

**Figure 4: X6901 SLS placed 1 meter from ISDC Model 2106 area blackbody**


**Figure 5: Camera and lens nameplates**


**Figure 6: Blackbody Head and controller nameplates. This ISDC Model 2106 blackbody has an emissivity rating of 0.96 +/- 0.02. It is perhaps not ideal for precise accuracy measurement of cameras.**

During this test, the blackbody (shown in Figure 6) was initially set to 100 °C and allowed to stabilize for 30 minutes prior to measurement. This settling time should never be skimped on during any type of test involving a blackbody, and the camera should also be powered on and allowed time to stabilize before any radiometric tests or measurements are done. This area blackbody is a bit slower to stabilize because it is a less expensive alternative to SBIR and CI Systems products, which are 50+% more expensive on average but settle faster. Note that sometimes the blackbody controller will say that it is stable at the controller setpoint, but the observing camera's ROI digital counts will still be increasing or decreasing. The controller reads out a temperature sensor, not the actual emitting surface temperature. I like to use the camera to confirm stabilization.

Cameras can take a while to stabilize themselves. I always install the lens I will use for the user calibration and leave it on to stabilize at the camera temperature, since heat will flow into the lens from the camera lens interface. Science camera bodies will continue to heat up for 60 minutes or more, unless the camera is bolted down to a big heatsink or a chiller plate. You can monitor the heating of the camera by looking in the "Misc" tab of the camera controller GUI and watching the Front Panel, Air Gap and Internal Chassis temperatures rise over time. These three temperatures can also be monitored using a live temporal plot of a measurement function in ResearchIR as well. When they have stabilized, then the user calibration can begin. If the user calibration is being performed for a short integration time, like 1msec or less, then the duration of this stabilization is not as important, because the imaging sensor will be much less sensitive to changes in self radiation if its integration time is set for a 100 °C or above scene temperature range, for example.

Figures 7 is a screen capture of the main ResearchIR screen with the statistics window showing statistics for digital counts, then temperature and radiance. Note the emissivity parameter for the

ROI is set to the nominal value of 0.96, though that does not affect the statistics when the units are counts.

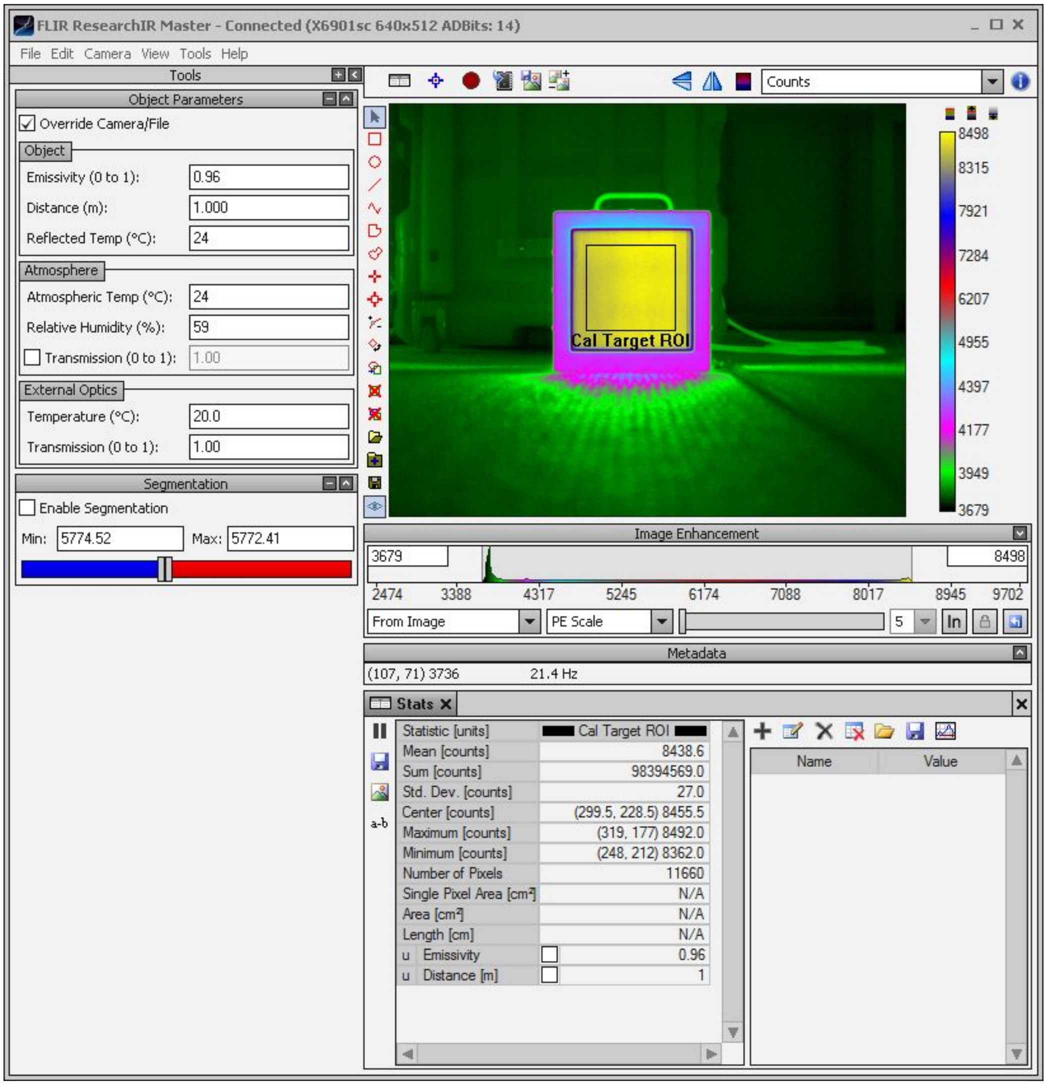


**Figure 7: The counts image has no processing done to it other than non-uniformity correction – it is the digital counts from the digitizer connected to the focal plane array (image sensor).**

## Temperature Image

The temperature image shown in Figure 8 has many processing steps done to the digital counts to convert them into apparent temperature. First, the raw digital counts are converted into radiance by ResearchIR, and the Tdrift correction is applied to adjust the radiance offset correctly to account for the operating temperature of the camera and the readings of the temperature sensors inside the lens interface and camera body.

The emissivity is set to 0.96 (or 96%) and the reflected temperature is set to 24 °C (measured by a tabletop thermometer/humidity meter) in the Object Parameters window on the left side of ResearchIR. ResearchIR computes an in-band radiance offset that is then applied to correct for the 4% reflection of the 24 °C ambient room. Thus, the radiance going forward is the emitted radiance only, without a reflected component. The distance is set to 1 meter, which was the distance used during these tests.

The air path transmission is calculated using a proprietary equation that computes the total water vapor in the air path, using the atmospheric temperature, the relative humidity and the distance (length) of the air path. For this experiment, the 1-meter target range is so short that the calculated atmospheric transmission is 1 and there is only a negligible path radiance to account for. Longer distances will result in lower transmissions, down to a software limit of 0.4 transmission. The user can override that computed value and put in even lower values that hopefully are based on air path transmission measurements, as this is a powerful knob to turn!

If there are no external optics (like a window into a chamber) in the optical path, then the external optics transmission parameter is also set to 1. If the external optics transmission is set to something other than 1, then the external optics temperature Object Parameter affects the temperature measurement. The thinking is this: Suppose you have a window that has a transmission of 0.9 and is at 40 °C temperature. Then the apparent temperature of the target as seen through that window is affected in two ways. The radiance is reduced by 0.9 and there will be an extra contribution of radiance from self-radiation of the 40 °C window that acts like an emitter with a (1-transmission) emissivity of 1-0.9 or 0.1. The software assumes there is no window reflection, which is not a bad assumption for AR-coated IR windows.

Here are all the correction steps in order, for the case of a 0.9 transmitting window at 40 °C:

1) The software calculates the in-band radiance of a 40C window with a 0.1 emissivity and subtracts that from the radiance measured by the FPA.
2) Then the software divides that corrected radiance by 0.9 (window transmission) to get the radiance before the window.
3) The software then subtracts off the atmospheric path radiance, and divides by the atmospheric transmission to get the radiance at the target. The air path radiance is the in-band radiance of an emitter at the atmospheric temperature, and an emissivity of (1-T), where T is the air path transmission.
4) The reflected radiance off the target is (1-emissivity)*R($T_{Reflected}$), and that is subtracted from the radiance at the target.

The radiance at the target that has been corrected for all these downstream effects is now scaled by dividing it by the surface emissivity. This is the radiance of the target as if it were a blackbody with an emissivity of 1. That fully corrected radiance is then used to determine the temperature of the target using a lookup table, or a polynomial fit to the curve of radiance versus temperature.

The fully corrected radiance has been corrected for all the systematic effects from the target to the FPA. These are, in order:

1. Emissivity of the target
2. Reflected radiance from the target
3. Atmospheric transmission (a function of distance, air temperature and relative humidity)
4. Air path radiance (not explicitly shown but calculated in the background)
5. External Optics transmission
6. External optics self-radiation (only if T is less than 1)

In the case of the 100 °C source I imaged, the measured apparent temperature is 98.1, which is an accurate measurement, as it is within 2 °C of the indicated temperature on the blackbody controller:

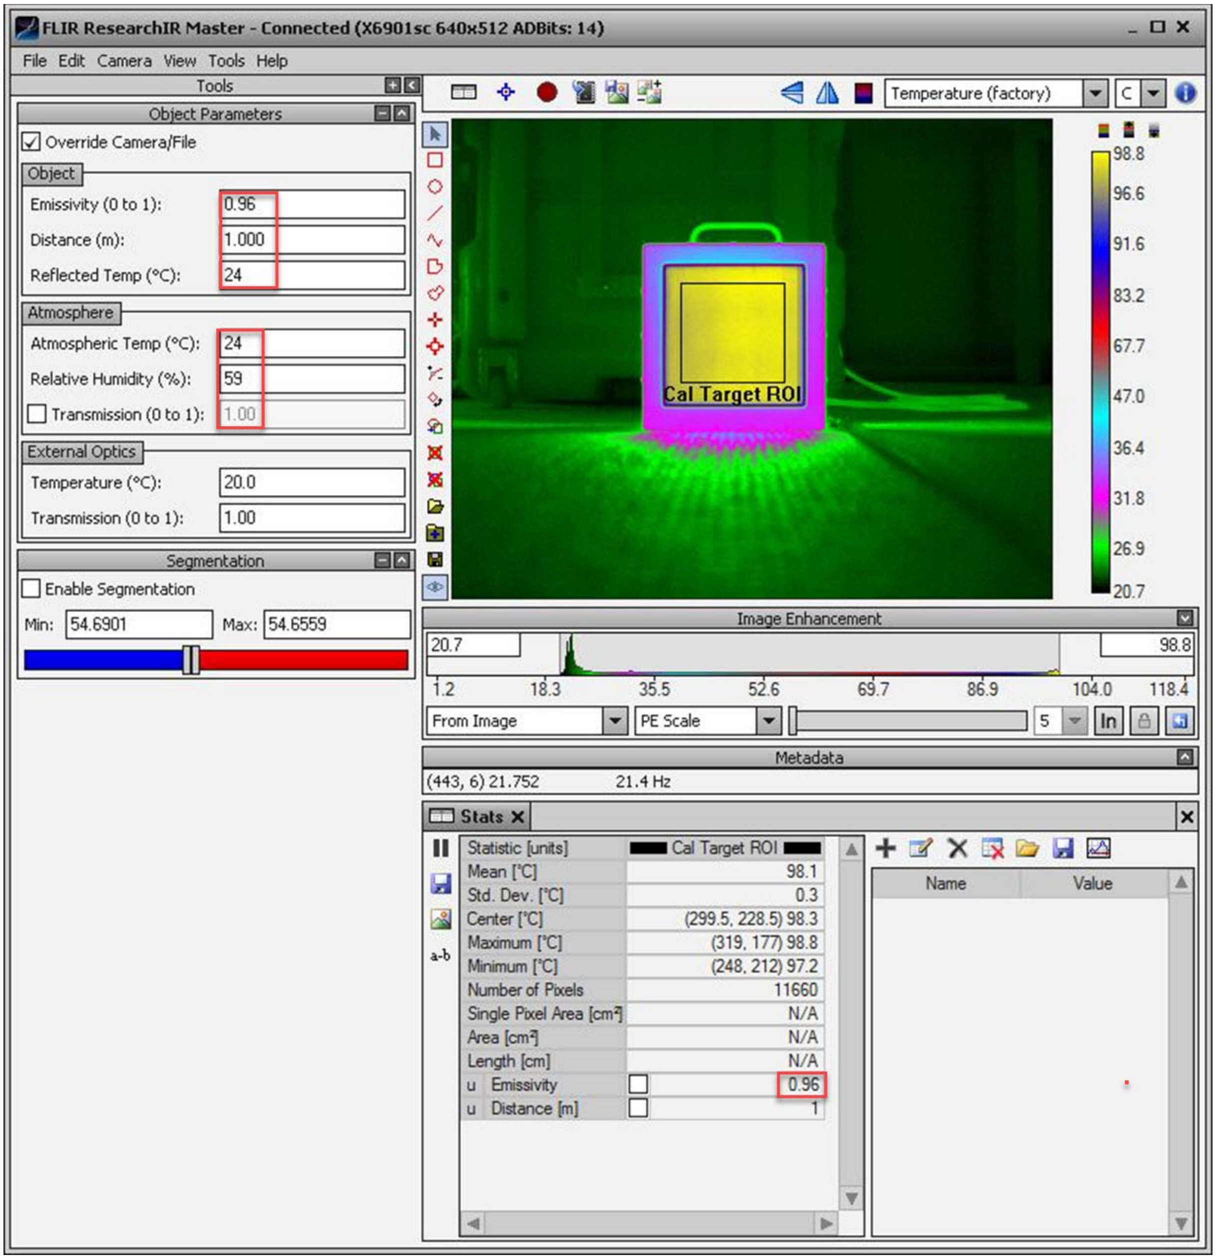


**Figure 8: An image of the blackbody in temperature units with an emissivity of 0.96 set in the Stats window and also in the Object parameters window on the left. The distance and reflected temperature values were measured and set correctly in the Object Parameters window, as were the atmospheric temperature and the relative humidity, both of which combine with the distance to give a calculated atmospheric transmission value, which is shown in grey text as 1.00, typical for this short of a distance. The External Optics transmission and temperature have default values of 1 and 20 °C respectively. If the transmission is set to anything other than 1, then the External Optics temperature setting will have an effect on the measurement.**

## Radiance Image

The radiance image in Figure 9 has Tdrift correction, Atmospheric Transmission correction and External Optics correction applied to it. Note that the object parameters Emissivity and

Reflected Temperature have no effect on the radiance measurements – they only affect temperature measurements.

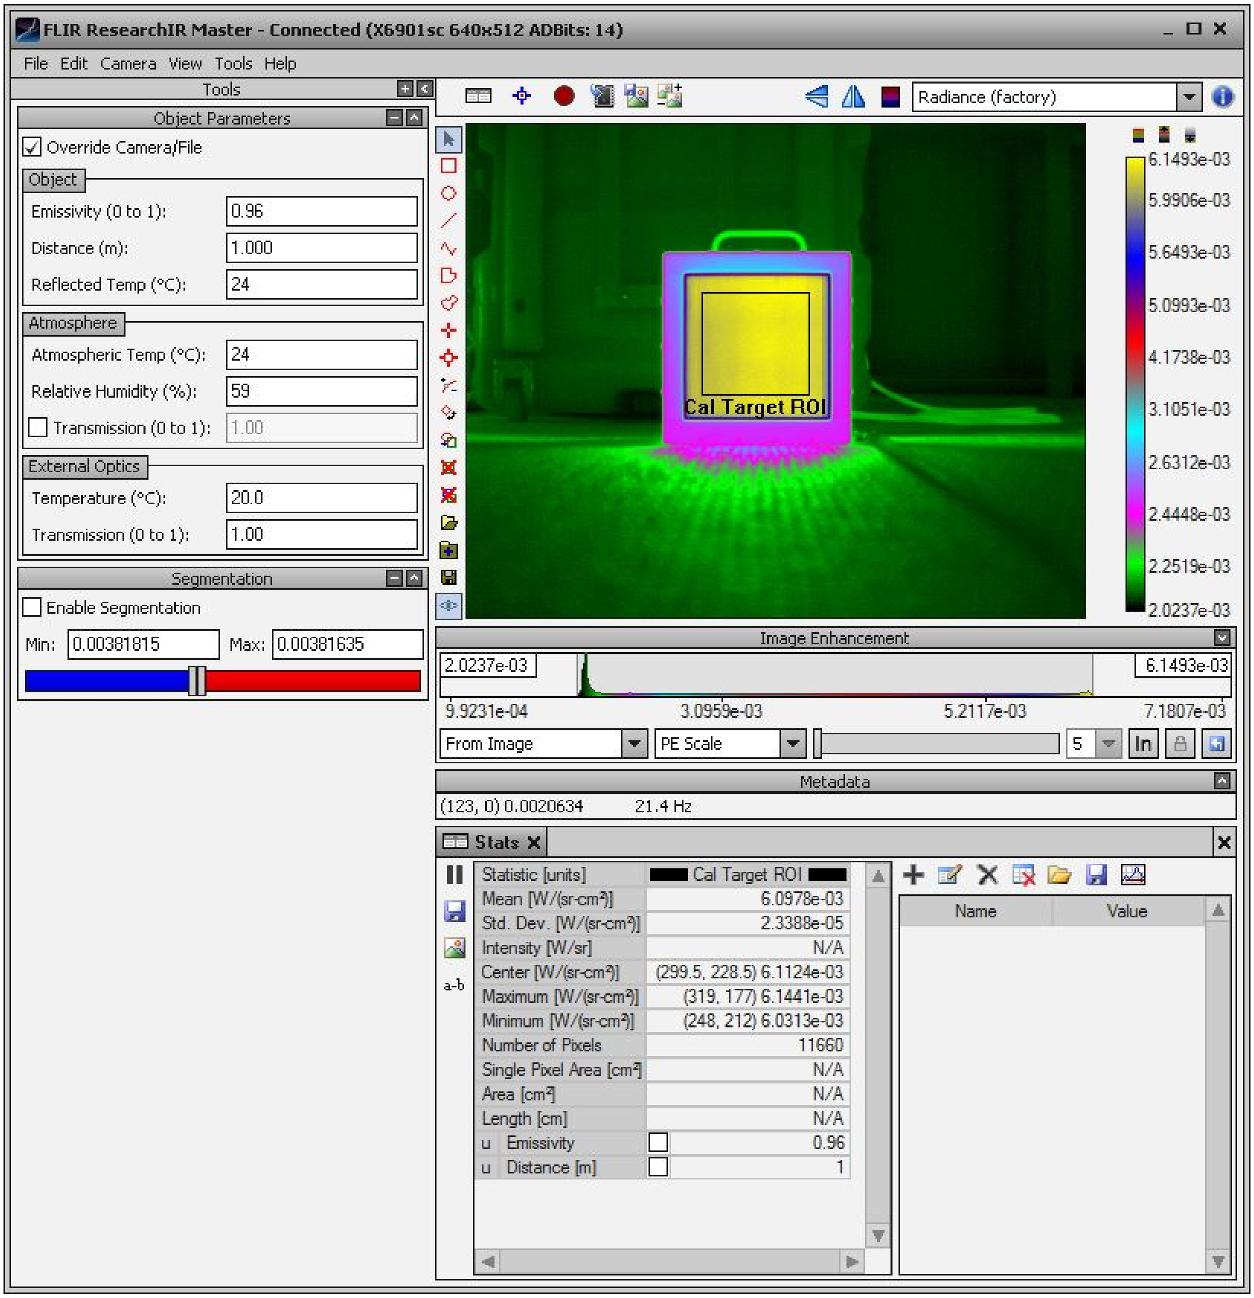


**Figure 9: The radiance measurement is the mean in-band radiance of the target in Watts/sq. cm/steradian, written as W/(sr-cm$^2$).**

In Figure 9, the Tdrift correction has been applied, but the results are not sensitive to the emissivity setting or the reflected temperature settings in the Object Parameters section of the Tools ribbon. Radiance is what you get in the pixel, and the emissivity and reflected temperature parameters are used to determine an apparent (or radiation) temperature that will hopefully be a good match to the actual physical surface temperature of a solid material.

However, the radiance measurement value is sensitive to the Object Parameters settings of the distance, atmospheric temperature and relative humidity, because those three parameters are used by ResearchIR to determine the atmospheric transmission and the air path radiance. The radiance values in the Stats table (mean, min, max) are the apparent radiance values corrected for atmospheric transmission and air path radiance. If the distance is set to something very long, like 1000m, and the air is hot and humid, then the computed atmospheric transmission can get very small, down to 0.4, and the air path radiance itself can become substantial. If the air temperature is set to something very high, like 200 °C, the radiance values can go off scale for the calibration. Care must be taken to enter in the correct values and there is definitely a limit to the Object Parameters with distance and temperature units! Emissivity and external optics transmission parameters should be limited by the user to the 0-1 range.[24]

The External Optics transmission setting will also affect the indicated radiance, as long as the transmission is set to something other than its default value of 1. If the transmission value is set to something other than 1, then the radiance will be scaled by dividing by that transmission value. Example: a radiance of X is divided by 0.9 which is the External Optics transmission. The reported value is 1.11X.

Note that the External Optics temperature setting will NOT affect the radiance measurements for a transmission value less than 1, only the temperature measurements.

## Radiance measurements are sensitive to spectral response

A critical thing to note is that the factory calibrations are always done with a spectral response file loaded in, rather than a simplified "tophat" response, as is the case for the vast majority of user calibrations. The radiance values are sensitive to the choice of spectral response. User calibrations that use a tophat response will yield very different radiance values for a given target temperature from the case where a spectral response file is used. We can supply customers with a spectral response file that is representative of the camera type they have. For an additional cost, we can measure the spectral response of their exact camera with their lens on it, and then give the customer a custom spectral response file for the particular lens they used. This is particularly desirable for custom cold filter cameras.

Note that temperature measurements are NOT very sensitive to the spectral response used in the calibration. The radiance can be thought of as simply a step along the way for ResearchIR to compute temperature, and since the spectral response is the same to convert from calibration source temperature to radiance, and then from radiance back to apparent temperature, you get the same temperature values in the end regardless of the spectral response used (within limits of course, it needs to be close to the right response). The best reason to not use tophat responses is they don't give as perfect a linear fit as when the spectral response file is used in the calibration, although the differences in the goodness of fit can be minute if the tophat limits are selected

[24] These parameters are not limited in software at the present time, though they should be.

correctly. We use the 50% points of the curve.  For example, Figure 10 shows the relative spectral response of the SLS camera used for the experiments described in this document.  The 50% points are 7.7µm and 10.4µm.  I can use either the curve itself, or the 50% points and the $R^2$ values are both 1.0000; it is really the same goodness of fit to a very high degree of precision.

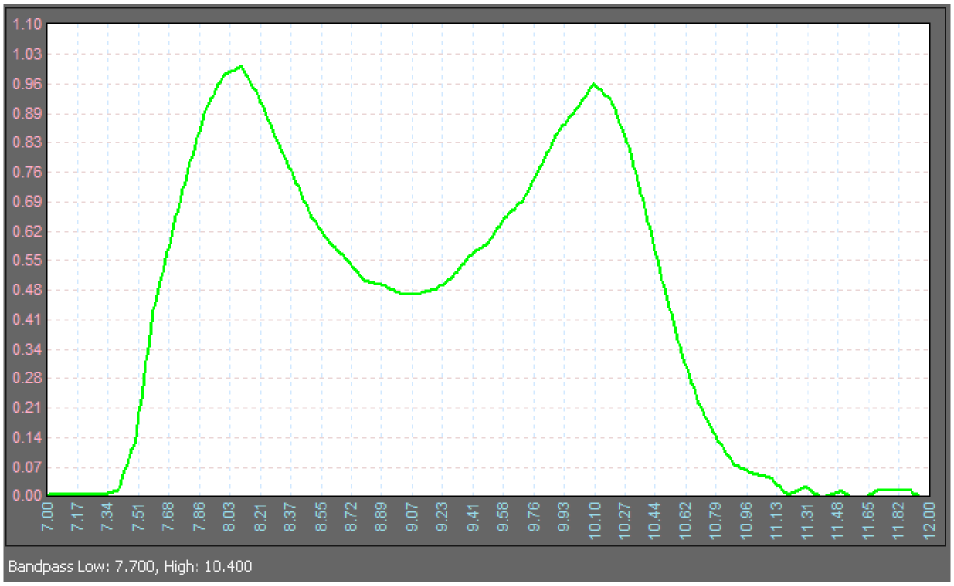


**Figure 10: Measured spectral response of Niceville SLS longwave IR camera**

## Using Two Different Ranges of Factory Calibrations to Measure the Same 100C Target

The X6900 SLS camera has three factory calibrations loaded into it, as shown in Figure 11 below.  Preset 3 has a factory NUC loaded, but there is no factory radiometric calibration, since the only detail in the name is the integration time.  When there is a radiometric calibration loaded, the name contains a temperature range.

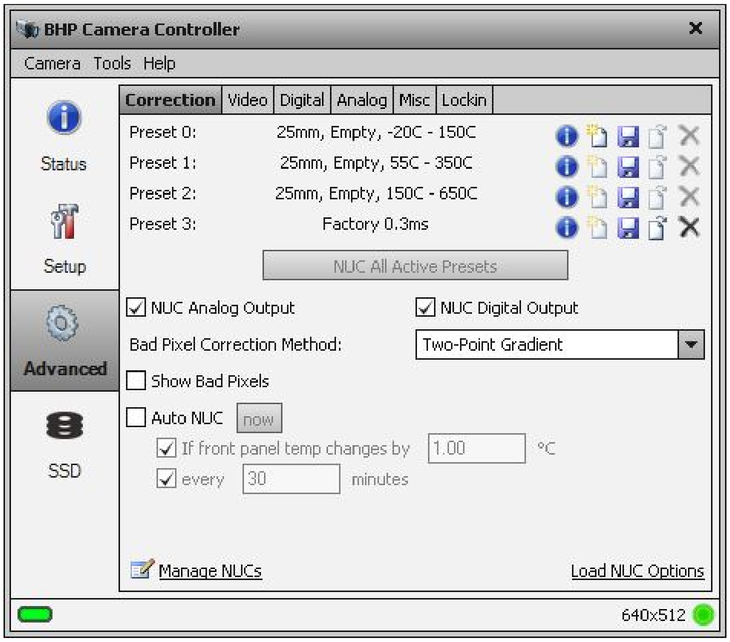


**Figure 11: Active presets on camera controller Advanced/Correction tab**

We just saw the results of using the -20 to 150 °C preset to measure the blackbody. The blackbody target is also within the range of the 55 to 350 °C preset, so I switched to that preset to see if the blackbody still measures out at ~100 °C and ~6.1e-3 W/sq. cm/sr. There ought to be good agreement, or there is a problem, since we are looking at the same target at the same temperature! Let's compare the results from the two presets, as shown in Figures 12 and 13 below. Note that the two different preset datasets were taken within 30 minutes of each other, and the camera and the blackbody drifted a little during that time interval, but as you can see in Table 1, the agreement between the temperature and radiance values is excellent. Note that the mean digital counts are quite different, as we would expect, given that the -20 to 150 °C preset integration time was 0.096msec, while the 55 to 350 °C preset integration time was 0.028 msec.

| Preset Calibration Range | Integration time (msec) | ROI Mean Digital Counts | ROI Mean Temperature (Celsius) | ROI Mean Radiance (Watts/sq.cm/sr) |
|---|---|---|---|---|
| -20C to 150 °C | 0.096 | 8438.6 | 98.1 | 6.0978e-3 |
| 55C to 350 °C | 0.028 | 2934.0 | 98.3 | 6.0807e-3 |

**Table 1**

The two digital counts images in Figures 12 and 13 look very similar, with slightly more noise apparent in the 55-350 °C preset image, as one would expect, since the sensitivity is much lower (~4X lower, or the ratio of integration times) in that preset.

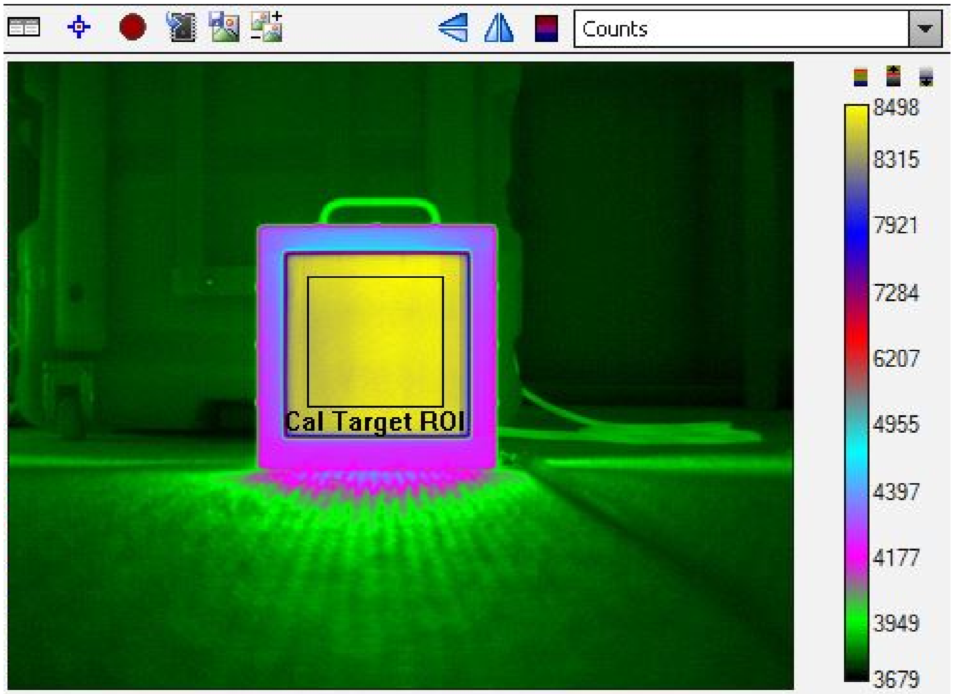


**Figure 12. An image made with a -20-150 °C Preset of a 100 °C blackbody in digital counts. The blackbody ROI mean is around 8500 counts, or about midscale for the 14-bit camera, which has a linear range up to 14,500 counts.**

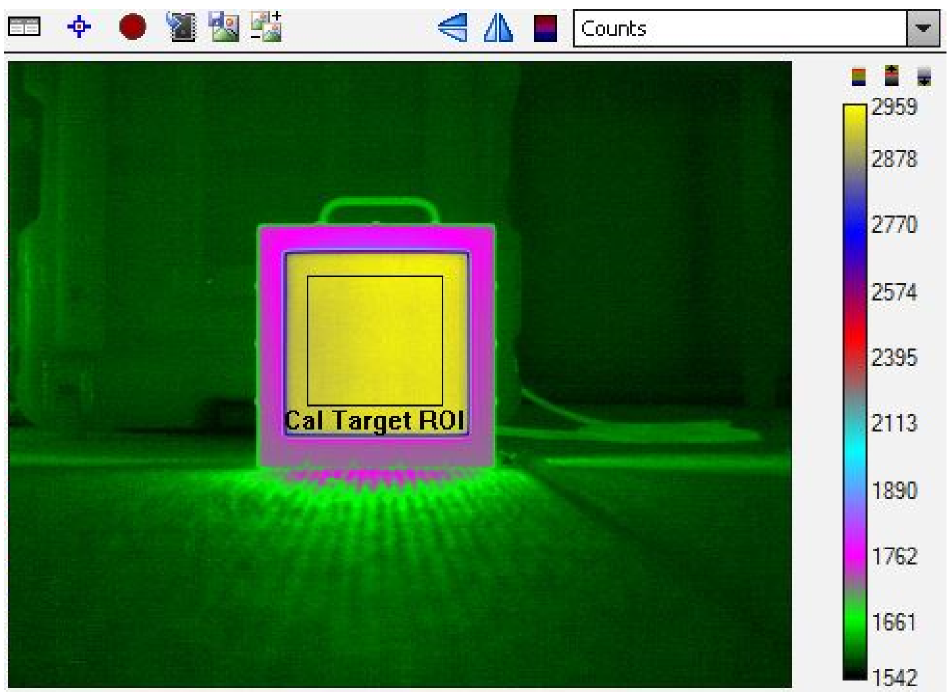


**Figure 13. 55-350 °C Preset image of 100 °C blackbody in digital counts. The blackbody is around 2960 counts, in the lower part of the scale for the 14-bit camera which has a linear range up to 14,500 counts.**

Figure 14 shows two temperature images of the 100 °C blackbody taken with the two different presets:

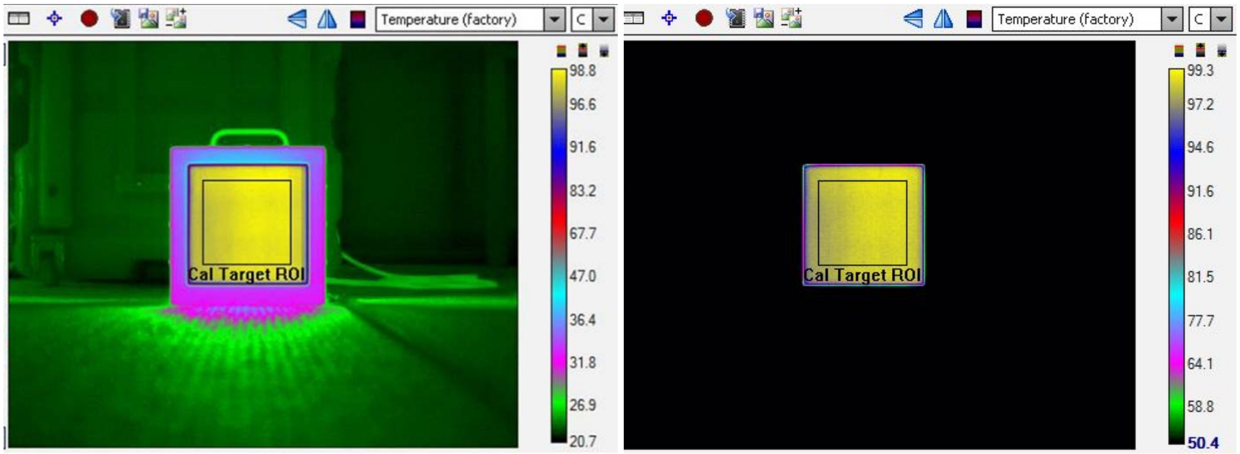


**Figure 14. -20-150 °C (left) and 55-350 °C (right) preset images of the 100 °C blackbody in temperature units. The ~24 °C ambient temperature room and carpet parts of the scene are below the lower temperature range limit of 50.4 °C on the righthand image, so those pixels are mapped to black by ResearchIR.**

Figure 15 shows images of the blackbody expressed in in-band radiance units for the two presets:

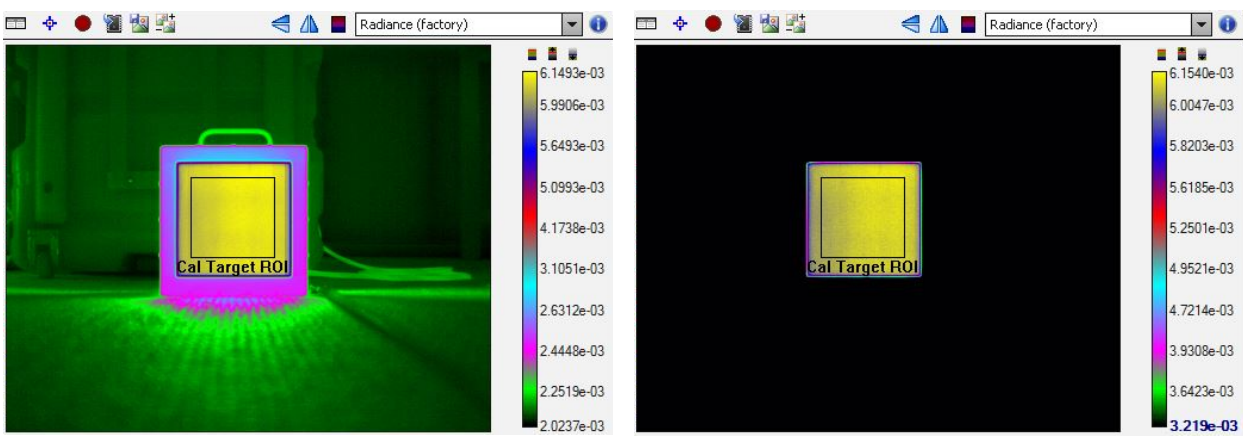


**Figure 15. -20-150 °C (left) and 55-350 °C (right) Preset images of the 100 °C blackbody in radiance units. The ambient background is below the lower radiance range limit of 3.219e-3 W/sq.cm/sr, so those pixels are mapped to black by ResearchIR.**

## Comparison of Factory Calibration to User Calibration

By now, I was convinced that the camera I was using had good factory calibrations that gave measurements that agreed with the experimental setup conditions and blackbody temperature. I then performed a user calibration on the camera using the same integration time (0.095msec) as the -20-150 °C factory calibration, and I used the factory NUC at that same integration time. Two different spectral responses were used for the user calibration to see the differences they yield in temperature and radiance measurement on the blackbody target at several temperatures.

I compared the results, as shown below. The step-by-step user calibration process is all detailed in Appendix 1 of this book.

| Calibration Type | BB Target Temperature (Celsius) | Calibration Range | Target Apparent Temperature (Celsius) | Target Radiance (ATM = 0.99) (W/sq. cm/sr) |
|---|---|---|---|---|
| Factory | 76 | -20 to 150 °C | 74.0 | 4.5369e-3 |
| User (Tophat 7.5-10.5µm) | 76 | -20 to 150 °C | 77.4 | 6.0284e-3 |
| User (*.prn File) | 76 | -20 to 150 °C | 77.4 | 4.3271e-3 |
| Factory | 111 | -20 to 150 °C | 108.9 | 6.8804e-3 |
| User (Tophat 7.5-10.5µm) | 111 | -20 to 150 °C | 112.2 | 9.2056e-3 |
| User (*.prn File) | 111 | -20 to 150 °C | 112.2 | 6.5862e-3 |
| Factory | 29 | -20 to 150 °C | 29.6 | 2.2787e-3 |
| User (Tophat 7.5-10.5µm) | 29 | -20 to 150 °C | 30.3 | 2.9335e-3 |
| User (*.prn File) | 29 | -20 to 150 °C | 30.0 | 2.1097e-3 |

**Table 2**

The factory and user calibrations both give decent results for temperature measurement that match the indicated temperature on the blackbody controller to within ~2 °C. I was somewhat limited with the testing in that I did not and do not know the exact in-band emissivity of the ISDC Model 2106 blackbody used for the testing. The factory specification for this blackbody says 0.96 +/-0.02 without specifying a spectral band.

The Factory calibration was done in Niceville using cavity blackbodies which have an emissivity of 0.99, due mostly to geometry, rather than an emissive coating. The user calibrations were done using the *same* area blackbody used for the subsequent testing, which perhaps is a bit of a circular argument, because by doing that, the emissivity is “baked in” during calibration, and the test results are insensitive to the emissivity value assigned in the calibration, as long as the same value is used in the ROI measurements.

The tophat radiance value is quite different from the factory or *.prn calibration values. This makes sense, because the tophat shape approximates a spectral response that is not actually very close to a tophat shape, due in part to the spikey nature of the SLS spectral response. Figure 16 shows the tophat response – a pure square wave or top hat shape with hard cut-on and cut-off wavelengths of 7.5 and 10.5µm. Those two limits are frequently cited as the spectral band of the SLS cameras in our marketing literature.

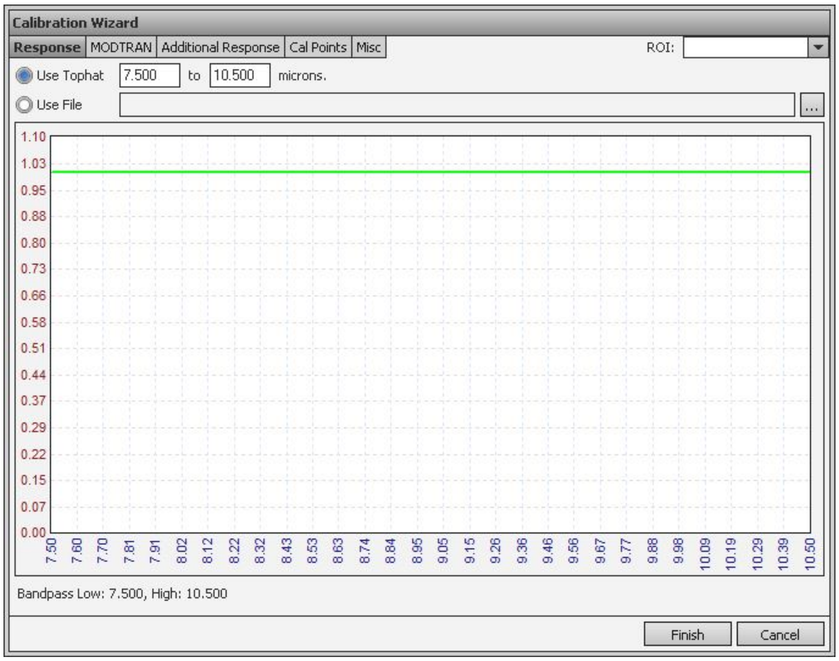


**Figure 16. Tophat spectral response for Niceville SLS longwave IR camera**

Here is the spectral response of an SLS camera. Note the big dip in the center, which definitely affects the resulting in-band radiance one gets for a given target temperature.

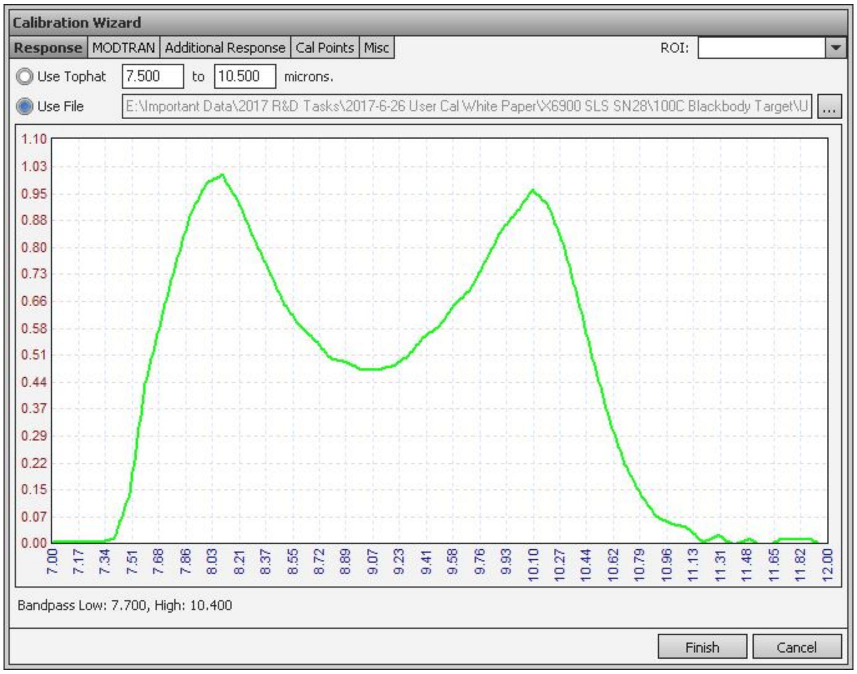


**Figure 17. Measured spectral response of Niceville SLS longwave IR camera**

The factory calibration data shown above had Object Parameters set such that the atmospheric transmission was 0.99. This 0.99 value was computed by the software by me entering into Object Parameters an atmospheric temperature of 24 °C, a relative humidity of 59% and a 1-meter distance or path length. The reflected temperature was set to 24 °C also, since the room was heated to 24 °C. See Figure 18 below for a screenshot of the settings used:

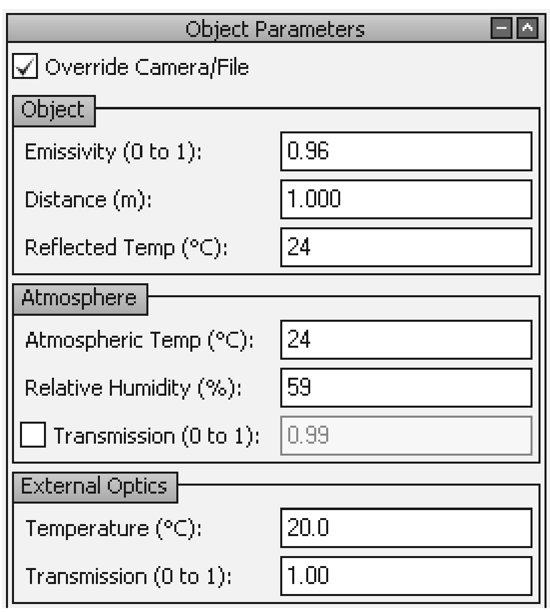


**Figure 18. Object parameters used for measurements in Table 2.**

# Chapter 11 : Changing Include Files with REdit

Introduction: ........................................................................................................................ 11—420

## Introduction:

Sometimes a customer collects data that is not radiometrically calibrated, but then they want to go back and calibrate it after the fact (we call that post calibration). Or they collected data with an active calibration file pair (include or *.inc file and calibration or *.cal file) that was generated by a laboratory calibration process done with a different lens than the one they used for the collection. Either way, they want to go back in and change the radiometric calibration parameters.

If the image data files were acquired in SFMOV format, this is easy.  If the data files are collected in .ATS format, then this is not easy and they have to be laboriously converted to SFMOV format.  That is why I never use ATS files unless I am using a factory calibration and don't ever want to change the calibration.

When one acquires data in SFMOV format, there are at least two other associated files that are created: the *.pod and the *.inc file.  POD stands for parametric object data and INC stands for included data.  The *.inc file contains the radiometric coefficients needed to convert the raw counts pixel data in the SFMOV files into radiance units.  The *.inc file also has the polynomial fit coefficients needed to convert from radiance to apparent temperature.

If the original data files were done with no user calibration in place, then there won't be a *.cal file created.  If there was a user cal in place at the time of acquisition, then you will get a *.cal file as well. This *.cal file is important and it "lives in the shadows" in that it cannot be seen as a separate tag in the REdit utility which is described below.

When you select a new include file in REdit, the associated *.cal file better be in the same directory as the SFMOV and the new include file or you will get problems if you try to measure temperatures in the image, and worst of all, there is no warning!!!

Here is a screen shot of an include file that was generated by a user calibration.  You can see that it has a tag in it that mentions the associated *.cal file. The tag is "CAFile".  The *.cal file has the same name as the include file.

```
HdSize auto
Group Instrument Info
IHFOV 0.000000e+00
IVFOV 0.000000e+00
SBPLo 3.000000e+00
SBPUp 4.900000e+00
SltRng 1.0
Group EUD Info
BgType_0 none
BgValu_0 0
Coeff0 -2.488694e-04
Coeff0_0 -2.488694e-04
Coeff1 9.254181e-08
Coeff1_0 9.254181e-08
DaUnit_0 W/(sr-cm^2)
EuRaw_0 Raw
PolyOrder_0 1
StdUnt_0 17
TempCoeff0_0 -3.276284e+01
TempCoeff1_0 8.251671e+05
TempCoeff2_0 -5.394117e+09
TempCoeff3_0 2.406512e+13
TempCoeff4_0 -6.291651e+16
TempCoeff5_0 8.710015e+19
TempCoeff6_0 -4.915163e+22
TempPolyOrder_0 6
TPFact 1.0
TPFact_0 1.0
Group Referenced Files
BgFile_0
CAFile 2msec 25mm Lens Radio Cal.cal
Group Misc Tags
TempCoeff0 -3.276284e+01
TempCoeff1 8.251671e+05
TempCoeff2 -5.394117e+09
TempCoeff3 2.406512e+13
TempCoeff4 -6.291651e+16
TempCoeff5 8.710015e+19
TempCoeff6 -4.915163e+22
DATA
```

The most important tags in the include file are the Coeff0 and Coeff1. These are the radiometric transfer coefficients that convert digital counts into radiance units. The units are defined by the tag DaUnit_0 which is written out as W/(sr-cm^2), or watts per steradian per square centimeter. The other tags are described in the ResearchIR user guide. If one has collected data with a calibration that was done with the wrong lens, for example, then the radiometric transfer coefficients will be incorrect. Switching to a new include file that was created with the correct lens will mean that there will be different coefficients there.

If the camera is radiometrically uncalibrated, then you will get an include file that should look something like this – note that it is quite devoid of tags:

```
HdSize auto
Group Misc Tags
FRate_0 124.25
ITime_0 6
xmrCameraName RS8313
Group Instrument Info
IHFOV 72.0906
IVFOV 72.0906
DATA
```

To change the refererenced include file, launch the REdit software which is located here:

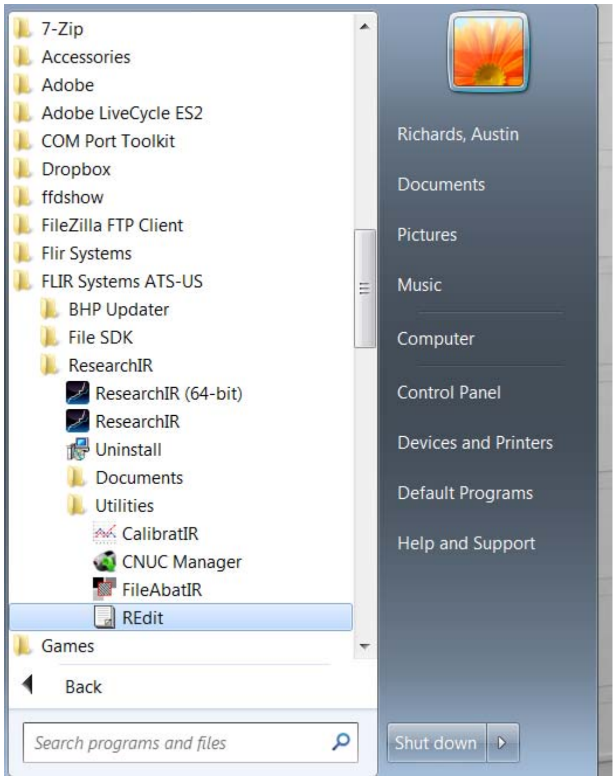


You will get the REdit main window:

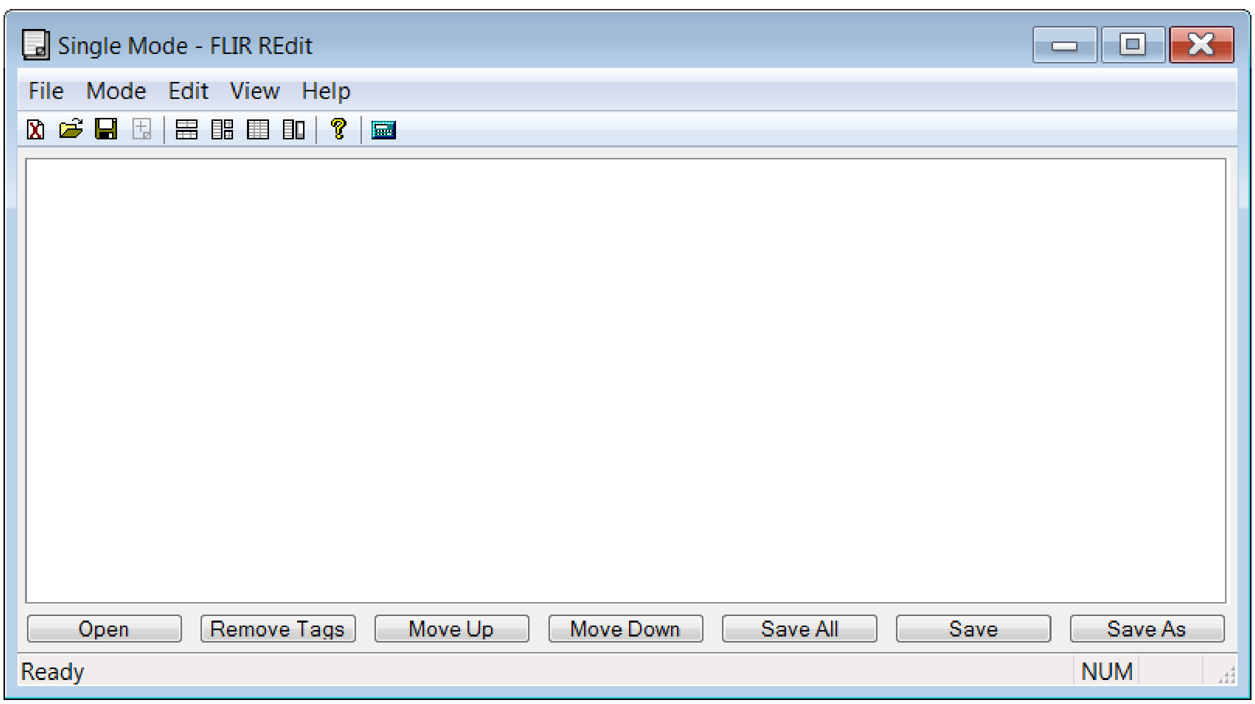

Open the SFMOV file you want to edit, but first, put the new include file and *.cal file in the same directory as the SFMOV. Then when you open the SFMOV in ResearchIR, ResearchIR looks in the same directory for the two files to properly convert raw counts to both radiance and temperature. If you leave out the *.cal file, radiance will still work but not temperature.

Here is an SFMOV file that is calling an include file with the same filename as the SFMOV. This SFMOV file is one of three presets that were acquired while the camera was in superframing mode which is why it has a -PS0 suffix on the file name.

There will always be an include file created when the movie file is acquired and it will have the same name. You will switch this out for a new include file.

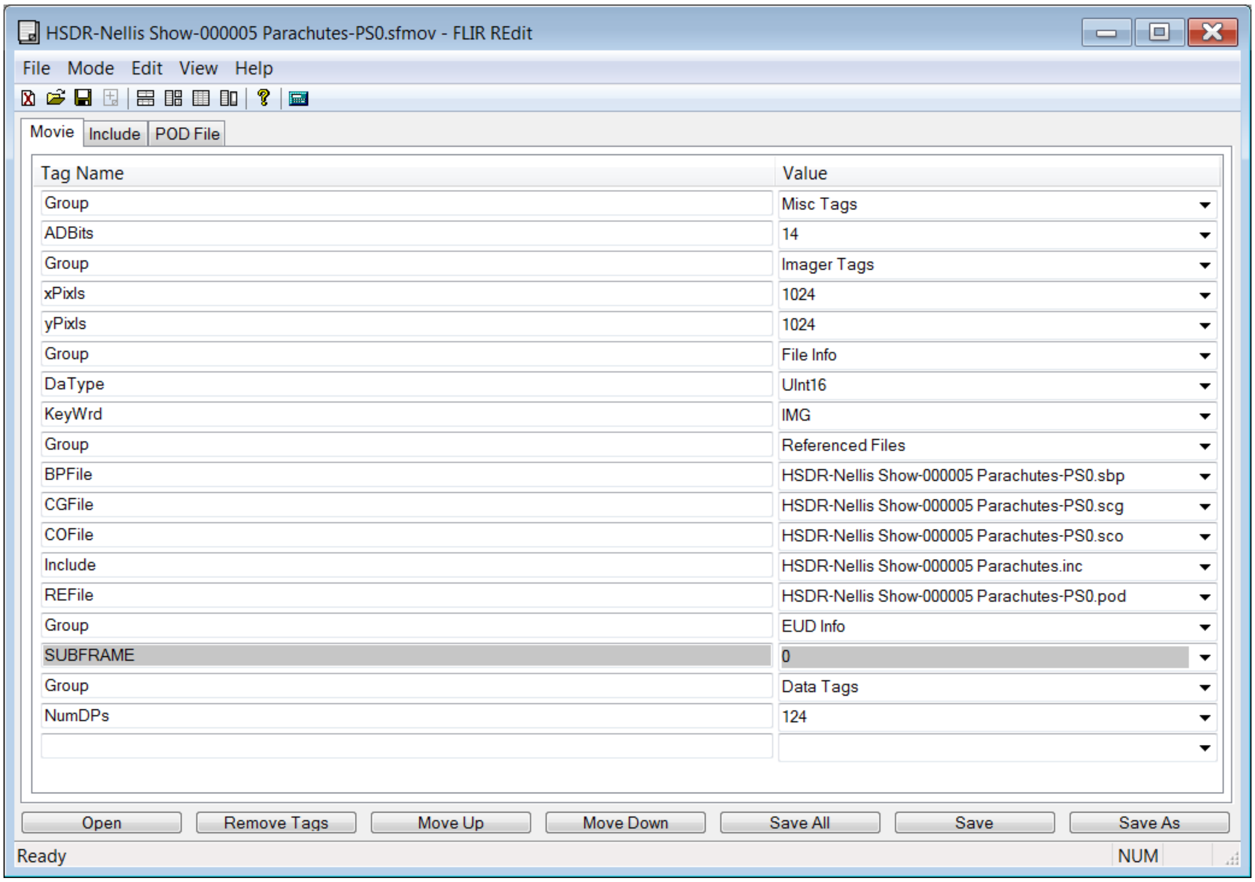


Here is a screen shot of the original include file which is called HSDR-Nellis Show-000005 Parachutes.inc. It does not have any radiometry tags in it. There are very few tags in it at all, in fact.

```
HdSize auto
Group Misc Tags
FRate_0 125.409
FRate_1 125.409
FRate_2 125.409
ITime_0 1
ITime_1 0.1
ITime_2 0.01
xmrCameraName SC8100
xmrSubframe true
DATA
```

After the data collection at Nellis Airforce Base, I went back home and did a laboratory radiometric calibration back in Santa Barbara on the camera with all the same settings and lenses, etc. Here is the new include file I generated which was for three presets. It is called 100mm Lens 1000 to 10 usec.inc and it was created at the same time as a *.cal file with the same name (the file is mentioned in the second line of the include file as the CaFile tag). When one generates a user calibration, the software creates a *.cal file which can be read by CalibratIR software or by ResearchIR, as well as a text file include file with the same name.

```
HDSIZE Auto
CaFile 100mm Lens 1000 to 10 usec.cal
SBPLo 3.000000
SBPUp 5.000000
TPFact_0 1.000000
TPFact_1 1.000000
TPFact_2 1.000000
SltRng 1000.000000
IHFOV 180.000000
IVFOV 180.000000
EuRAW_0 Raw
StdUnt_0 17
DaUnit_0 W/(sr-cm^2)
BgType_0 none
BgValu_0 0.000000e+000
BgFile_0
EuRAW_1 Raw
StdUnt_1 17
DaUnit_1 W/(sr-cm^2)
BgType_1 none
BgValu_1 0.000000e+000
BgFile_1
EuRAW_2 Raw
StdUnt_2 17
DaUnit_2 W/(sr-cm^2)
BgType_2 none
BgValu_2 0.000000e+000
BgFile_2
PolyOrder_0 1
Coeff0_0 -0.000138578
Coeff1_0 1.47644e-007
TempPolyOrder_0 6
TempCoeff0_0 45.4733
TempCoeff1_0 22797.8
TempCoeff2_0 -957413
TempCoeff3_0 2.08587e+007
TempCoeff4_0 -2.29423e+008
TempCoeff5_0 1.22483e+009
TempCoeff6_0 -2.5233e+009
PolyOrder_1 1
Coeff0_1 -0.000638071
Coeff1_1 1.53094e-006
TempPolyOrder_1 6
TempCoeff0_1 45.4733
TempCoeff1_1 22797.8
TempCoeff2_1 -957413
TempCoeff3_1 2.08587e+007
TempCoeff4_1 -2.29423e+008
TempCoeff5_1 1.22483e+009
TempCoeff6_1 -2.5233e+009
PolyOrder_2 1
Coeff0_2 -0.00396617
Coeff1_2 1.07939e-005
TempPolyOrder_2 6
TempCoeff0_2 45.4733
TempCoeff1_2 22797.8
TempCoeff2_2 -957413
TempCoeff3_2 2.08587e+007
TempCoeff4_2 -2.29423e+008
TempCoeff5_2 1.22483e+009
TempCoeff6_2 -2.5233e+009
DATA
```

Here I switch the old include file for the new one using REdit and the pulldown on the right which opens a Window file browser:

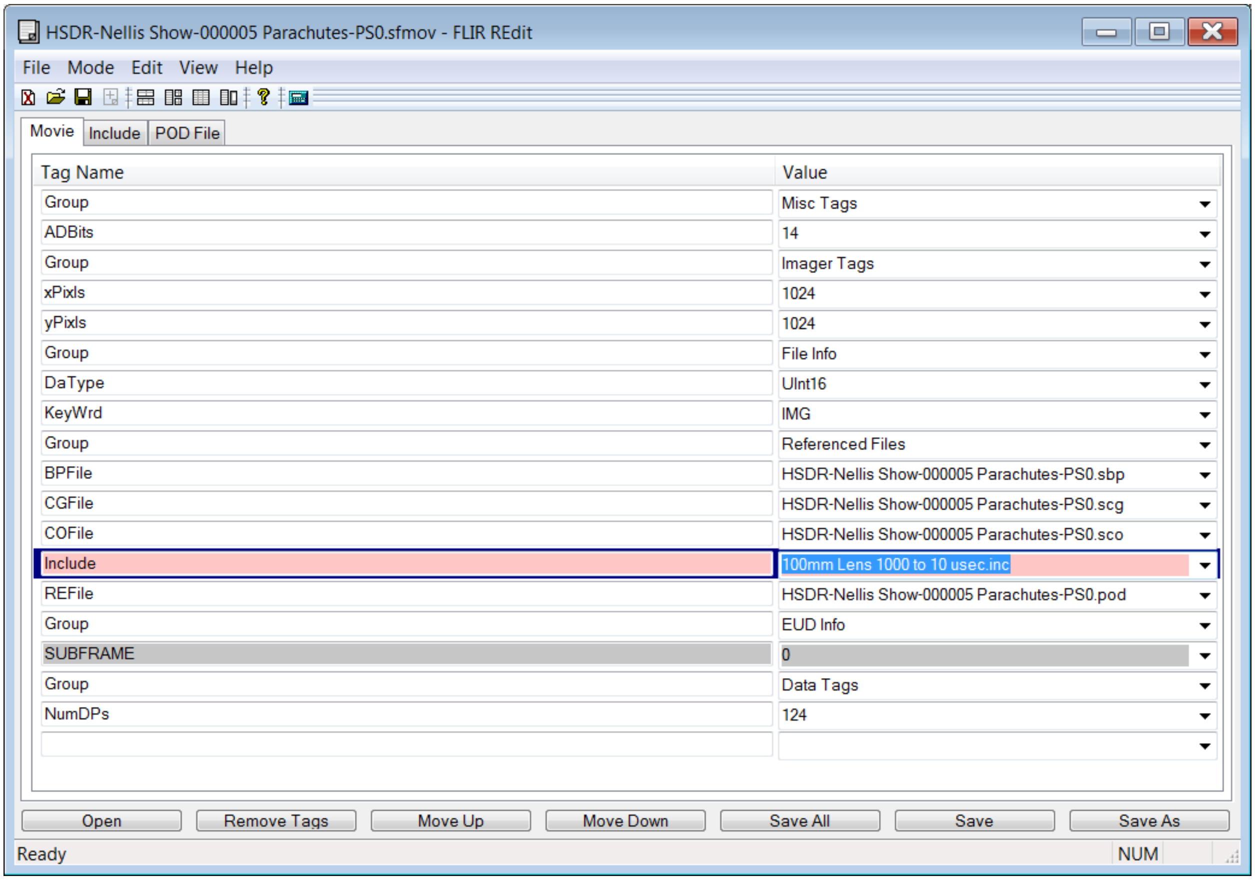


Now I hit Save and it will create a new SFMOV file that has the include file pointer directing it to the new include file 100mm Lens 1000 to 10 usec.inc. It will also create a .bak file which is the original unsullied file with the original include file referenced.

**If the new include file is not in the same directory as the SFMOV, then you will get an error message like this:**

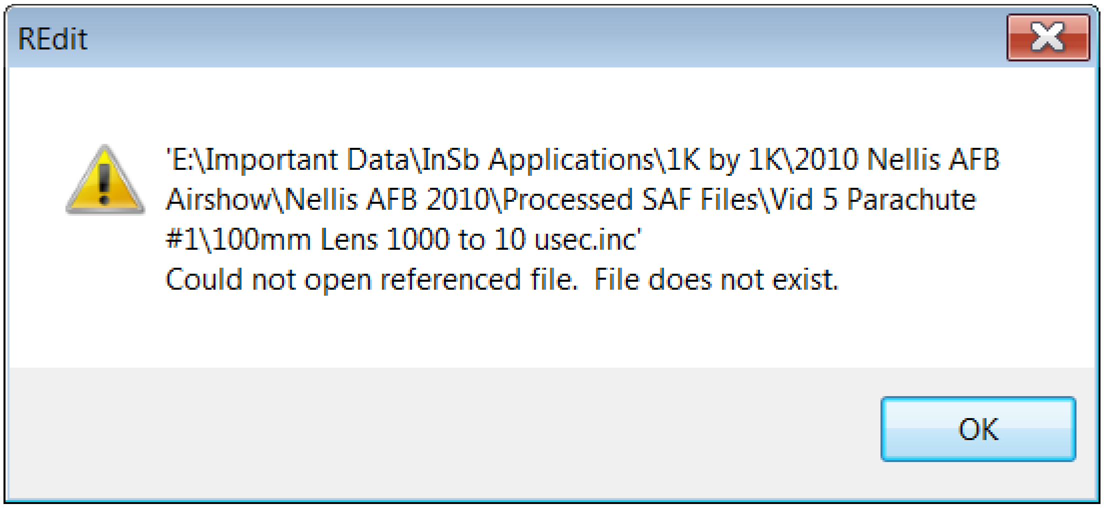


**NOTE: You need to make sure the *.cal file is also in the same directory! If you don't, then when you go to look at images in temperature units, the images won't have meaningful temperature data! But there is no warning that there is a problem and radiance images still work. You just get grey temperature images like this:**

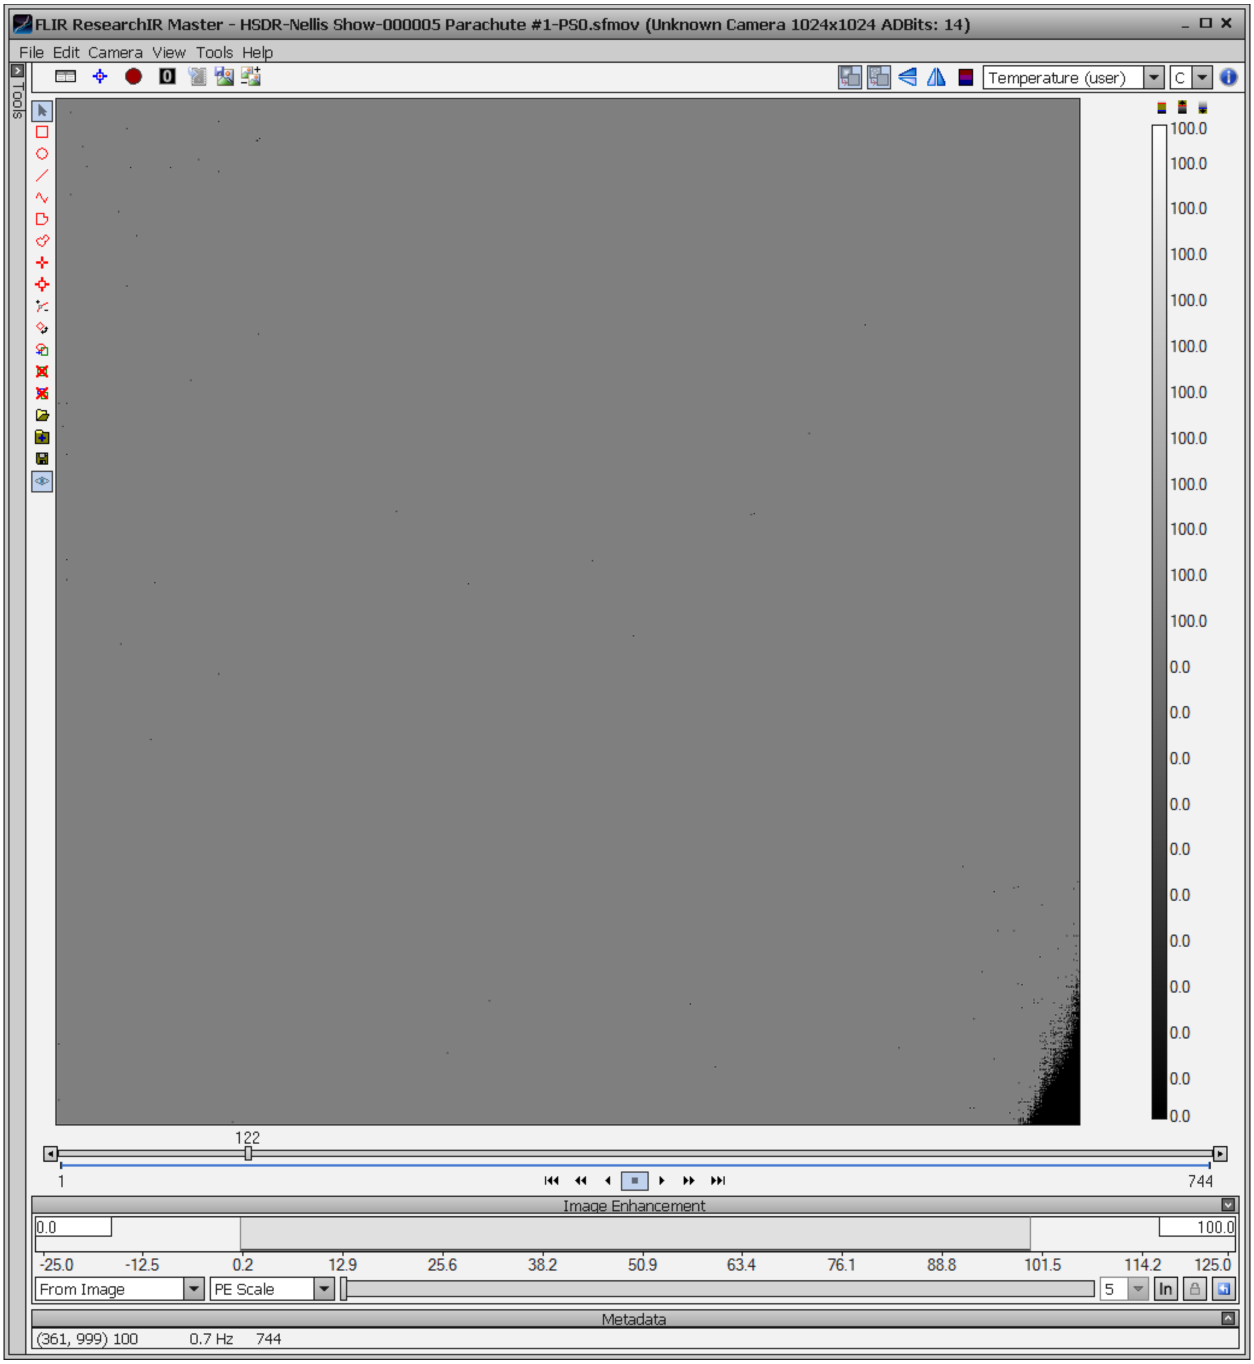


Here I put the *.cal and the include file in the same Windows directory as the SFMOV and all is well again. I can measure a temperature of 25C on the leg of the skydiver.

Moral of the story: do not delete any files associated with the SFMOV and make sure they are all carried around. That is why people like the ATS format, because everything is in one file which is acting like a wrapper for all the discrete files inside.

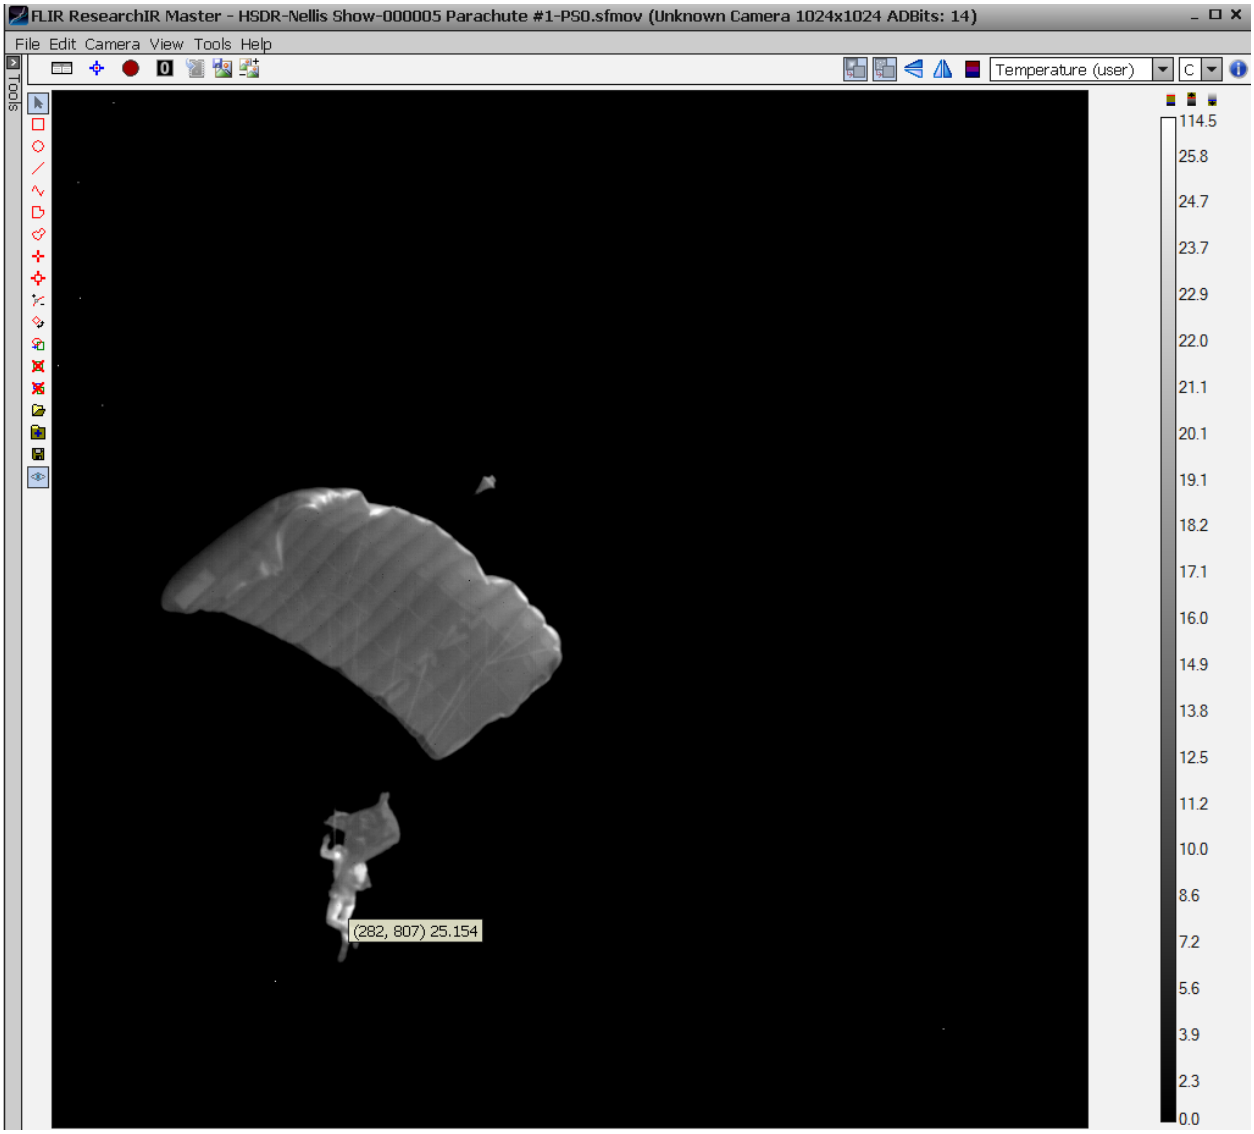
FLIR ResearchIR Master - HSDR-Nellis Show-000005 Parachute #1-PS0.sfmov (Unknown Camera 1024x1024 ADBits: 14)
File Edit Camera View Tools Help
Temperature (user)
C
114.5
25.8
24.7
23.7
22.9
22.0
21.1
20.1
19.1
18.2
17.1
16.0
14.9
13.8
12.5
11.2
10.0
8.6
7.2
5.6
3.9
2.3
0.0
(282, 807) 25.154

# Chapter 12 : Generating NUC Tables

Introduction ........................................................................................................................ 12—431

Gain Table Generation ........................................................................................................ 12—432

Gain Table Results ............................................................................................................. 12—434

Offset Table ........................................................................................................................ 12—437

## Introduction

It is useful to understand how non-uniformity tables for IR cameras can be generated. Some camera users want to make their own NUCs "from scratch" using software tools like FLIR ResearchIR and MATLAB. These types of advanced users often will acquire 14-bit uncorrected image data with a PC-side NUC applied to the video. They then make better NUC tables to apply to the image data during post-processing. This document explains how one can generate NUC tables from scratch.

The equation for the application of a 2-point NUC is straightforward. Let U(i,j) be an uncorrected image, G(i,j) is the gain table, O(i,j) is the offset table. Then the corrected image C(i,j) is calculated as follows:

$$\mathbf{C(i,j) = G(i,j) * U(i,j) + O(i,j)}$$

The multiplication operator used here is an array operator that computes the element-by-element product.

We start with sequences of images taken while looking at a laboratory blackbody or other uniform source like a lightbox if the camera operates in the SWIR band. The image sequences can be 16 frames or longer. Each sequence gets averaged to reduce the temporal noise. In this example, 64 frames were captured. One can calculate G and O using two sequences or three. In the case of two sequences, the offset is calculated from the colder of the two sequences. In the case of three sequences, the gain is calculated from the cold and hot sequences and the offset from the warm sequence. The example given here used three different blackbody temperatures (20, 25 and 30C) to generate three image sequences. The NUCs generated in the factory for FLIR Niceville cameras are done with two blackbody temperatures. There does not seem to be noticeable differences between the two methods, particularly with recent production (since 2007) cameras and their focal plane arrays built on very linear ROICs.

Figure 1 is the first frame of a 64-frame image sequence taken with a 3-5 micron bandpass InSb camera that has no NUC applied to it. The blackbody was set to 25 C, the integration time was 1.6 msec and the coldshield f/number was 2.5. There was a 25mm lens on the camera.

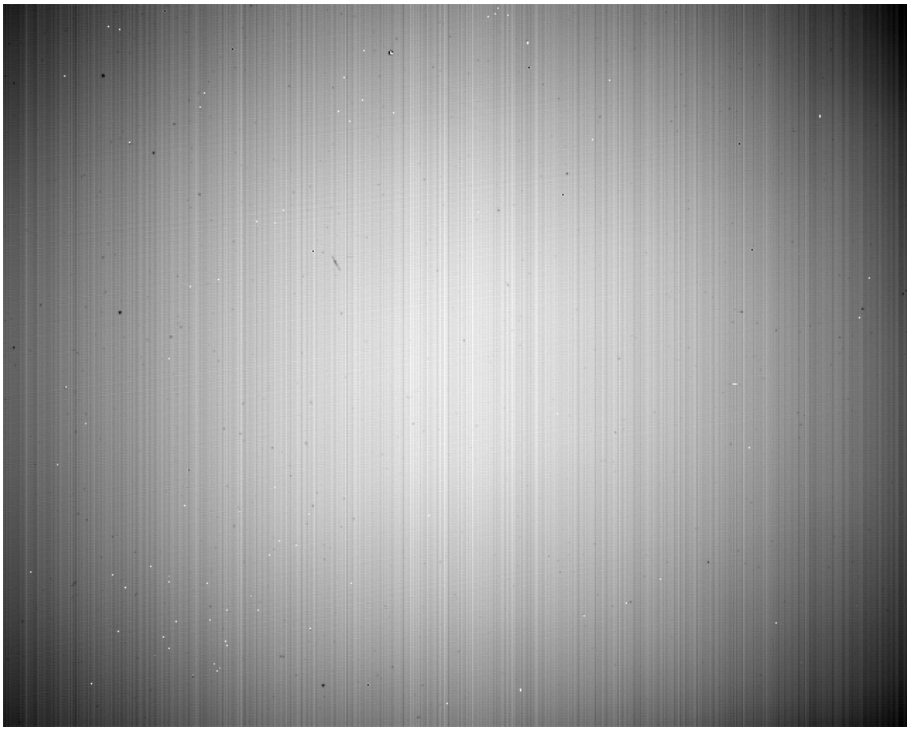

**Figure 1. An uncorrected image of a 25C uniform area blackbody source. This is the first frame of a 64-frame image sequence.**

According to the statistics tool in ResearchIR, the standard deviation of this first frame is 477.4 counts and the mean is 8798.9 counts. After the NUC is applied, that standard deviation value should be less than 10 counts.

I will now show you how a NUC is generated from the means of three image sequences and applied to this above image. We start by generating the gain table.

## Gain Table Generation

Here are the means of 64 -frame sequences of the blackbody at 20 C and at 30 C which will be used to generate the gain table:

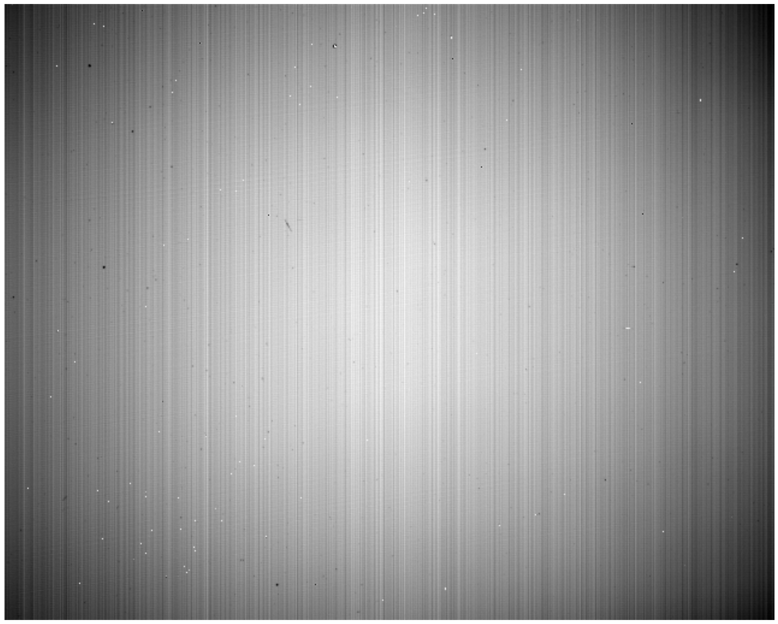

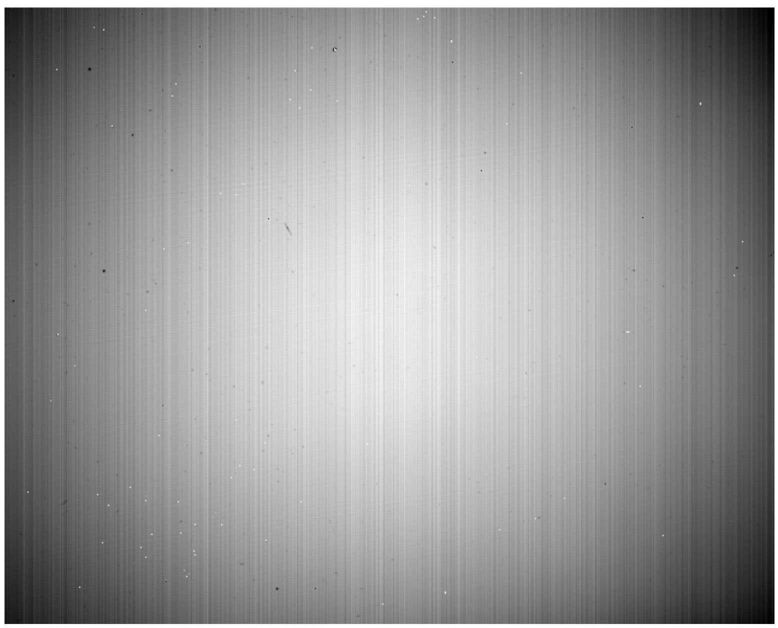

**Figure 2a-b. Frame averages of 64-frame sequences taken of 20C and 30C uniform area blackbody sources.**

The mean count values of the 20C and 30C images are 7774 and 9964 counts respectively. This is about 2200 counts of signal swing. The two images will be used to calculate the gain table. First, we take the difference between the hot and cold average frames. This is called the responsivity table and is itself a frame or image. The mean count value of this responsivity table is computed, giving a single number, which in this case is 2189.9 counts. Then the gain table is computed. It is the global mean responsivity count value divided by the responsivity of each pixel. The values will tend to be between 1.2 and 0.8. It is useful to put in some logic to prevent dividing by zero if the responsivity of a particular pixel is zero and to mark the pixel as bad in a bad pixel map, which is an image where good pixels are indicated by 1 and bad pixels indicated by 0, or some other scheme of the user's devising.

The gain image is brighter on the corners than in the center, as shown in the line profile taken across the diagonal of the gain table. When an un-NUCed image is multiplied by the gain table, the coldshield shading of the un-NUCed image that makes the corners darker and the center lighter will be compensated.

Here is a step-by-step procedure to generate the gain table:

1. Point the camera at a uniform temperature ambient scene. The calibration plate included with FLIR Niceville science cameras works well.
2. Using ResearchIR software, set the camera integration time to 0 and measure the mean count value of the resulting image using the Statistics tool. Note the value.
3. Set the integration time of the camera to the desired value for the NUC tables that are to be generated.
4. Take 64 frames of image data while looking at a uniform area blackbody source set to a temperature so that the mean digital counts are ~1000 counts above the "zero integration time" image mean.
5. Take the average of this sequence using MATLAB or Python to generate a 14-bit Tiff image referred to as C(i,j), where C stands for cold.
6. Now point the camera at a uniform area blackbody source set to a temperature that gives a mean count value that is around 12,000 digital counts. There is leeway here. Even a thousand counts of signal swing will work, but more swing is better.
7. Take 64 frames of image data and average it. This image will be referred to as H(i,j) . H stands for hot.
8. The responsivity table is computed: R(i,j) = H(i,j) – C(i,j). This is the signal swing for each pixel in response to the two uniform sources.
9. Compute M, the mean count value of the responsivity table R(i,j)
10. The Gain table G(i,j) = M/R(i,j). In cases where the pixel has zero responsivity, the pixel coordinates (i,j) should be noted and added to a bad pixel map.

## Gain Table Results

The gain table generated from this InSb camera is shown as an image in Figure 3. The image has bright corners and a dark center, which means that when the uncorrected images are multiplied by this gain image, the dark corners in the images will be "gained up" and the bright center will be "gained down", which will improve the uniformity. The corner pixel gain values are on the order of 1.2 and the center pixels gain values are on the order of 0.9.

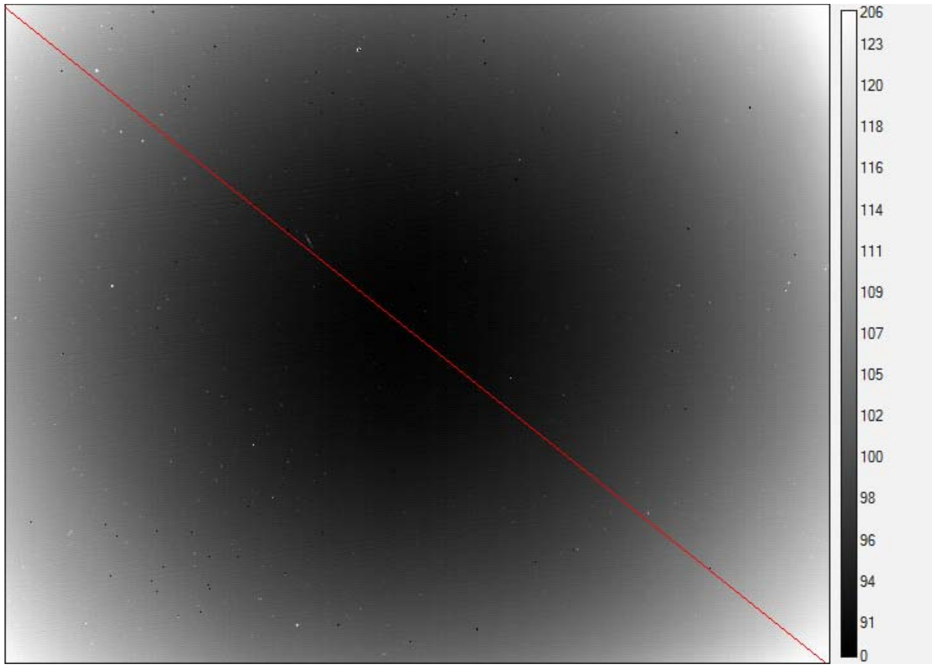


**Figure 3a. Gain table with line profile across diagonal.**

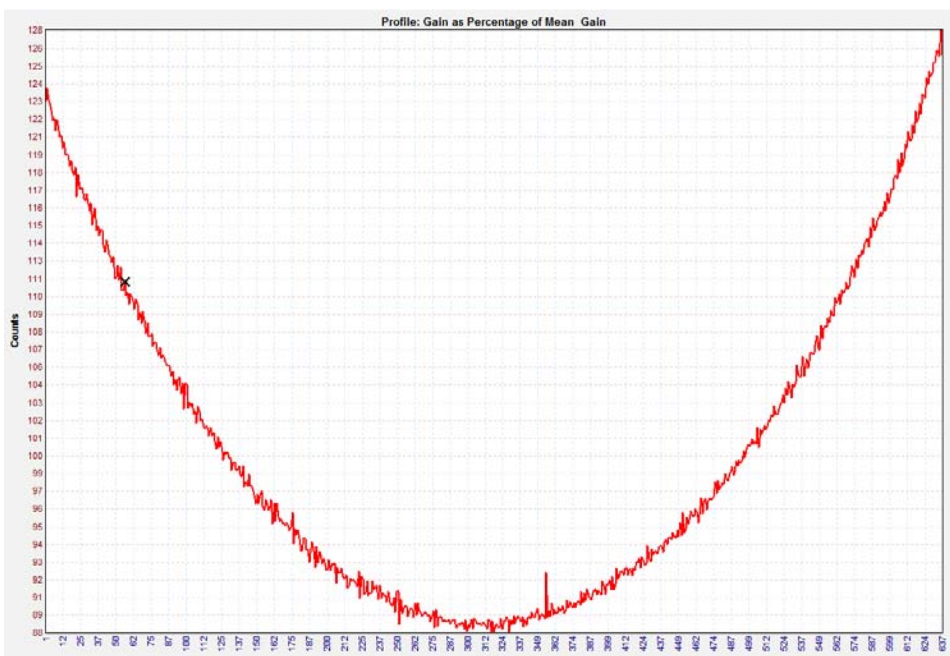


**Figure 3b. A line profile plot of the diagonal line across the gain image. Gains vary from 1.28 to 0.88.**

Here is the first frame of the 25C blackbody un-NUCed sequence. This image has darkening of the corners which is due to coldshield shading. There is also a line profile taken across the diagonal of the image:

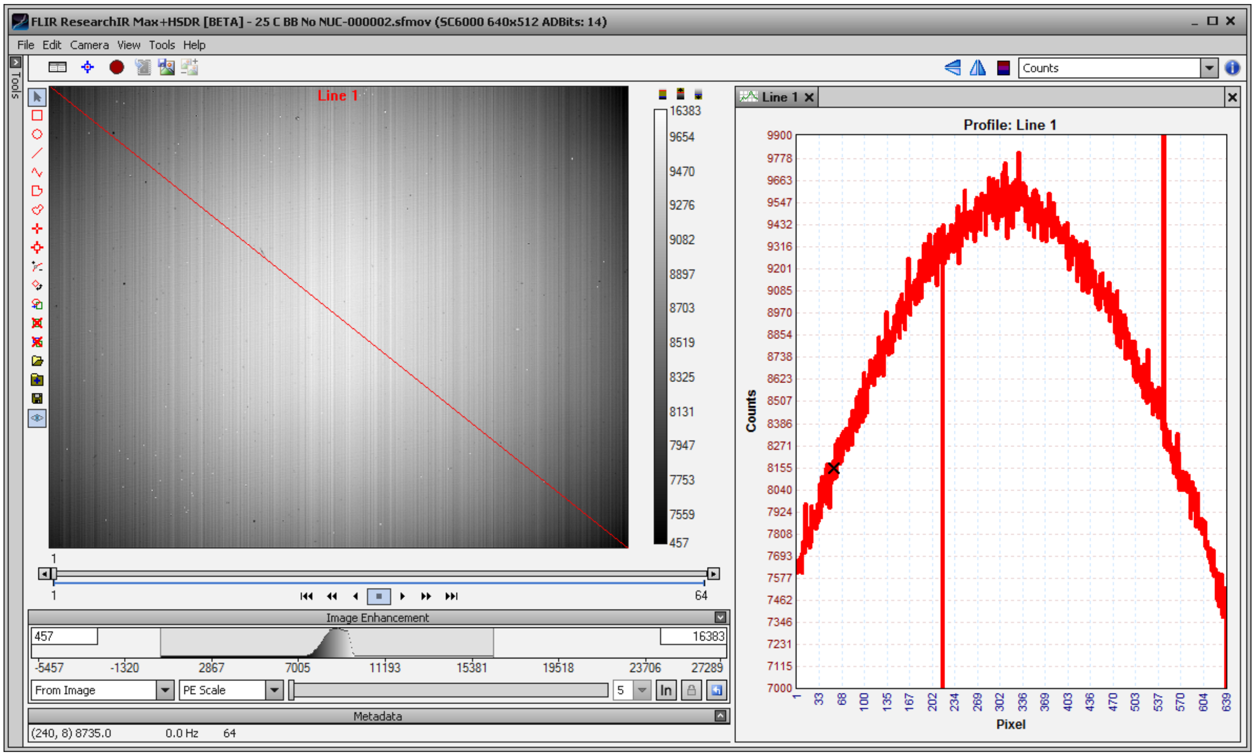


**Figure 4. Un-NUCed image of 25C blackbody**

You might imagine that multiplying this un-NUCed image by the gain image would take out most (if not all) of the “bowl” shape due to cold shield shading. It does help, but it does not remove all of it. In fact, it overcompensates for the darkened corners and makes the corners lighter. Here is the product of the gain image and the first frame of the 25 C sequence:

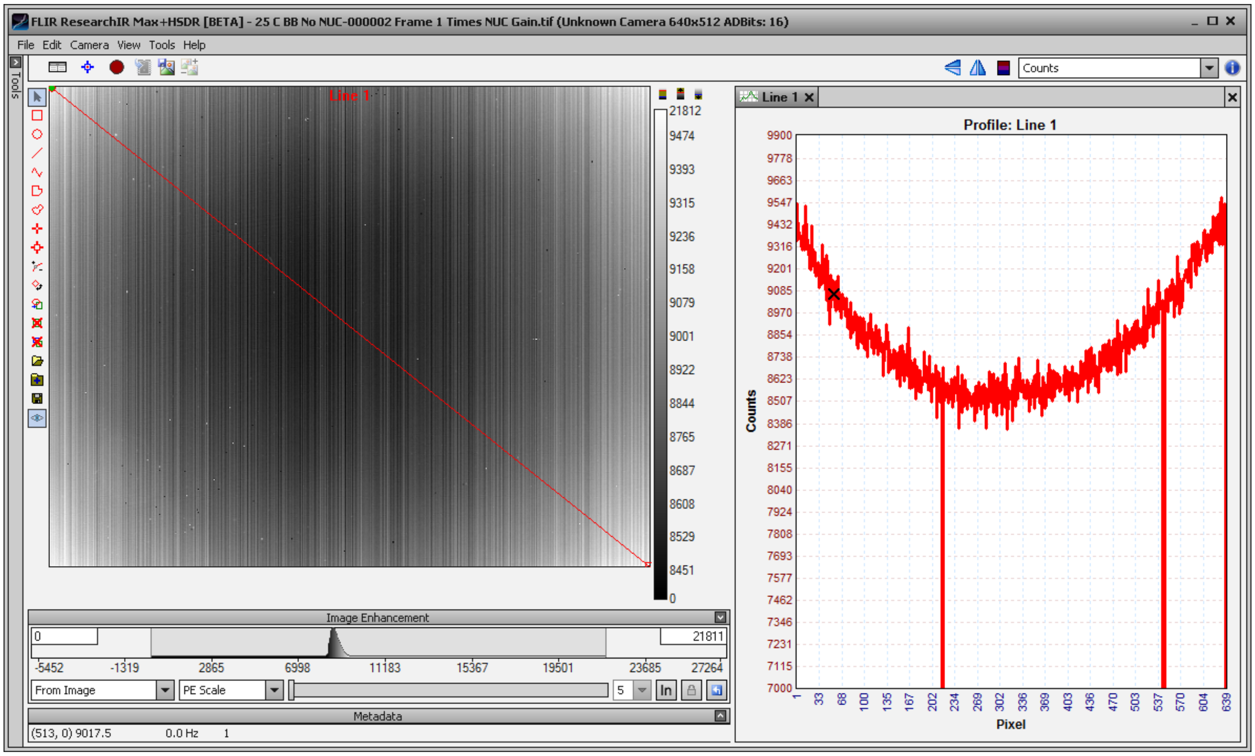


**Figure 5. Un-NUCed image of 25C blackbody multiplied by gain table.**

## Offset Table

The remaining nonuniformity can be removed by adding an offset table to the product of the gain table and the un-NUCed image. The offset table is calculated as follows:

11. Take 64 frames of image sequence while looking at a uniform source. This offset source is usually at a temperature halfway between the two source temperatures used to generate the gain table G(i,j). In this example, the offset will be generated by 64-frame image sequences of a 25C blackbody source.
12. Average the 64 frames into a single average frame to reduce the temporal noise. This average frame is called X(i,j)
13. Multiply each pixel in the image X(i,j) by its corresponding gain value in the gain table G(i,j), and call this image Y(i,j).
14. The offset table is the image O(i,j) = mean(Y(i.j)) – Y(i,j)

Figure 6 is the offset image O(i,j). Note the bowl shape is inverted:

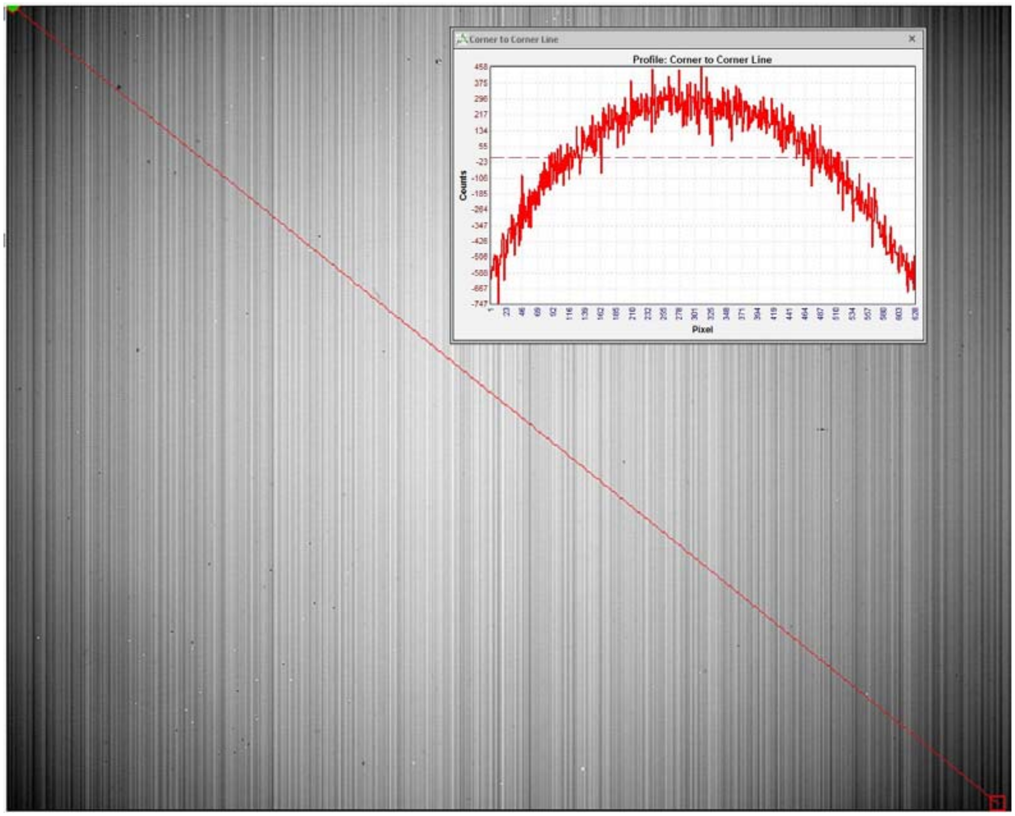


**Figure 6. NUC table offsets presented as an image**

Figure 7 is the image formed by taking Frame 1, multiplying it by the Gain image, and then adding the offset image. The standard deviation goes down to 4 counts, and the image is very well NUCed:

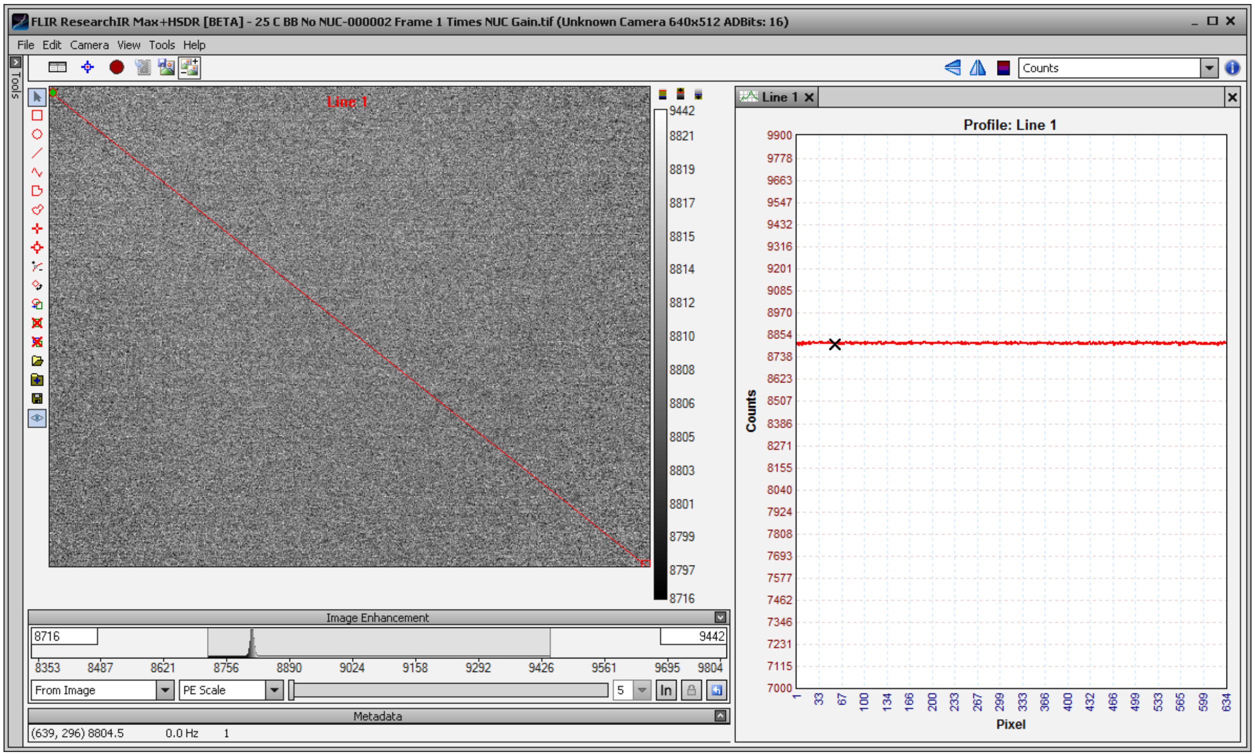


**Figure 7. Fully NUCed image of 25C blackbody**