A method for infrared temperature measurements of thin film materials with a low, unknown, and/or variable emissivity at low temperatures

by

Jason Neal Jarboe

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

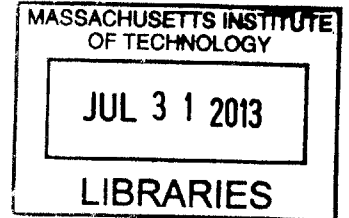Bachelor of Science in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013

© Jason Neal Jarboe, MMXIII. All rights reserved.

Author . . . .
Department of Mechanical Engineering
May 17, 2013

Certified by . . . .
Alexander H. Slocum
Pappalardo Professor of Mechanical Engineering
Thesis Supervisor

Accepted by . . . .
Anette Hosoi
Professor of Mechanical Engineering
Undergraduate Officer

A method for infrared temperature measurements of thin film materials with a low, unknown, and/or variable emissivity at low temperatures

by

Jason Neal Jarboe

Abstract

Accurate non-contact temperature measurements of objects using thermal radiation is often limited by low emission of IR radiation because of low temperatures and/or emissivities, or by the unknown or changing emissivity of the material being measured. This thesis covers an effort to build a practical, inexpensive, and widely applicable non-contact system for accurately measuring the temperatures of materials of low, unknown, and/or variable emissivity. The method to be used is intended specifically for those objects at low temperatures (below 100 degrees Celsius), which are conventionally the most difficult to accurately measure.

Thesis Supervisor: Alexander H. Slocum
Title: Pappalardo Professor of Mechanical Engineering

Acknowledgments

For Almow.

And I'd like to thank Frank.

And Alex.

And Janette.

And my parents. Without whom this would never have happened.

Contents

List of Figures

10

Chapter 1

Introduction

1.1 Problem Background

Occasionally the need arises to make a non-contact temperature measurement of objects that are nigh impossible to measure using conventional infrared thermometry techniques. One example of such objects are low temperature films with low, unknown, and/or variable emissivity.

Existing solutions for such problems tend to be custom for the particular material and geometry, and also tend to be expensive. A common approach to films is to use a custom bandpass infrared (IR) window that corresponds to a wavelength where the material being measured is unusually adsorptive/emissive. This approach subsequently struggles with the resulting low energy levels, and will not accurately measure a film with a different composition than the material it was designed for, which is often a requirement for converting applications.

What follows is an attempt to find a single solution to at least a portion of these problems that is inexpensive, widely applicable, and practical to implement.

1.2 Why IR temperature detection?

For most applications, the primary benefit to using IR radiation to measure temperature is that it requires no physical contact with the surface being measured. This is even an absolute requirement for some applications, such as those with a moving target, that require electrical isolation for safety, that require no contamination, marking, or heat sinking effects

from contact temperature probes, or that measure very hot targets at temperatures beyond what a contact probe can withstand. For some applications, another benefit is that it is often faster to measure temperature using IR than a contact probe, both in setup time, and in the response time of the sensor.

1.3 Basic methods of infrared radiation measurements

There are several sensing methods commonly used to measure infrared radiation. Broadly speaking, these can be separated into quantum devices that react directly to incoming photons, and devices which measure thermal effects of the incoming radiation falling onto an absorber.

Most infrared devices intended to measure surface temperatures of cooler than incandescent objects use a long wave pass filter to block "high energy" photons from things like reflections from the surface of sunshine, room lighting, or any nearby objects at high temperatures. This is done to minimize the disturbance of the temperature measurements caused by these reflections.

Photon counting/semiconductor photon detectors

Generally, photon sensitive IR detectors are used to detect "hot" objects, at temperatures far above that of the sensor assembly and any optics. Otherwise, the output will be dominated by the photons emitted by the sensor itself. In some cases, it may be possible to dig out the signal due only to photons from the target by modulating the incoming signal (using an optical "chopper", typically a mechanical assembly involving a rotating wheel), and measuring the AC output of the sensor.

Of the crystalline detectors, only MCT (Mercury Cadmium Telluride) sensors are sensitive to the wavelengths of infrared radiation emitted by low temperature targets, and they require liquid nitrogen for cooling, at the relatively long wavelengths I'll be coping with for cool targets. Sensors constructed from PbS (Lead Sulfide) and PbSe (Lead Selenide) are even less sensitive to low energy photons emitted by room temperature objects.

Treating IR photons like microwaves, and using an antenna for direct detection doesn't quite seem to stretch in the long wave infrared frequencies emitted as thermal radiation near room temperature. (Although I read at least one brave paper where the authors tried

14

to use emitted microwaves to measure brain temperatures.[1])

One of the newest types of infrared detectors, QWIPs (Quantum Well Infrared Photodetector) offer good sensitivity, but require cooling to work well near room temperature.

Thermal detectors

As a class, one nice characteristic of thermal devices is that they absorb a very wideband spectrum of radiation, and can therefore give a very flat spectral response.

Bolometer detectors are not fundamentally sensitive, but their performance has been improving with time. The most common construction for these devices uses pairs of matched thermistor detectors, measuring opposite sides of a thermal insulator to determine the heat flow across the insulator, which is proportional to the amount of radiation absorbed. Such bolometers require a reference current to operate, and it is difficult to make absolute measurements of very small amounts of infrared radiation.

There are two common types of sensors used for the detection of IR using polarization of crystals in response to thermal changes due to absorbed radiation. These pyroelectric and ferroelectric devices don't do DC signals, and so require optical chopping for temperature measurement applications. (They are commonly used with simple multifaceted lenses or mirrors for proximity detection of people for presence detection for automatic doors and security systems, wildlife camera traps, and such.) Ferroelectric sensors want to be run at *elevated* temperatures for best sensitivity. (That is, at temperatures just below their Curie temperatures, where they undergo a phase transition in their crystal lattice.)

The Golay cell[2] is a hardly used but rudely sensitive thermal system using pressure generated by temperature changes, but it is normally AC only and so generally requires mechanical chopping of the incoming radiation. As far as I can tell there is no fundamental reason that the pressure measurement couldn't be made absolute and work as a DC detector, though.

Thermopiles have several advantages over some of the other options for infrared detectors. Thermopiles are cheap and widely available from a number of vendors. They are fairly sensitive to incoming radiation, and linear over many magnitudes of signal. Like all thermal

---

[1] Reference the paper or delete the reference to it. However microwave antennas can be used to convert IR to heat for use with other thermal devices. This gives the ability to tune the device to nearly arbitrary wavelengths, and would be a neat way to make an inexpensive multicolor detector.

[2] Golay and Zahl (United States Patent US2424976), after the work of HV Hayes (1934)

detectors, thermopiles can react to low energy photons, and so can work well for targets at relatively cool temperatures.

They also require no bias voltage or current, which simplifies things a bit, and eliminates one source of errors and noise in a completed system. The voltage output of a thermopile for a given heat flow is not particularly sensitive to changes in temperature of the thermopile itself. And finally, thermopiles have identically zero signal when target and thermopile are at the same temperature, which is the key property I intend to exploit in this paper to make temperature measurements of surfaces with very low emissivities.

1.4 Basic thermopile physics

At the heart of a thermopile sensor, is the thermocouple. When two dissimilar metals are connected, a voltage is generated that is dependent upon the temperature of the junction. To use this "Seebeck Effect" to measure temperature of the dissimilar metal junction, a separate reference measurement of temperature is needed. When working with thermocouples, this is referred to as the cold junction, and traditionally an ice bath was used to establish a reference temperature of 0° Celsius. While the use of a physical ice bath is still not unknown, it's now much more common to use a electrical temperature sensor (thermistor, Resistance Temperature Detector (RTD), or a band-gap temperature sensor) to measure the reference temperature at the cold junction and then use that reference measurement, along with the voltage developed by the thermocouple to compute the temperature of the hot junction.
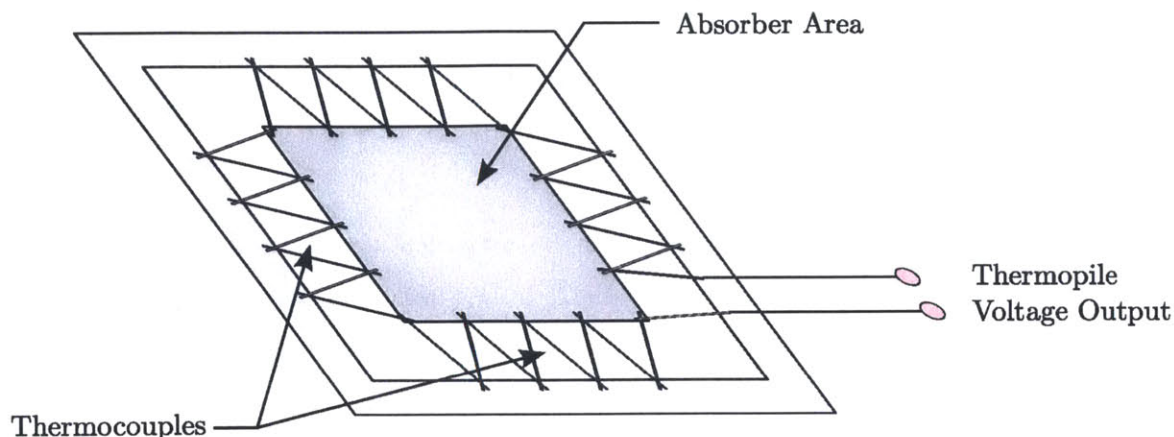
Figure 1-1: A diagram of a thermopile.

16

"Thermopile" is actually a generic term to describe a collection of thermocouples connected in series so that the voltages generated by each thermocouple junction are added. In the context of this paper, "thermopile" refers to a type of thermal infrared detector with an active area that absorbs (and emits) infrared radiation efficiently, and a series of thermocouple junctions that develop a voltage based on the temperature difference between this absorber, and a reference location in the sensor. There is a good deal of variation in the geometry and materials used in commercially available thermopile sensors. A diagram of the typical construction of an infrared thermopile sensor is shown in Figure 1-1.

Like the individual thermocouples it's comprised of, in order to make a practical measurement of temperature, a thermopile requires a reference temperature measurement of its "cold junction".

1.5   Infrared temperature detection using thermopiles

For simplicity's sake, the entirety of this paper assumes that the whole world is made of nothing but diffuse emitters. A diffuse emitter is a surface that emits and absorbs radiation from all directions equally well, and is sometimes termed a hemispherical emitter. None of the results presented here depend on this assumption, I'm just ignoring directionality as a matter of convenience. (And I'm also ignoring geometrical view factors as a convenience, since the geometry will be stationary for this application, and is multiplied by the sensitivity of the thermopile. Since the output sensitivity of the thermopile must be individually calibrated for each device as a practical matter, the gain adjustment can be made to account for both the sensitivity of the thermopile, and the geometry of the sensor.)

A blackbody is a theoretical object that emits thermal radiation at the maximum possible rate at all wavelengths, (and absorbs all radiation falling upon it). This functions as the standard to which the emission and absorption of real surfaces are referenced to. (A blackbody can be well approximated for experimental purposes by a small opening into a large interior cavity in a solid body.) The energy radiated per area by a blackbody at a given temperature and wavelength is given by the Planck Distribution:

$$E(\lambda, T) = \frac{2\pi h c^2 \lambda^{-5}}{e^{\frac{hc}{(k\lambda T)}} - 1} \tag{1.1}$$

where h is the Planck Constant (6.6262×10^{-34} J s), c is the speed of light ($2.9979246 \times$

$10^8 \ m/s$), λ is the wavelength of the radiation, k is the Boltzmann Constant ($1.380662 \times 10^{23} \ J/K$), and T is the temperature of the surface in Kelvin.

Emissivity, ε, is a property of a physical surface, and refers to the efficiency of emitting and absorbing infrared radiation, as a fraction of the radiation that would be emitted or absorbed by an equivalent blackbody. Therefore, the values of emissivity for real surfaces vary between zero and one. Emissivity can vary as a function of wavelength, temperature, and even with time as the condition of a surface changes.

The total amount of energy emitted by a real object (non-blackbody) can be found by integrating the product of the emissivity of the object and the energy emitted by a blackbody, Equation 1.1 over all wavelengths:

$$Q = \int_0^\infty \varepsilon(\lambda, T) E(\lambda, T) \mathrm{d}\lambda \tag{1.2}$$

An assumption often made is that emissivity is independent of wavelength (and everything else); the surface is then referred to as being a "greybody". In the case of a greybody, the emissivity is a constant between zero and one, and evaluating the above integral gives:

$$Q = \varepsilon \sigma T^4 \tag{1.3}$$

where σ is the Stefan-Boltzmann constant.

Unlike the direct measurement of IR radiation given by photon counting devices, the output voltage of a thermopile is a measure of the net heat flow delivered to the absorber area of the thermopile sensor by absorption and emission of infrared radiation by the absorber. As illustrated in Figure 1-2, this is a combination of energy that is reflected, transmitted, and emitted by the target surface, minus the energy that is emitted by the sensor itself. This net heat flow actually measured by the thermopile is termed Q_{NET} in Equation 1.4:

$$Q_{\text{NET}} = -Q_{\text{SENSOR}} + Q_{\text{TRANSMITTED}} + Q_{\text{EMITTED}} + Q_{\text{REFLECTED}} \tag{1.4}$$

For any given radiation that falls on a surface, it is either absorbed, reflected, or transmitted through the object. I'm simply going to assert that absorption is always equal to emissivity, so therefore:

18
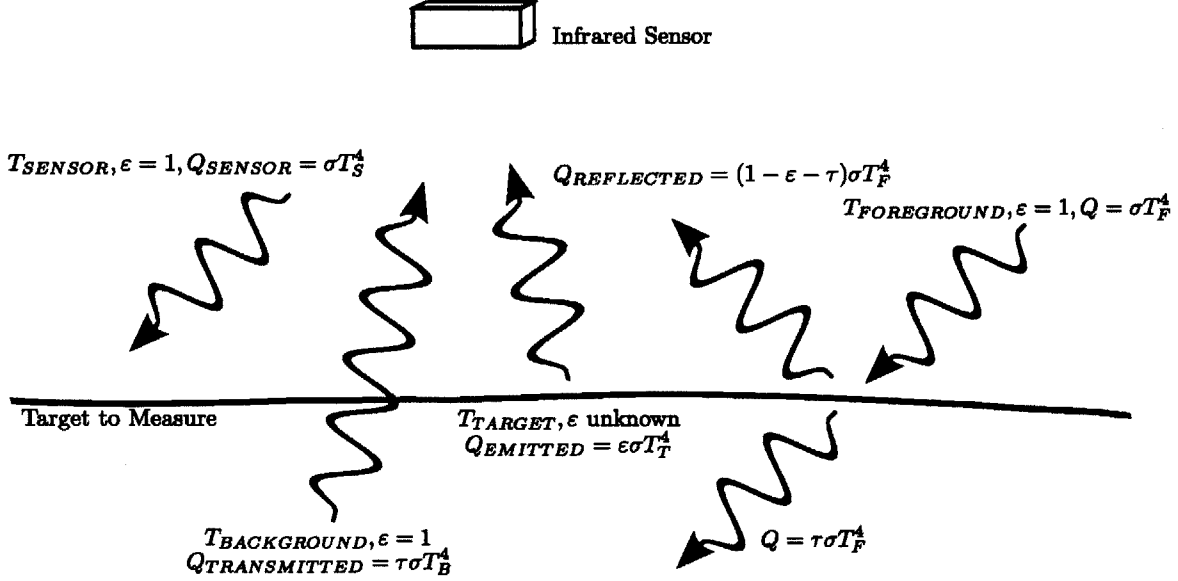
Figure 1-2: An illustration of the sources of thermal radiation absorbed and emitted by an infrared sensor, using greybody assumptions.

$$\varepsilon + \tau + r = 1 \tag{1.5}$$

where τ is the transmissivity, and r is the reflectivity of the target.

As previously mentioned, in order to use the output signal of a thermopile to compute the temperature of the target object under measurement, an independent measurement of the ambient temperature of the thermopile is needed. This reference temperature is required to compute the value of the heat flux emitted by the absorber area of the thermopile, Q_{SENSOR} in Equation 1.4.

In the case where the target is a greybody, all sources of reflections from or transmission through the target are at the same temperature as the sensor, the background, foreground, and sensor behave as blackbodies, and the sensitivity of the thermopile is invariant with respect to wavelength, then starting from Equations 1.4 and 1.3, the temperature of the target can be computed as:

$$Q_{\text{NET}} = -\sigma T_{\text{SENSOR}}^4 + \tau \sigma T_{\text{SENSOR}}^4 + \varepsilon \sigma T_{\text{TARGET}}^4 + r \sigma T_{\text{SENSOR}}^4$$

$$Q_{\text{NET}} = -(1 - \tau - r)\sigma T_{\text{SENSOR}}^4 + \varepsilon \sigma T_{\text{TARGET}}^4$$

Taking note of Equation 1.5, this can be simplified to:

$$Q_{\text{NET}} = \varepsilon\sigma T_{\text{TARGET}}^4 - \varepsilon\sigma T_{\text{SENSOR}}^4 \tag{1.6}$$

Rearranging this to solve for the temperature of the target gives:

$$T_{\text{TARGET}} = \sqrt[4]{\frac{Q_{\text{NET}}}{\varepsilon * \sigma} + T_{\text{SENSOR}}^4} \tag{1.7}$$

In the real world, it is rare for all of those assumptions to be true. Since sensors built for measuring the infrared radiation from cool targets normally have a filter window installed whose transmission of IR varies with wavelength, the integration of Equation 1.2 requires an additional term in the integrand for the variable transmission with wavelength, and becomes:

$$Q = \int_0^\infty \tau_{\text{Filter}}(\lambda)\varepsilon(\lambda, T)E(\lambda, T)\mathrm{d}\lambda \tag{1.8}$$

If the surface to be measured is assumed to grey, then the constant emissivity of the in Equation 1.8 can be pulled out of the integrand, and the equation simplifies to:

$$Q = \varepsilon \int_0^\infty \tau_{\text{Filter}}(\lambda)E(\lambda, T)\mathrm{d}\lambda \tag{1.9}$$

Now, assume an emissivity for (or measure the emissivity of) the target, assume the emissivity of the sensor's active area to be one, and additionally assume that the ambient background temperature is uniform and identical to the temperature of the sensor. With this set of assumptions, Equation 1.4, Equation 1.5, and Equation 1.9 can be combined and give:

$$\int_0^\infty \tau_{\text{Filter}}(\lambda)E(\lambda, T_{TARGET})\mathrm{d}\lambda = \frac{Q_{NET}}{\varepsilon} + \int_0^\infty \tau_{\text{Filter}}(\lambda)E(\lambda, T_{SENSOR})\mathrm{d}\lambda \tag{1.10}$$

The integrations can be done numerically, and the results used to generate a function f where the target temperature is given as a function of the other variables involved:

$$T_{TARGET} = f\left(\frac{Q_{NET}}{\varepsilon} + \int_0^\infty \tau_{\text{Filter}}(\lambda)E(\lambda, T_{SENSOR})\mathrm{d}\lambda\right) \tag{1.11}$$

Since all of the terms on the right hand side of Equation 1.11 are easily measurable, this

20

form is convenient to use in computing the target temperature when building an infrared sensor.

1.6 Overview of methods currently in use for measuring thin films.

Brightness (one color) with assumed ε, dark shadows

The assumptions made in this method are generally that the temperature of the background matches the temperature of the sensor, and that the background acts as a blackbody. For an ideal sensor with no IR filter, Equation 1.7 is used to compute the target temperature. For a sensor with an IR filter, Equation 1.11 can be used to compute the target temperature, for a given set of sensor inputs. The majority of infrared devices available for temperature measurement are based on the infrared brightness. This is the simplest, easiest and cheapest method to obtain a temperature using infrared detectors. These devices often come with a way for the user to adjust the emissivity used in the calculation of temperature to match the emissivity of the surface under measurement. Pitfalls for measurements taken this way include that it requires *a priori* knowledge of the emissivity of the surface to be measured in order to obtain an accurate temperature, they are somewhat sensitive to violations of the greybody assumption, and these sensors are also disturbed by reflections from background objects that are not at the same temperature as the sensor. It is possible to add another adjustment so that the user of the device can input a value for the device to use for the temperature of the background radiation.

One technique sometimes used when measuring a surface of low, unknown, or variable emissivity is to replace as much of the background as possible with a reflective surface. This causes some portion of the radiation reflected off of, or transmitted through the target to be photons originally emitted by the target, raising the effective emissivity as seen by the sensor. When dealing with most non-conductive materials ($\varepsilon \geq \sim 0.85$), this can raise the effective emissivity to the point that the surface can be treated as a blackbody and produce accurate measurements of the surface temperature. While this technique does help increase the measurable signal when dealing with surfaces of low emissivity, it doesn't help eliminate the uncertainty of the computed temperature related to an unknown value of emissivity, unless the emissivity was high ($\varepsilon \geq \sim 0.8$) to start with.

Use a bandpass filter that blocks radiation on the longwave side of the Planck Distribution for minimal dependence of calculated target temperature upon variations in ε, at the temperature range of interest

The temperature calculations in this case are identical to those of the one color brightness measurement, since this is just a special case of a one color brightness detector. The basic idea for this method is that the error of the measured temperature is less sensitive to errors in assumed emissivity when the wavelengths of radiation we are measuring are shorter than the wavelength of peak intensity (as given by Wien's Displacement Law) at the temperature of interest. So by choosing the wavelengths of a bandpass filter appropriately, we can reduce the measurement errors due to differences between the assumed emissivity of the target and the actual emissivity. This technique offers modest improvements, and like many established methods for improving measurements in the face of low or unknown emissivity, comes at the expense of a lower measurable signal. It also requires that the IR filter used be chosen for a specific target temperature range, for maximum effect.

Use a very narrow window at which the target is highly emissive ($\varepsilon \approx 1$)

This method of temperature measurement, is also essentially the same as that of the one color brightness measurement, and shares the solution given by Equation 1.11). In this case the bandpass filter's passband is very narrow, and the signal therefore is wee small. Over this narrow band of wavelengths, the assumed emissivity of the surface is ≈ 1.

This measurement affords all of the normal difficulties of a conventional one color brightness sensor, except that for the materials it is intended to be used on, the emissivity is effectively known to be one. The price to be paid for this is that the sensor signal is very small, and therefore the device is more expensive, with quality optics and electronics to perform acceptably. A drawback for this technique is that it will not work for materials that do not have unusually strong emissions in the sensor's pass band. This technique may also be limited by commercially available filters.

Two color ratio measurement, using two narrow bandpass sensors, and assuming the target is a greybody

This is the classic emissivity independent two-color method of infrared temperature measurements. By taking two measurements in two different bands of wavelength, you get two equations with two unknowns. If the target is assumed to be a greybody, then the two independent Equations 1.6 can be (not-quite-correctly) combined to give:

$$\frac{Q_{\text{NET1}}}{Q_{\text{NET2}}} \approx \frac{(\varepsilon\sigma)}{(\varepsilon\sigma)} \frac{(T_{\text{TARGET1}}^4 - T_{\text{SENSOR1}}^4)}{(T_{\text{TARGET2}}^4 - T_{\text{SENSOR2}}^4)} \tag{1.12}$$

Because of the wavelength dependence of the filters for the two sensors, each σT^4 in the equation above will actually be an integration over all wavelengths:

$$Q = \int_0^\infty \tau_{\text{Filter}}(\lambda)\varepsilon(\lambda)E(\lambda, T)\mathrm{d}\lambda$$

So the exact equation is:

$$\frac{Q_{\text{NET1}}}{Q_{\text{NET2}}} = \frac{\int_0^\infty \tau_{\text{Filter1}}(\lambda)\varepsilon(\lambda)E(\lambda, T_{\text{Target}})\mathrm{d}\lambda - \int_0^\infty \tau_{\text{Filter1}}(\lambda)\varepsilon(\lambda)E(\lambda, T_{\text{Sensor1}})\mathrm{d}\lambda}{\int_0^\infty \tau_{\text{Filter2}}(\lambda)\varepsilon(\lambda)E(\lambda, T_{\text{Target}})\mathrm{d}\lambda - \int_0^\infty \tau_{\text{Filter2}}(\lambda)\varepsilon(\lambda)E(\lambda, T_{\text{Sensor2}})\mathrm{d}\lambda} \tag{1.13}$$

Notice that if the target is assumed to be a greybody ($\varepsilon_1(\lambda) = \varepsilon_2(\lambda) = \varepsilon$), then the emissivities on the right hand side of Equation 1.13 can be carried outside the integral and then cancel out, and the target temperature only depends on the ratio of the two sensor outputs, and the (easily measurable) temperatures of the individual sensors.

This method is normally used for targets with very high temperatures, in which case the second term depending on the ambient temperature for each sensor becomes small enough to ignore, and a working device can use the simpler:

$$\frac{Q_{\text{NET1}}}{Q_{\text{NET2}}} = \frac{\int_0^\infty \tau_{\text{Filter1}}(\lambda)\varepsilon_1(\lambda)E(\lambda, T_{\text{Target}})\mathrm{d}\lambda}{\int_0^\infty \tau_{\text{Filter2}}(\lambda)\varepsilon_2(\lambda)E(\lambda, T_{\text{Target}})\mathrm{d}\lambda} \tag{1.14}$$

In cases where the greybody assumption for the surface is practical, the emissivities cancel out leaving the computed target temperature dependent only on the ratio of the surfaces:

$$\frac{Q_1}{Q_2} = f(T_{\text{TARGET}})$$

This $f(T_{\text{Target}})$ is, near as I can tell, most often empirically generated from test measurements. If the target is a greybody and if the emissivity is high enough for the noise level in the output to be tolerable, then the two color ratiometric sensor will give an accurate measurement. If the target surface is not a greybody, then ratiometric two color sensors can be provided with a "slope" adjustment for corrections to the $\varepsilon_1(\lambda)/\varepsilon_2(\lambda)$ ratio to allow for accurate measurements of non-greybodies after adjustment by the user. This requires an *in situ* calibration of the measurement, and the slope, $\varepsilon_1(\lambda)/\varepsilon_2(\lambda)$, itself can vary with temperature, which will still cause errors in the computed temperatures.

According to something old, but that I recently read [3], two-color devices were actually less accurate during testing than one color (brightness) detectors when used to measure a selection of random surfaces.

Because of the narrowed waveband, ratiometric sensors must deal with a smaller signal. They can be even more susceptible to reflections than one channel brightness sensors, especially with low emissivity targets. On the plus side, if the target is hotter than everything else in its environment, and the target is well-approximated as a greybody, a ratiometric sensor will accurately measure it even if it doesn't fill the field of view of the sensor, and if the target moves around in the field of view. Two-color sensors work well to overcome atmospheric adsorption, so long as the adsorption acts uniformly at both sampled wavelengths. (That is, if the atmospheric adsorption is itself "grey".)

For very demanding applications, the idea of the two-color pyrometer can be extended three IR wavebands, and extended far beyond, to the extreme of using IR spectrographs with thousands of channels, with measurements effectively being made by correlation to a known set of reference data for the material under measurement. Although this should eliminate slope type errors for characterized materials, the approach still has many of the normal difficulties of infrared temperature detection. And it brings some of its own difficulties, in that it requires a precise hardware for the coping with ever smaller signals due to the tiny wavebands, lots of optical hardware with moving parts for separating out the different

---

[3]Paul Nordine, "The Accuracy of Multicolor Optical Pyrometry", High Temperature Science, 1985.

wavebands, and substantial processing power for making the correlations. Even then it can not be guaranteed to be an improvement over a simpler system for the measurement of unknown surfaces and environments.

Chapter 2

A Disappearing Thin Film Pyrometer

Perhaps the oldest and simplest pyrometric technique is to estimate the temperature of incandescent objects by judging the color of the light they emit in the visible range. This method was later refined into a quantitative sensor, the optical disappearing wire pyrometer. In a disappearing wire pyrometer, a sight tube to the object to be measured is crossed by a filament wire whose temperature is monitored by measuring its electrical resistance. The user of the pyrometer manually adjusts the amount of electrical power dissipated in the wire until the color and brightness of the wire matches the color and brightness of the target being measured. In this state, the measured temperature of the wire is equal to the temperature of the target, the wire can no longer be discerned from the target in the background when looking down the sight tube, and is said to have disappeared. Figure 2-1 illustrates the concept.

The concept being explored here for the measurement of thin films of unknown, changing or very low emissivities is inspired by, and similar in concept to a disappearing wire pyrometer. In this technique, which I refer to as a disappearing thin film measurement, the field of view of the sensor in the foreground is filled by the film to be measured, and a target beyond the film fills the background. The background target plate is driven to a temperature such that the intervening presence of the film to be measured is undetectable. Once in this state, the film is the same temperature as the target in the background, which can be easily measured, unlike the film. In practice, since many thin films are also somewhat

27

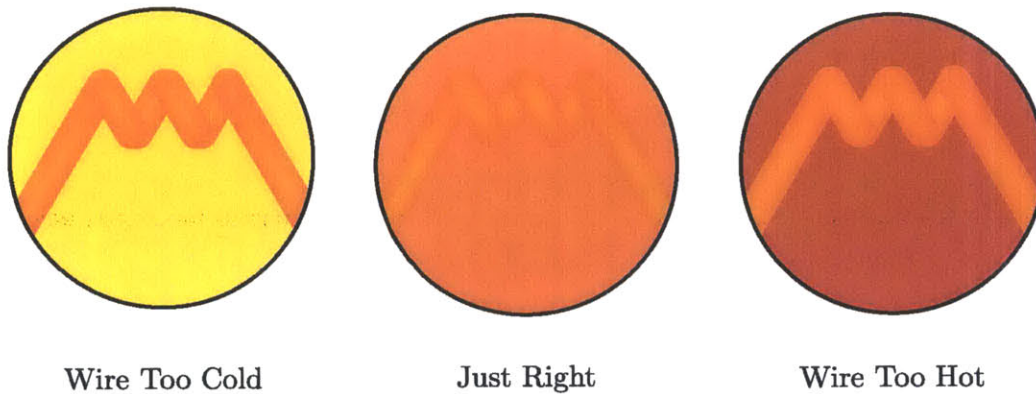| Wire Too Cold | Just Right | Wire Too Hot |

Figure 2-1: The concept behind an optical disappearing wire pyrometer. Note that the filament in all three cases is exactly the same color, and only the background color varies.

reflective, the front surface of the sensor facing the film must also be controlled to be the same temperature as the background (and therefore the film).

A large downside to this technique is that it is limited to target temperatures of no more than the maximum operating temperature of the thermopile itself. Most commercial thermopiles are limited to a maximum operating temperature of around 70°C to 85°C, although at least one manufacturer has made a model claimed to operate at temperatures up to 200°C.[1] As perhaps luck would have it though, these lower temperatures are exactly the conditions that are the most troublesome for traditional methods of IR thermometry.

This technique is also generally applicable to opaque surfaces of very low, unknown or changing emissivities [2], and is actually simpler to implement for this case, which is why I used a flat plate for the proof of concept testing in Chapter 3. This technique will also work for objects which are too small to completely fill the field of view of the sensor, change size or geometry, or move around inside the field of view of the sensor, and it will also work with materials that have strong emissivity vs wavelength variations, or strong variations of emissivity with direction.

A sensor using this technique must deal with small signal levels since the fundamental problem I'm trying to solve is that the material emits very little infrared radiation. The

---

[1] E Kessler, et al.,"High-Temperature Resistant Infrared Sensor Head" Proceedings of Sensor 2005, 2005, pp. 73-78 (micro-hybrid over in the former East Germany)

[2] (Note: someone's just HAD to have done this for opaque targets, I need to go search the literature for it and make a reference here.)Well, I found US Patent 4,900,162 by IVAC, they are doing heating/cooling for a tympanic membrane ear thermometer. Although that's more about minimizing the disturbance to the measured surface than dealing with low, variable, or unknown emissivity.

power draw will be higher than a conventional IR sensor, since the temperature of the sensor itself needs to be actively controlled.

From Equation 1.6 (or Equation 1.11 for sensors with a dependence on wavelength), when $Q_{\text{NET}} \Rightarrow 0$, then $T_{\text{TARGET}} \Rightarrow T_{\text{SENSOR}}$. Sticking with the idealized case to keep the math simple, from Equation 1.7, the errors due to emissivity become negligible when:

$$\left| \frac{Q_{NET}}{\varepsilon_{\text{ASSUMED}}\sigma} - \frac{Q_{\text{NET}}}{\varepsilon_{\text{ACTUAL}}\sigma} \right| \ll T_{\text{SENSOR}}^4$$

Re-arrange this as:

$$\left| \frac{1}{\varepsilon_{\text{ASSUMED}}} - \frac{1}{\varepsilon_{\text{ACTUAL}}} \right| \ll \frac{\sigma T_{\text{SENSOR}}^4}{Q_{\text{NET}}} \tag{2.1}$$

This inequality can be used to get an idea of the practical limits of this system.

Limits of radiated energy from thin films

The energy emitted by a film drops as the film's thickness is reduced. This will also place an ultimate limit on the thinnest films whose temperatures can be accurately measured by an infrared system.

To the first order, reflectivity is independent of emission and adsorption. Absorption (and emission, since I am asserting that they are equal in value) is a volumetric phenomenon, and there will be a constant fraction of the energy lost per distance traveled through the volume (for a given material, direction, and wavelength). Considering an infinitesimal block, the amount of energy absorbed as it travels through the block is the product of this specific absorption (which has units of m^{-1}) and the energy traveling through the volume:

$$dE = E_{out} - E_{in} = -\frac{E_{in} + E_{out}}{2} A(\lambda) \, dx$$

Separating the terms, and taking the integral of both sides gives:

$$\int_{E_{in}}^{E_{out}} \frac{1}{E} \, dE = - \int_0^{Thickness} A(\lambda) \, dx$$

29
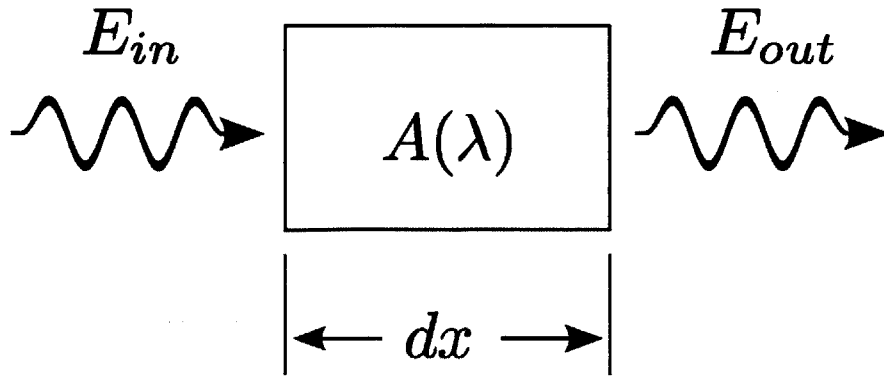
Figure 2-2: Infrared absorption is a wavelength and depth dependent material property.

$$\ln(E)\Big|_{E_{in}}^{E_{out}} = -A(\lambda) \times Thickness$$

$$\frac{E_{out}}{E_{in}} = e^{-A(\lambda) \times Thickness}$$

where $A(\lambda)$ is the specific absorption of the film's material.

As the thickness of the material goes to zero, the ratio of E_{out}/E_{in} becomes one, and the material absorbs (and by implication emits) no energy.

Chapter 3

Preliminary Testing

3.1 General Setup

For simplicity, the initial proof of concept testing was done using an opaque flat plate as an IR target. Minimizing the net heat flux detected by the thermopile of the sensor is easier when there is no transmission through the target, since only the sensor assembly and any IR reflections off of the target surface need to be controlled. See Figure 3-1 to see this simplified setup, compared to that shown in Figure 1-2.

Infrared Sensor

$T_{SENSOR}, \varepsilon = 1, Q_{SENSOR} = \varepsilon \sigma T_S^4$

$Q_{REFLECTED} = (1 - \varepsilon)\sigma T_F^4$ $T_{FOREGROUND}, \varepsilon = 1, Q = \sigma T_F^4$

Target to Measure $T_{TARGET}, \varepsilon \text{ unknown}, Q_{EMITTED} = \varepsilon \sigma T_T^4$

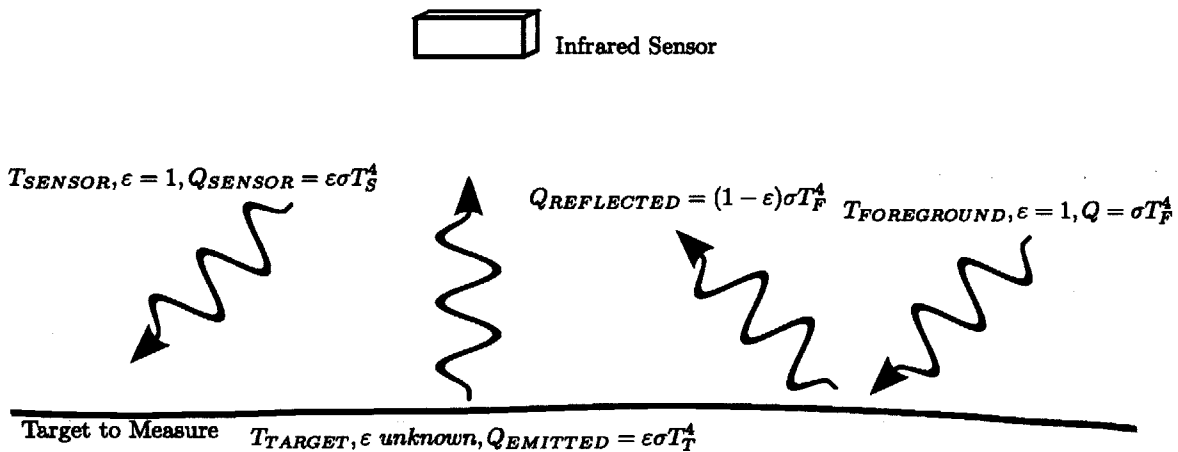Figure 3-1: An illustration of the simplification made possible by using an opaque target.

The aluminum target plate pictured in Figure 3-2, with one half painted and the other half roughly polished to have emissivities of approximately 0.9 and 0.07, respectively, was used to test the response of the system to changing the emissivity of the surface being measured while maintaining a uniform target temperature.
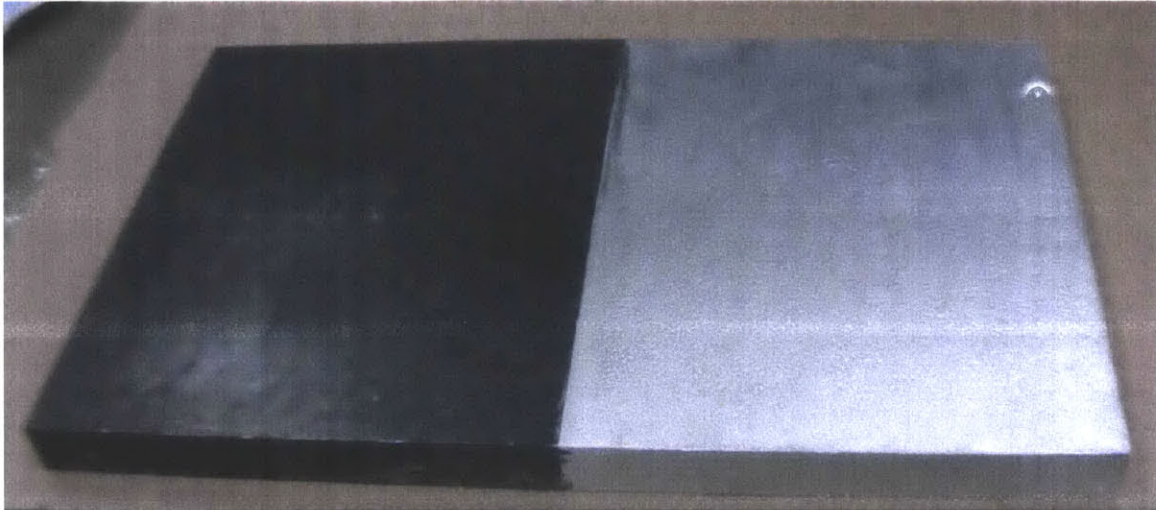
Figure 3-2: The target used for testing the proof of concept.

3.2   Initial Proof of Concept Testing for Servocontrolled Ambient Temperature

During testing, the aluminum target plate was elevated above the ambient temperature by approximately 15° C. Two infrared sensors were built, one as a reference using the conventional one color brightness method using an assumed emissivity of 0.9, and the other of identical construction and firmware except for the addition of a string of resistors embedded in the housing of the device used as heating elements. The assumption that the background radiation is uniform and the same as that emitted by the sensor was ensured to be reasonable by the design of this test setup; the two IR devices were placed very close to the surface being measured, so that nearly all of the background radiation falling on the section of the target plate being measured was emitted by the sensor device itself. See Figure 3-3 for an overview of the hardware and configuration used during this testing.

The raw sensor outputs for ambient temperature as well as for the thermopile voltage were transmitted to the computer controlling the experiments, by test firmware written for this experiment. The resistors acting as heating elements were driven by an external DAC according to current set by the controlling computer which was also parsing and recording the data generated during the experiment. The heater current was set using a proportional only control loop which was trying to drive the net infrared heat flux between the heated sensor and the target plate to zero. In the final device, this is expected to be a full blown

PID loop for better performance, to speed up the response time, and to eliminate the droop of the control variable noticed during this initial testing.
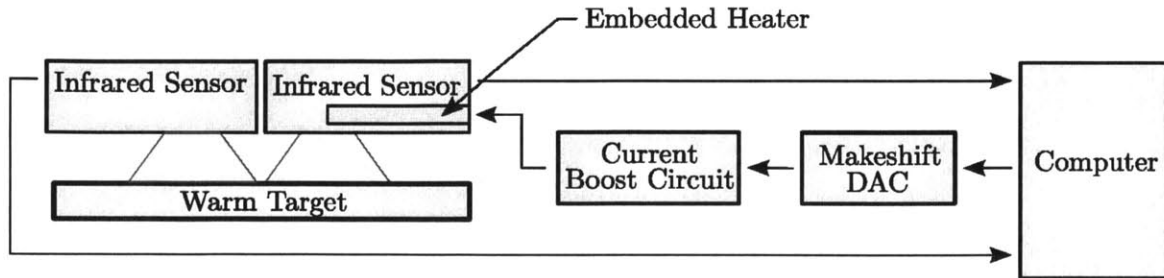

Figure 3-3: Experimental setup used for proof of concept testing.

The two devices were placed side-by-side above the target plate. For testing the effect of variations in emissivity on the devices, both devices were moved back and forth from the shiny to the dark side of the target plate while recording the device outputs.

Figure 3-4 shows that the change in IR signal is reduced by about a factor of seven during this experiment by controlling the temperature of the device itself. For the target and temperatures used in this experiment, this performance would be acceptable for many low temperature applications dealing with materials with unknown emissivities. But for the intended application of very thin films that I am targeting, it would be useful to accurately measure materials with emissivities at least as low as 0.01. For thin films, the setup will be complicated by the (mostly) transparent nature of such films, as that will require controlling both the temperature of the sensor, and also the temperature of the background of the material on the other side of the film. Maintaining a uniform temperature between the background, the foreground, and the thermopile itself is expected to be challenging.

I would expect that a general system intended for use on thin films would require an improvement of perhaps greater than a factor of forty beyond the performance achieved by this experimental setup, while dealing with many newly introduced problems. A correctly tuned PID control loop should be an improvement, but I also expect that, to get the performance that I really want, great care will need to be paid to the thermal aspects of the thermal control system and especially their relation to the housing of the thermopile sensor. If the heaters (or coolers, if present) cause gradients in the sensor assembly then those gradients *will* have a drastic affect on the output voltage of the thermopile.

The improved temperature measurements given by this method are evident when the calculated temperatures are plotted versus time. In the graphs shown in Figure 3-5, the
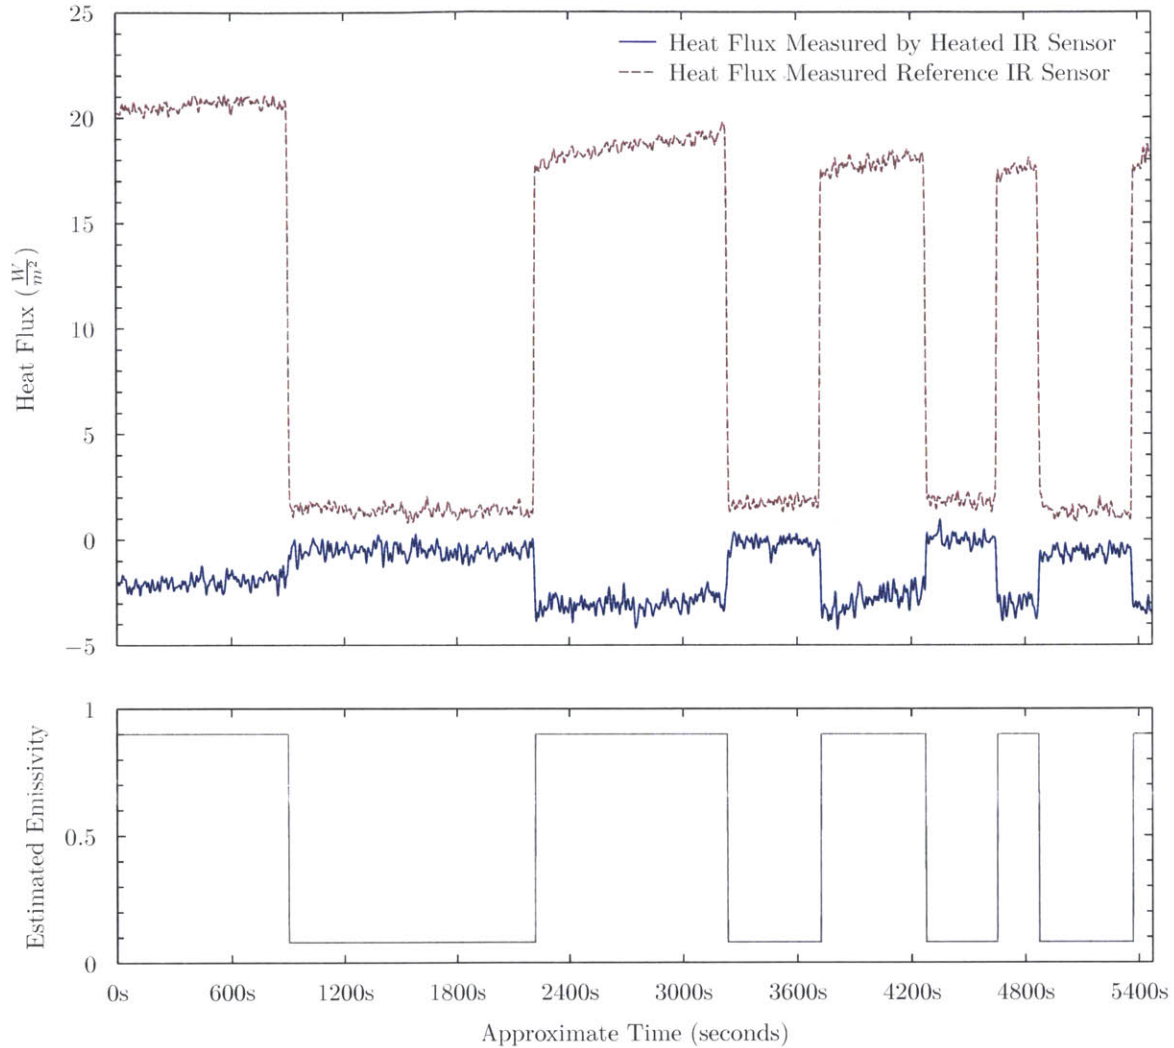
Figure 3-4: Measured heat flux of the two devices under test as the emissivity was varied over time.

device with its ambient temperature controlled to minimize the IR signal computes a far more consistent target temperature than the reference device, as the pair of sensors are moved back and forth between a highly emissive surface, and a surface with low emissivity. This improvement in performance is the expected improvement in this experimental setup, based on the heat fluxes shown in Figure 3-4. An example of the possible near ideal performance of this sort of system is shown in Figure 3-6.

These initial results of experiments from some cobbled together prototype hardware and software are encouraging, and indicate that the approach is quite feasible. A careful attention to details will be required to achieve the improvements needed to accurately measure the temperature of very low emissivity thin films with any reasonable response

Figure 3-5: Measured temperature of target plate, along with measurement of identical reference device. This is the corresponding output of the system to the heat fluxes shown in Figure 3-4, during the same test

time, but the approach should be feasible.
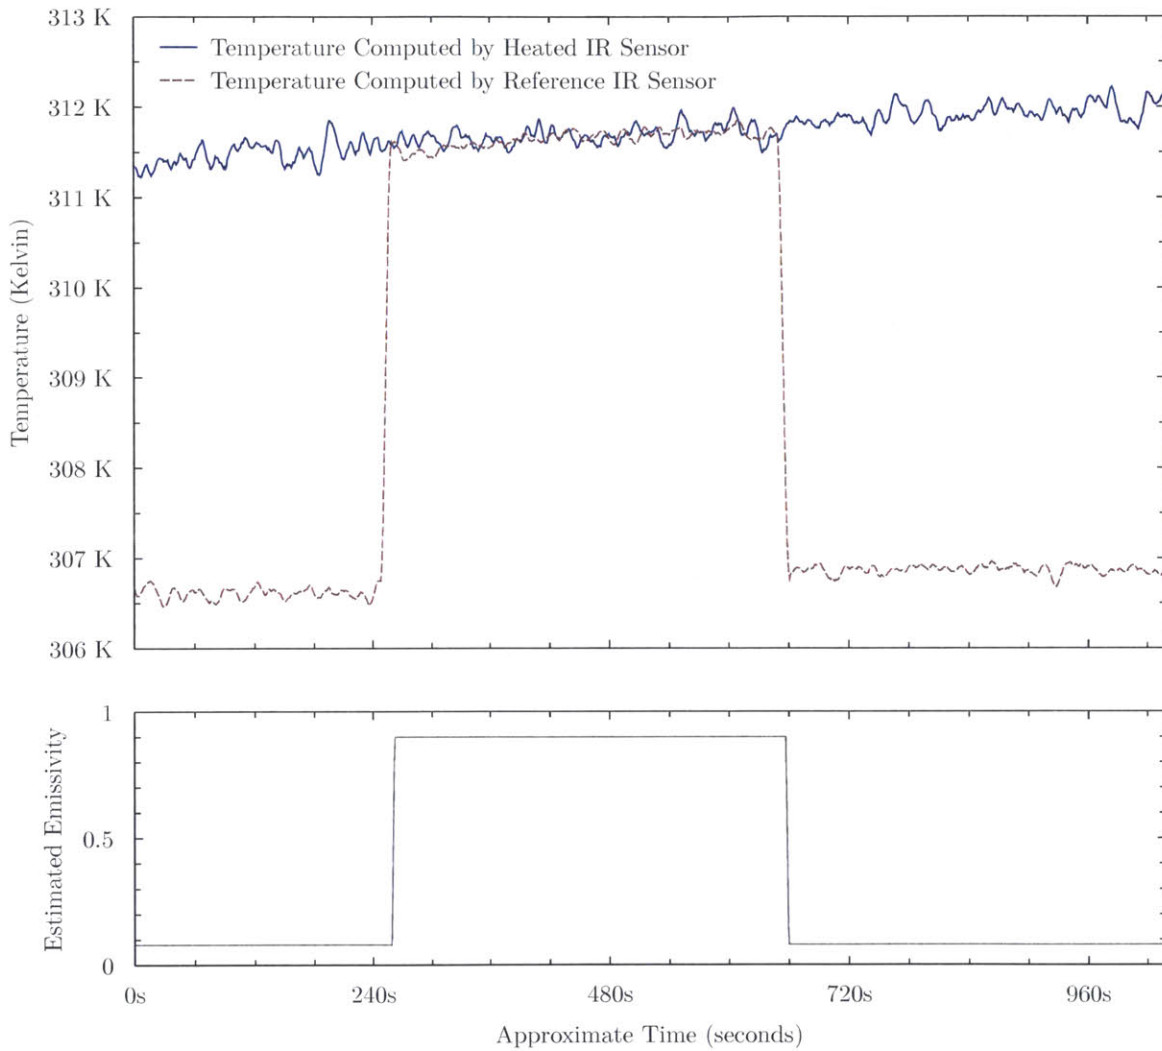
Figure 3-6: Measured temperature of target plate, along with measurement of identical reference device. This is a cherrypicked example, where I had been directly tweaking the control loop, to force it to give a near ideal output.
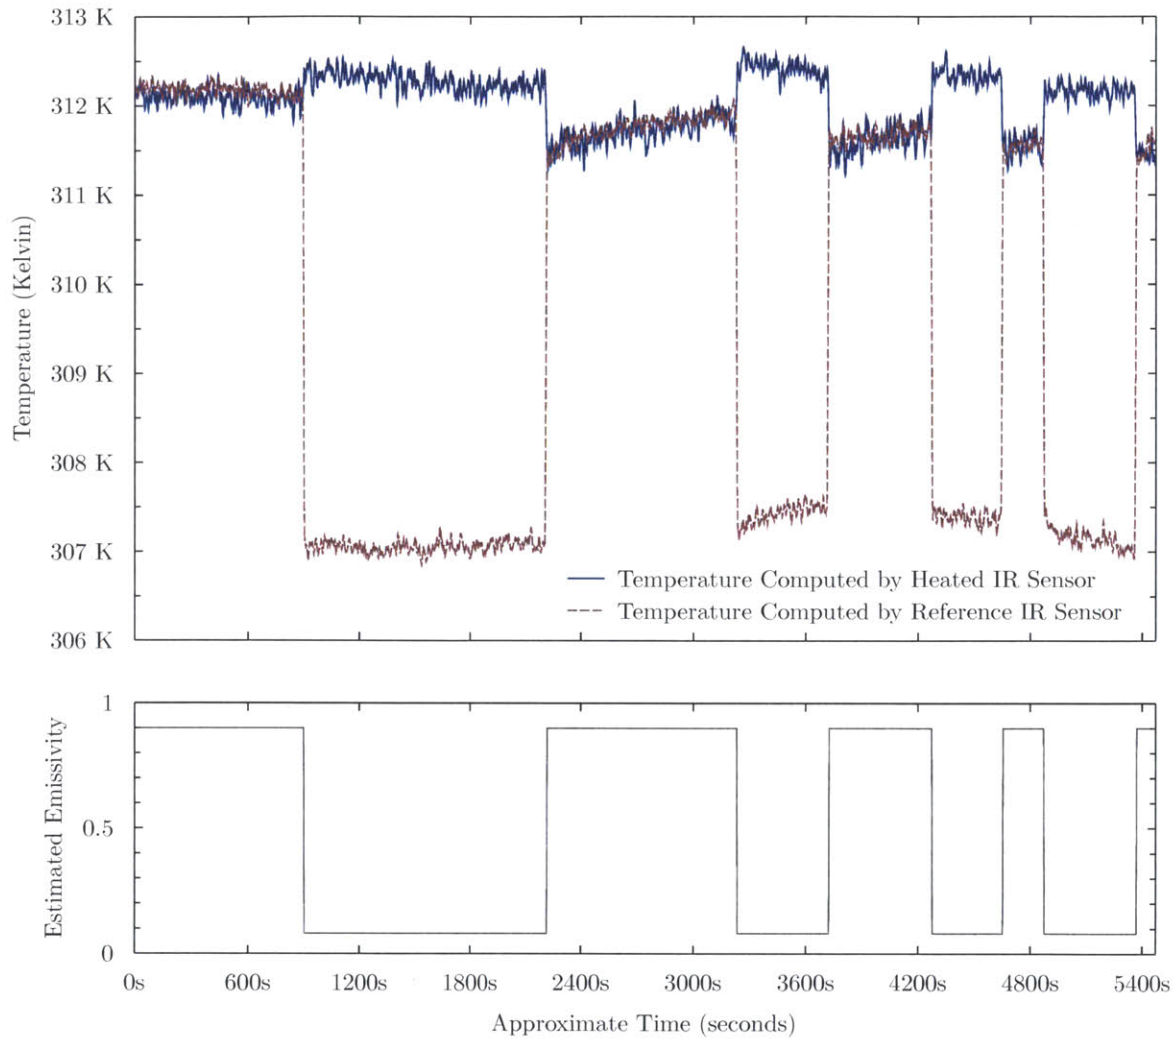
Chapter 4

Design choices and component selection

4.1 Electrical Design

The choice of amplifier topology was fairly arbitrary. Since the application requires an amplifier with a high input impedance, I chose a simple non-inverting amplifier. Since the design also required a fair amount of gain, and needed a different set of trade-offs for the output driving the analog to digital converter (ADC) than for the interface for the thermopile, I used two stages of gain.

I want the device to take steady state temperature measurements of surfaces with emissivities of down to $\varepsilon = 0.01$ with a RMS error of less than 1°C. To accomplish this, the aspirational design goal is for the measurement to have an accuracy of 0.060 μV and a resolution of 0.015 μV, while maintaining the maximum bandwidth attainable by the final system. The design endeavors to be the lowest noise, highest performance amplifier I can build in the time allotted.

In order to quantify the expected noise of the amplifier, an estimate of each possible noise source is made, and they are summed together in a root mean sum (RMS) fashion since the noise sources are all assumed to be independent.

Every resistor has "Johnson Noise" or "thermal noise" that is given by:

$$V_{JohnsonNoise} = \sqrt{4kTRB} \tag{4.1}$$

where k is again the Boltzmann Constant, T is the temperature of the resistor in Kelvin, R is the resistance of the resistor in Ohms, and B is the measurement bandwidth in Hz. Another source of noise found in resistors is called "shot noise". This noise is caused by the discrete nature of the electrons carrying the current, and the formula for computing it in a conductor with uncorrelated charge carriers is:

$$I_{ShotNoise} = \sqrt{2qI_{DC}B}$$
$$V_{ShotNoise} = \sqrt{2qI_{DC}B} * R$$

(4.2)

Another type of noise is referred to as "flicker", "1/F", or "pink" noise. The amplitude of this type of noise has an inverse relationship to frequency. This sort of noise, with an increasing amplitude as frequency goes down is typically prominent in active electronic components, such as voltage references and amplifiers. For a measurement bandwidth between the frequencies of B1 and B2, this noise can be approximated as:

$$V_{FlickerNoise} = \int_{B1}^{B2} v_n(F) \, dF$$

(4.3)

where $v_n(F) \propto 1/Frequency$ here is a constant for each device. The amount of flicker noise present in a given active electronic component typically needs to be estimated from a graph of noise versus frequency presented in the datasheet for the device. Amplifiers also have additional noise sources termed voltage noise, e_n, and current noise, i_n. These are referenced to the input of the amplifier, and these values can also normally be found in the data sheet for the amplifier.

The analog to digital converter itself will introduce some errors in the conversion. One unavoidable source of errors in that the digital levels are discrete, and so the conversion is "quantized" by rounding the analog signal to (hopefully) the nearest bit. This effect is completely negligible here, since the 24-bit sigma delta converter used has far more resolution than the other noise present in the system would allow measurement of. (In contrast, quantization noise is quite noticeable in the measurement of the thermistor value in the final assembly.) The other major error term for ADCs is variously measured as "integral nonlinearity", "ENOB" (Effective Number Of Bits), or "SINAD" (Signal to Noise and Distortion) in the device datasheets, and basically acts as a catch-all for all the other errors in the device.

For the purposes of estimating the noise expected in the amplifier chain, I will simply estimate the measurement bandwidth to be limited by the combination of digital filtering and through action of the autozero scheme to be a range of 0.5 Hz to 1.5 Hz.

Source	Type	Magnitude
Thermopile		
	Johnson Noise	39 nV
	Shot Noise	11 nV
First Op Amp		
Resistor Feedback Network	Johnson Noise	7 nV
	Shot Noise	2 nV
Reference voltage	Flicker Noise	145 nV
Voltage noise	e_n	~0V
Current noise	i_n	~0V
Second Op Amp		
Resistor Feedback Network	Johnson Noise	10 nV
	Shot Noise	2.9 nV
Reference voltage	Flicker Noise	6.3 nV
Voltage noise	e_n	~0V
Current noise	i_n	~0V
AtoD Converter		
	19 ENOB	9.5 nV
Reference voltage	Flicker Noise	290pV
2^{nd} Reference voltage	Flicker Noise	290pV
Total expected noise with a one Hertz measurement bandwidth		**210 nV**

Table 4.1: Noise estimates of the completed amplifier circuit, all referenced to the input of the amplifier chain.

The total noise estimated for measurements, shown in Table 4.1 indicates that to achieve the accuracy I would like from the amplifier may require a filter with a time constant of at least a couple of seconds.

4.2 Mechanical Design

The temperature of the sensor assembly in the steady state condition can be guaranteed to be practically uniform by the large thermal conductivity of the aluminum plates if they are well insulated from external heat loss/gain. The correct temperature of the plates can be guaranteed through the action of the PID control loop acting to minimize the value of the thermopile's voltage output. The more challenging design issue for the performance of

the device is the transients caused by the heating of the sensor. The basic design tradeoff is response time versus bandwidth and the possibility of wildly inaccurate readings if the thermal system is driven hard enough to cause substantial temperature gradients in the foreground and background plates, but especially if driven hard enough to cause substantial temperature gradients inside the thermopile. The maximum heating and cooling rates may need to be artificially limited to prevent unacceptable errors while the device is actively changing temperature.

The thermal design of the sensor was initially done as two halves. The design of the background plate is simple enough, in that the goal is simply to have the surface of the plate be controlled to a uniform temperature by heating the back of the plate. Applying a uniform heat flux to the back of the plate will change the temperature of the front surface accordingly. Since the geometry of the plate is prismatic, there is symmetry in two axes, and the front surface will always be maintained at a uniform temperature. The real world intrudes a bit, as I'll be heating with strings of resistors, so the surface won't be precisely uniform, but I'll spread enough heating elements far enough apart, to expect the front face to be acceptably uniform.

To guide the design for the thermal setup, the intended geometries were modeled and subjected to finite element analysis (FEA). As the primary concern was the dynamic response of the sensor device, and the minimization of any gradients the thermopile was subjected to during transient applications of heat, I simply enforced a condition of a constant heat flux into the body at the location of the heaters, and no heat loss through all other external surfaces. (In reality, there will be some heat loss through these surfaces when the device differs from the ambient temperature, but they should not significantly effect the dynamic response of the device, nor cause large thermal gradients in the device.)

While the outputs of the FEA program are believable and appear correct, some sort of sanity check of the results is prudent. As an engineering sort of sanity check, it should be obvious that since we are steadily adding energy to the plate, with no escape for it in the FEA model, the steady state solution will be a linear ramp controlled by the mass and heat capacity of the plate. With an input power of 11.378 Watts (The same loading used during the simulation, representing 16 Volts supplied to a 22.5 Ohm resistor network ($16 \ V^2/22.5 \ \Omega$).) applied to a half of the device, and using a density of 2700 kg/m^3, a specific heat capacity of 897 $J/(kgK)$, the expected steady-state rate of increase is:

$$6.25 \; in.^3 \times \frac{(0.0254 \; m)^3}{in.^3} \times \frac{2700 \; kg}{m^3} \times \frac{897 \; J}{kg \; K} / \frac{11.378 \; J}{second} = 21.801 \frac{s}{K} \; \text{or} \; 0.04587 \; \frac{\text{Kelvin}}{\text{second}} \quad (4.4)$$

The terminal slope shown by the FEA modeling output in Figure 4-1 shares the first five significant digits with the computation above.

I will now attempt a closed form solution to a model approximating that in the FEA analysis. The big simplification made here to enable an analytic solution is to assume that the heat flux applied to the back of the plate is uniform, rather than the distributed but discrete heating elements in the FEA model. So I'll be solving the heat equation for the transient solution to uniform heat flow applied to one face of a $^1/_2$ " thick aluminum plate, with no heat flow through the opposing face. The symmetry of the problem ensures that this is now a problem of heat transfer in a single dimension, so the heat equation reduces to:

$$\frac{\partial^2 T}{\partial x^2} = \frac{1}{\alpha} \frac{\partial T}{\partial t} \quad (4.5)$$

where the thermal diffusivity of the material, α, is:

$$\alpha = \frac{k}{\rho c_p}$$

where the thermal conductivity of the material is k, the density is ρ, and the specific heat capacity is c_p.

Attempting to solve Equation 4.5 by the method of separation of variables, I assume that the solution has the form of:

$$T(x,t) = A(x) * B(t)$$

Stuffing this proposed solution back into Equation 4.5 gives:

$$\frac{\partial^2 A(x)}{\partial x^2} B(t) = \frac{1}{\alpha} \frac{\partial B(t)}{\partial t} A(x) \quad (4.6)$$

Rearranging to group up all the dependencies on x and t yields:

41

$$\frac{\partial^2 A(x)}{\partial x^2} \frac{1}{A(x)} = \frac{1}{\alpha} \frac{\partial B(t)}{\partial t} \frac{1}{B(t)} \tag{4.7}$$

The two terms on each side of the equality in Equation 4.7 are completely independent, yet equal. So it can be reasoned that any solutions that exist to the equation must be equal to a constant, which I shall dub K. This realization results in a pair of ordinary differential equations which can hopefully be solved.

$$\frac{\mathrm{d}^2 A(x)}{\mathrm{d}x^2} - KA(x) = 0 \quad \text{and} \quad \frac{\mathrm{d}B(t)}{\mathrm{d}t} - K\alpha B(t) = 0 \tag{4.8}$$

The second equation in Equation 4.8 is the easy part, and has a general solution of:

$$B(t) = e^{K\alpha t} \tag{4.9}$$

Assume a solution to the first equation in Equation 4.8 has a form of:

$$A(x) = e^{\sqrt{K}x} \tag{4.10}$$

where all that's left to do is to solve for K by applying the boundary conditions, and verify that the resulting answer is actually a valid solution. The three boundary conditions for this problem are that the initial temperature is uniform, there is a uniform heat flux applied to one face of the plate, and that there is no heat flux on the other face. In equation form:

$$T\Big|_{x,t=0} = 300 \text{ K} \tag{4.11}$$

$$\frac{\mathrm{d}T}{\mathrm{d}x}\Big|_{x=0,t} = -\frac{q}{k} \tag{4.12}$$

$$\frac{\mathrm{d}T}{\mathrm{d}x}\Big|_{x=L,t} = 0 \tag{4.13}$$

For Equation 4.10, depending on the sign of K, three distinct cases exist:

$$K > 0 \quad A(x) = 0, \text{ No solution satisfies the boundary in Equation 4.11}$$
$$K = 0 \quad A(x) = Cx + D \tag{4.14}$$
$$K < 0 \quad A(x) = e^{\sqrt{-K}x} = C\sin\left(\sqrt{-K}x\right) + D\cos\left(\sqrt{-K}x\right)$$

Here I note that if I reverse my coordinates, so that x=0 is the insulated face, and that x=L is the side with the imposed heat flux, then the derivative of the $\cos(\sqrt{-K}x)$ term in the last line of Equation 4.14 is always zero at $x = 0$. This means that the derivative of the term $C\sin(\sqrt{-K}x)$ must also always be zero, since $\frac{d}{dt}\sin\sqrt{-K}x$ is not zero at $x = 0$, this means that C must be zero, and this entire $C\sin(\sqrt{-K}x)$ term is zero, and can be dropped. So my boundary conditions are now:

$$
\begin{aligned}
T\Big|_{x,t=0} &= 300 \text{ K} \\
\frac{dT}{dx}\Big|_{x=L,t} &= -\frac{q}{k} \\
\frac{dT}{dx}\Big|_{x=0,t} &= 0
\end{aligned}
\tag{4.15}
$$

And my solution for $A(x)$ is:

$$
\begin{aligned}
K = 0 \quad & A(x) = Cx + D \\
K < 0 \quad & A(x) = D\cos\left(\sqrt{-K}x\right)
\end{aligned}
\tag{4.16}
$$

Inspecting the boundary equation for $\frac{dT}{dx}\big|_{x=L,t} = -q/L$ here, with respect to the cosine term in Equation 4.16, the only way the derivative can always be a constant is if the cosine term is always evaluated at a multiple of π. From this, and the knowledge that K is a negative number, I deduce that:

$$
K = -\left(\frac{n\pi}{L}\right)^2
\tag{4.17}
$$

I also know, following the logic used in the gross engineering estimate earlier, that this solution will have a steady state portion that is a linear ramp with respect to time. This means that the solution isn't quite the $A(x)B(t)$ form that I assumed to start with, but is rather of the form $A(x)B(t) + Ft + G$, where F and G are constants. From knowledge about other solutions to the heat equation, I expect the particular solution to be an infinite series.

At this point I looked at a couple of references, and found no solution to this precise problem. So I guessed at the form of the infinite series. After finding something that seems correct, based on the boundary conditions, I revert the coordinate system back to the original sense, which yields what should be the exact closed form solution:

$$T = \frac{q}{k} \left[\frac{\alpha t}{L} + \frac{x^2}{2L} - x + \frac{L}{3} - \frac{2L}{\pi^2} \sum_{n=1}^{\infty} \frac{\cos(n\pi \frac{x}{L})}{n^2} e^{\frac{-(n\pi)^2 \alpha t}{L^2}} \right] + T_{init} \qquad (4.18)$$

Since I'm only really interested in checking this versus the data on the surface temperature output by the FEA modeling, I can ignore everything except for the value of x $(x = L)$ that holds my interest. After this simplification, the final form to use for checking the FEA output is:

$$T\Big|_{x=L} = \frac{q}{k} \left[\frac{\alpha t}{L} - \frac{L}{6} + \frac{2L}{\pi^2} \sum_{n=1}^{\infty} \frac{(-1)^n}{n^2} e^{\frac{-(n\pi)^2 \alpha t}{L^2}} \right] + T_{init} \qquad (4.19)$$
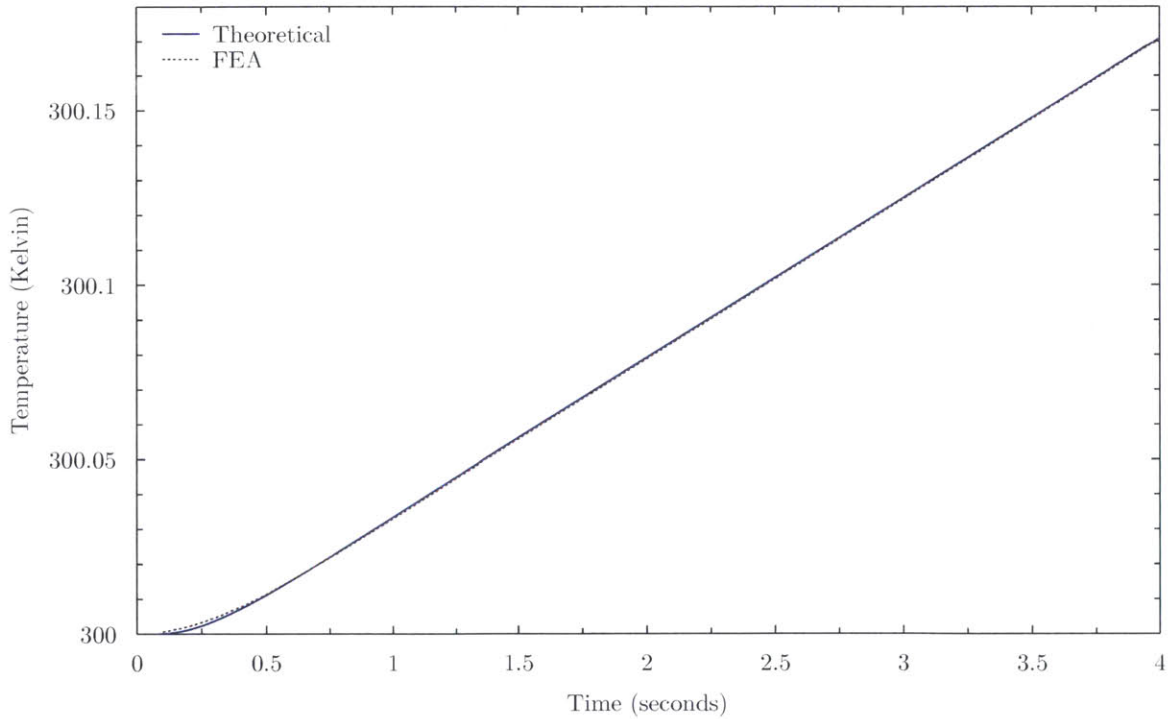

Figure 4-1: Transient predicted by FEA, along with the theoretical exact solution to a slightly simplified version of the same problem.

As a sanity check on this sanity check, as time (t) goes to infinity, the infinite series term becomes $\pi^2/12$, and the whole thing reduces to:

$$T\Big|_{x=L,t=\infty} = \frac{q\alpha t}{kL} + T_{init} \qquad (4.20)$$

The numeric slope of this equation, for this geometry and material is:

44

$$\frac{q\alpha}{kL} = \frac{1410.9 \ W/m^2 \frac{237 \ W/m \ K}{897 \ J/Kg \ K \times 2700 \ Kg/m^3}}{237 \ W/m \ K \times .0127 \ m} = 0.04587 \ \frac{\text{Kelvin}}{\text{second}} \qquad (4.21)$$

Which is a comfortingly familiar number, matching the result found in Equation 4.4, as well as fitting the steady state ramp in the FEA results.

The numeric values obtained from Equation 4.18 at the measured surface as a function of time is shown in Figure 4-1, along with the mean average for that surface from the FEA program with the same heating power applied. The numeric solution shown in the figure used the first million terms of the infinite series, which should be enough to give at least eight significant digits. The only notable deviation between the two sets of solutions is during the initial transient response, where the FEA model has the average mean temperature of the surface under consideration warming up faster than the exact analytical solution. The reason this is so (and why I've pointed out the type of average being used twice in this paragraph), is obvious in the animated output of the FEA; because the heat is being added in discrete locations, some parts of the surface are subject to higher heating loads, and change temperature much faster than the rest of the surface, which is further away from the heating elements. These early "hot spots" will pull up the mean temperature of the surface during the startup transient, as shown in Figure 4-1.

The time delay in the real world data is expected to be quite a bit slower than either of these models. In the models, the heat flux is applied instantaneously directly to the aluminum plate; in the real world, it is generated in the interior of resistors, and must be conducted to the surface through the material of the resistor and then through the thermal epoxy the resistor is buried in.

The design of the sensor housing half of the device is complicated by the necessary geometry for embedding the thermopile in the plate, and also by the need for not just the entire front face to be of a uniform temperature, but for the need to have the temperature of all the internal surfaces of the thermopile to be uniform and also at the same temperature as the front face of the sensor at all times. To try to minimize gradients inside the thermopile, I decided to keep a radial symmetry in the geometry and heat application near the thermopile. The optimization of the geometry to minimize the effects of thermal transients on the output of the sensing system was based on the results of FEA modeling.

For the control of the sensor assembly, I decided to use a PC "in the loop" to do

45

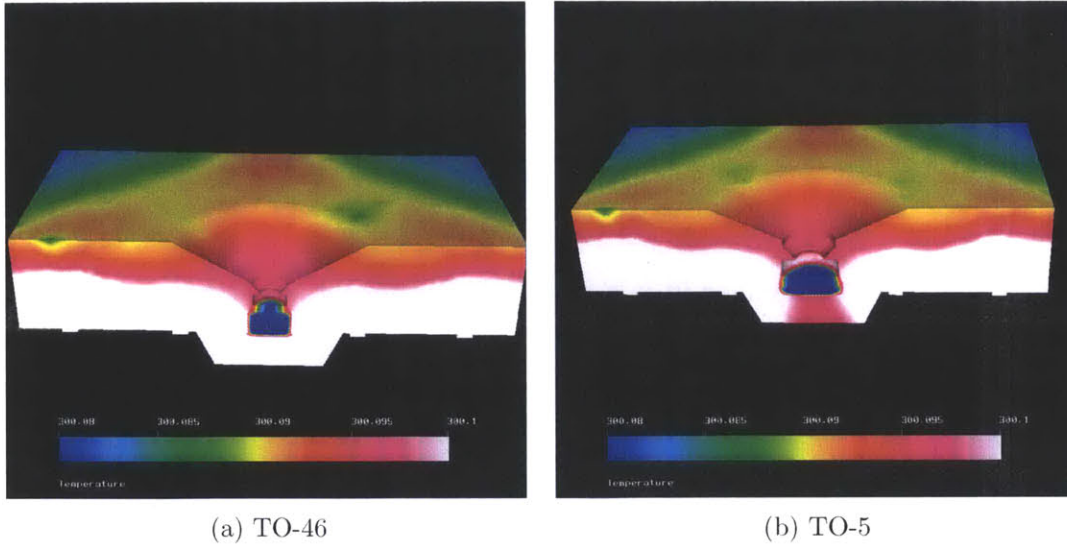(a) TO-46   (b) TO-5

Figure 4-2: FEA results of temperature distribution in the sensor housing during heating of device with a thermopiles in a TO-46 versus a TO-5 sized can. Note that the gradients on the interior surface of the thermopile are reduced by about 60% for the thermopile in a TO-46 can.

the computations for the PID control. This decision was made for a couple of independent reasons, one being that doing the PID loop in hardware would result in ungainly component values, given the slow response time of the system. The resistors and capacitors with the very large values required to deal with timescales on the order of many seconds are extremely susceptible to leakage currents on the PCB, and require careful guarding and compensation schemes to have any chance of performing suitably for this type of application. The other reason was that implementing the control system in software affords greater flexibility, potentially allowing for more precise tuning, and perhaps implementing more complicated algorithms, such as feedforward in the loop, if they are later deemed necessary.

There is a whole lot of unused circuitry on the amplifier PCB: various fine adjustments, variable gains, two 20 bit analog outputs, a pair of amplifiers intended to be used with normal or differential thermocouples, two hardware PID loops intended to keep the sensor housing thermally uniform if needed, an optical serial output in the event that too much noise was conducted into the amplifier housing during testing, and I forget what else. Mostly this was for future use, or just in case I ran into problems, and needed some of the included functionality to cope.

4.3   Component Selection

The operational amplifier (op-amp) used in the first stage amplifier/filter was selected on the basis of having a very low e_n, i_n, and $V_{OFFSET\ DRIFT}$.

The autozero switch in front of the thermopile amplifier needs to switch fairly fast, have low charge injection, and it needs to have very low leakage currents to and from the analog lines. The part selected comes in a SOT23-6 package, with the three analog pins all on one side of the device, which is nice for allowing a complete guard ring around the low level signals.

For a reference temperature measurement, I chose to use a thermistor since they are sensitive, are easy to use, and there is one built into the thermopile I'm using. Given the expected gradients as indicated by FEA modeling shown in Figure 4-2 , it would probably be best to use the reference temperature thermistor built into the thermopile, to try to minimize the errors in the reference temperature due to differences between the ambient measurement and the temperature of the absorber area of the thermopile.

My selection criteria for the analog to digital converter used for sampling the signal from the thermopile sensor was that it be accurate, fast, high resolution, and easy to use. In this context, fast just means a much higher than 120Hz sampling rate, since I know I will want to filter out 60Hz line noise in the digital domain. And I will want the response time of the whole analog chain for the thermopile to be in the millisecond range, to allow for rapid autozero cycles.

For the digital sampling of the thermistor, I chose to use the 10-bit ADC converter built into the microcontroller (with an additional two bits of active dithering of the signal, combined with lots of oversampling), mostly on the basis of basis of ease of implementation. I did not pay a lot of attention to the details of this section of the design, since "good enough" performance should be easily achievable. Nevertheless, testing of the amplifier on the benchtop showed better than expected performance. (This topic will get revisited.)

For an overview of the critical hardware in the amplifier design, see Figure 4-3. The operation of the autozero scheme is conceptually simple; measurements are alternately taken of the thermopile signal, and of a zero reference signal. The difference between these two measurements is then considered to be the actual thermopile signal.

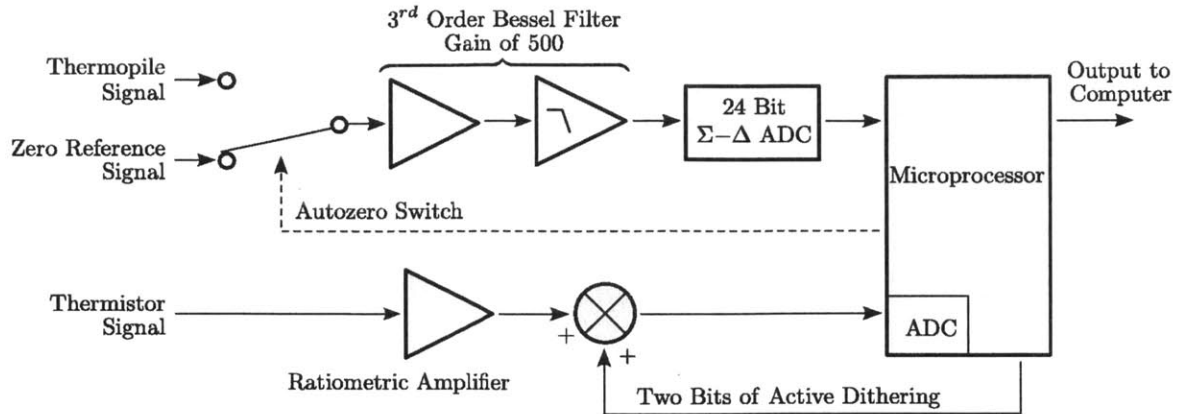The analog signal sampling is done at a rate of 7.3728 MHz, and an output rate 19200

Figure 4-3: A diagram of the amplifier design.

samples per second. The input to the ADC needs to be low pass filtered to prevent frequencies above the Nyquist frequency (half the sampling rate) from aliasing down into the lower frequencies. For this style of sigma-delta converter the required anti-aliasing filter should be designed to prevent aliasing at both the output rate, as well as the sampling rate. In the layout of the amplifier PCB, I placed pads for components to create a pair of cascaded single pole anti-aliasing filters, and left the actual design of the filter to be determined after sending off the gerber files for manufacture of the PCB. At that time, key bits of knowledge like how fast I would actually run the ADC were yet to be determined. When the filter was eventually designed, a sharper roll-off than two cascaded single pole filters was desired. This led to a third order Bessel filter being designed, and a few non-SMD components roached onto the PCB when it was stuffed. The Bessel response for the filter was chosen for its low distortion of waveforms in the time domain, which eases the implementation of an autozero system.

A cut-off frequency of 6kHz was chosen as a trade off between having an acceptable attenuation at the 9600 Hz Nyquist frequency of the output filter on the ADC, and allowing the maximal sampling bandwidth of the amplifier. The strategy chosen for the amplifier was to keep all of the bandwidth that the microcontroller on the amplifier PCB could eat, and do most of the required filtering of the many times oversampled signal in the digital domain. This way, the trade-off between response time and noise is not set in stone, and so the final system is flexible, and I should be able to change the firmware and software around in order to balance the bandwidth and noise without modifying the hardware. The digital filtering ultimately used during testing was extremely simple, being a boxcar moving

48

average.

Originally, I had envisioned that the signal to noise ratio of the completed sensor would be the be-all/end-all for the thermopile selection in this application. However, some initial experimentation with the FEA modeling suggested that it would be crucial for the sensor to react fairly well to thermal transients. This calls for a thermopile with a fairly fast response time, and as further FEA experimentation revealed, one that is as small as possible for a given construction style. I chose to use an Excelitas thermopile with a 1.2mm x 1.2mm absorber area and a wide bandpass filter for its high output signal, and in a miniature TO-46 can for minimized thermal gradients inside the can.

The printed circuit board (PCB) was designed and fabricated as a four layer board, with enough copper connected to ground that it was a pain to solder, even with thermal reliefs present on all ground connections. The board was mounted in an aluminum housing, and I wanted all connectors to have fully shielded coaxial cables. For the current testing, two connections (one for the thermistor, and the other for the serial output to the computer running the experiment) used panel mount SMA connectors. For the remaining connection to the two leads for the thermopile, neither conductor is ground. My first thought for this was to use triaxial cables, with a ground shield on the outside, and the two leads on the inner conductors. The availability of triaxial cables seems rather poor, and the prices rather shocking. So I improvised a plan B, which was to steal a metal overbraid from an available thermocouple, painstakingly thread it onto a thin coaxial cable, directly solder the cable at both ends, and ground it using a metal strain relief mounted on the same metal panel on the amplifier housing as the SMA connectors. This arrangement is not as handy as a disconnectable cable, but compromises must be made somewhere, if your budget is finite. In an effort to minimize the errors introduced by leakage currents in the high impedance and very low signal level ($\sim$ zero Volts), these signals were guarded on the PCB by surrounding them on the surface and all inner layers with a low-impedance signal maintained within a few mV of the same voltage level as the low level signals. This can cut down on the leakage currents, since with (hopefully) negligible voltage potential to drive them, leakage currents should be very small.

As another measure to minimize noise, the amplifier circuit is powered by a battery located inside the solid metal amplifier enclosure. The RS-232 level shifting was done external to the amplifier, since the switching noise of the MAX232 chip in the level shifter

is substantial.

The firmware onboard the microcontroller was written to electrically autozero the amplifier, and send updated data at a rate of 960 Hz, in order to be synchronous with 60Hz for the optimal rejection of 60Hz line noise in a simple manner. The sampling of the thermistor signal was done at a lower rate, around 11 Hertz, since that signal does not change rapidly. Both of these signals are subjected to a large amount of oversampling, to allow for filtering of the noise by error averaging.

4.4 Verification of Design Elements

After the design and construction of the filter was completed, an op-amp based differential amplifier followed by a resistor ladder attenuator was built to enable the output of a function generator to be scaled to the microvolt to millivolt level for testing the analog front end of the thermopile amplifier circuit. The test results illustrated in Figure 4-4 and Figure 4-5 show the achieved amplifier and filter response in the frequency and time domains, along with the simulated results that served as the basis for the design. I find the agreement agreeably agreeable.

After the device assembly was completed and some shakedown testing of the system was conducted, the reduction of noise (when compared to the benchtop testing) reduced the effectiveness of the applied dithering to the thermistor signal. A prominent stairstep tendency was present in the measurements of the ambient temperature, with a magnitude of $1/4$ of a Least Significant Bit, as shown in Figure 4-6. This is merely of academic interest, as a case of "so much better, it's worse!" The performance of this subsection of the design exceeds the requirements for this application, as each quarter LSB step only represents $\sim 0.017°$ C, and the errors in the surface temperature computed by the sensor introduced by this quantization are not exaggerated by the computations.

As a test to see if the two halves of the sensor would have issues staying thermally balanced when driven, I took the two halves of the sensor, mounted them with the two working surfaces facing each other, but with no thermal connection between them, wired the two sets of heating resistors in series, and measured the transient output of the thermopile sensor when a step voltage input was applied to the heaters. The resulting transient was well below 100 μV, which is well into my happy range.
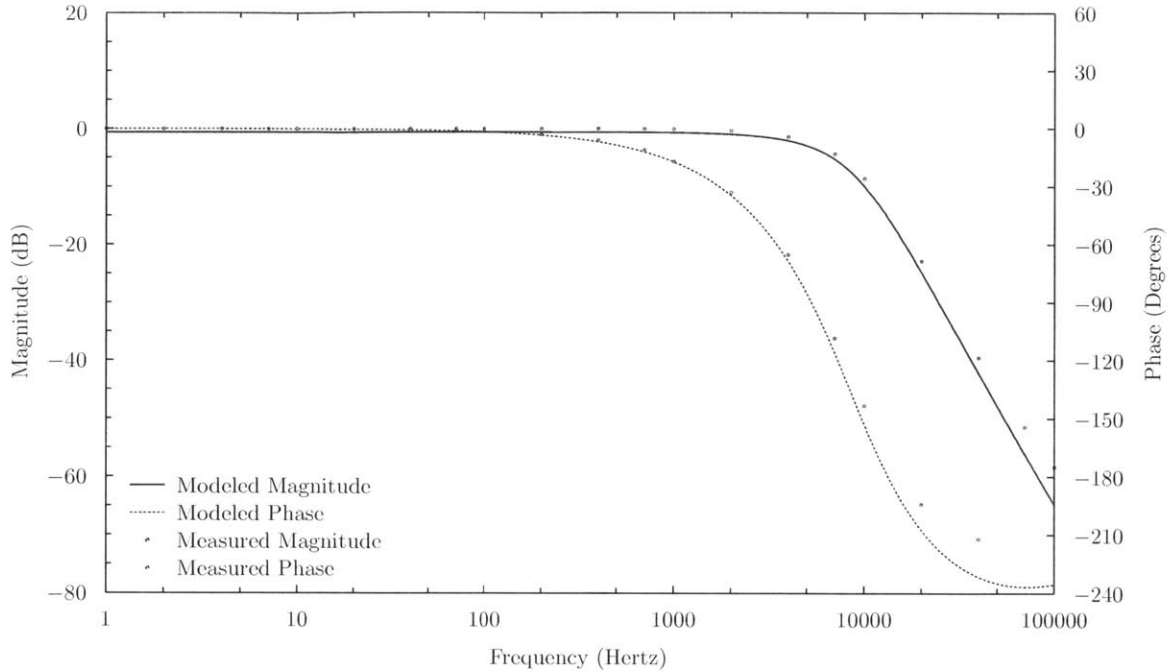
Figure 4-4: A Bode plot showing the frequency response of the analog filter on the front end of the amplifier circuit, that functions as an anti-aliasing filter, as well as bandwidth limiting for noise reduction. The modeled results shown were extracted from LTspice during the design of the filter.

I expected to have some issues with the electrical leads of the thermopile conducting heat into and out of the thermopile housing, and the thermal gradient that this would set up. My contingency plan to try to make a more solid thermal connection between the sensor housing by bonding the shield grounds to the aluminum body of the sensor housing was swamped. After that failed to even come close to working, I was forced to get a little extreme. I drilled four small holes, approximately 0.100" in diameter and 0.250" deep, in a drill circle near the sensor itself. I wrapped about six inches of very fine magnet wire around a short section of teflon tubing. One end of the wires was routed through the center of the tubing, the tubing was then inserted into the hole, and potted in place with a thermally conductive epoxy. The coaxial cables leading to the amplifier were connected to the end of the wires emerging from the teflon tubes, and the other end of the wires were soldered to the thermopile leads. This arrangement added about two Ohms of resistance to each connection, and no noticeable heat sinking issues through the thermopile leads remained.

Initial testing of the completed amplifier circuit on a bench shows that the peak to peak range of the readings, was about 15 μV to 20 μV. As can be expected, the dominant
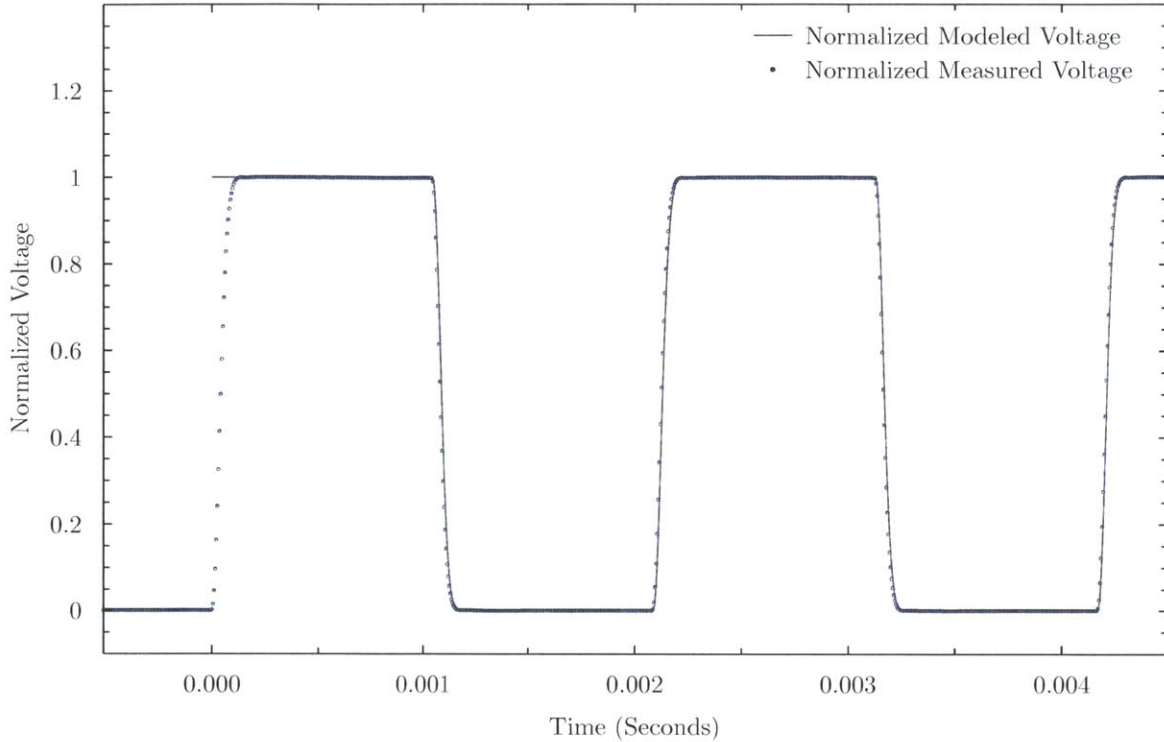
Figure 4-5: Here's a plot of response of the analog filter on the front end of the amplifier circuit to a 480Hz square wave. The modeled results shown here were also generated by LTspice during the design of the filter.

noise was 60Hz noise, with an amplitude of about 12 μV. The amplifier was designed to electrically chop the signal at 960 Hertz, and gather data at a rate of 19200 samples per second. Both of these were chosen to be synchronous with 60Hz, to enable an easy digital filter to reject 60Hz noise. The only other prominent narrowband noise in the analog output of the amplifier circuit was an approximately 4kHz signal caused by the switching frequency of the chopper stabilized op-amp used as the first stage of gain/filtering. Since the frequency of this noise is below the Nyquist frequency of our sampling rate, and very far above the near-DC signals of interest for this application, effectively filtering it out in the digital domain should present no problem.
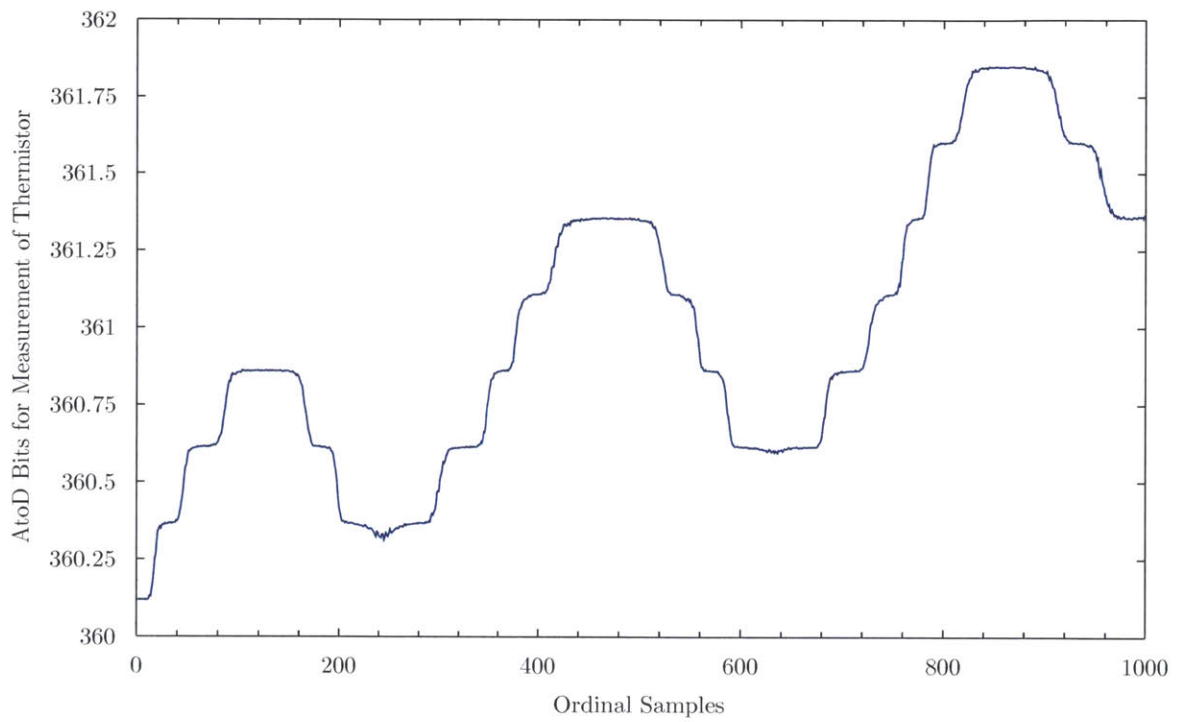
Figure 4-6: Measured bits of the thermistor amplifier output, showing quantization errors as a stairstep function with a height of 1/4 LSB.

Chapter 5

Design of Experiment

5.1 The Dark Side, and The Bright Side

Two separate sets of tests were done with the completed sensor. During the first round of testing, only the sensor housing half of the complete device was used, and the sensor was tested with the same half painted aluminum target that was used for the proof of concept testing. The test setup for these tests is shown in Figure 5-1.
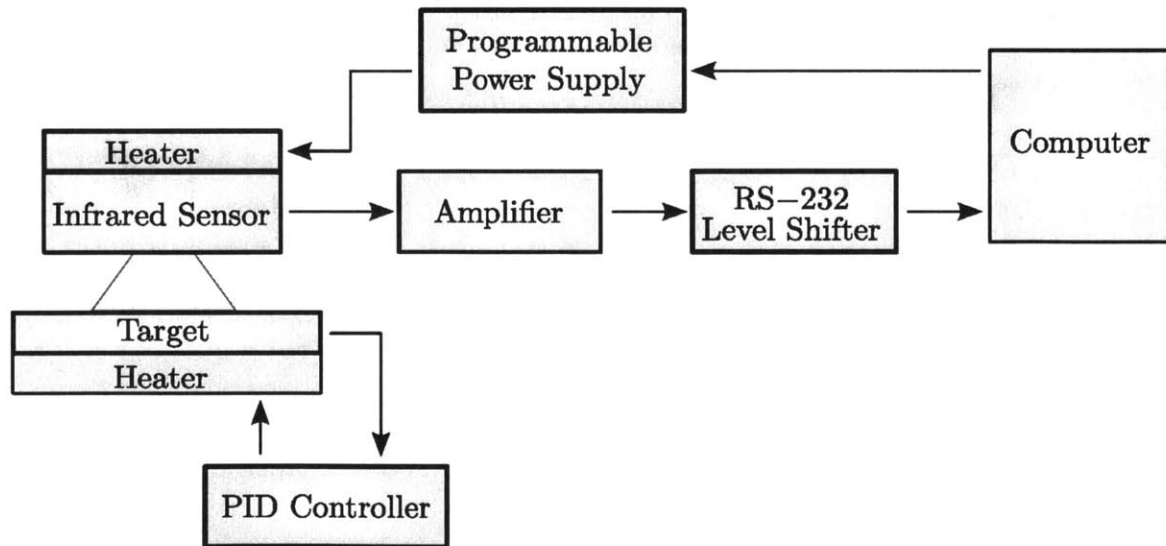

Figure 5-1: Experimental setup for testing the new sensor with the half-shiny aluminum target plate.

The computer controlling the experiments was set up to control the output of the sensor to zero, by controlling the amount of power delivered to the heaters on the back of the sensor body using a plain old PID control loop. For this testing the reference measurements of the

target plate were taken with a handheld thermometer at irregular intervals, and typed into the computer. This data was recorded to file, along with the outputs of the amplifier PCB.

When the power applied to the heaters in the device were first tested under working control as a closed loop, the control variable (the heat flux as measured by the thermopile) was not very stable. It did not go into a sinusoidal oscillation as it would if the loop gain were set too high, but rather was wandering aimlessly around the control setpoint, and very "surge-y". After a period of puzzlement, I realized that I was controlling the system by driving the voltage supplied to the heaters, while power dissipated is proportional to the supplied voltage squared. Correcting this to supply a voltage proportional to the square root of the desired power output did not improve the control loop's behavior to a satisfactory level of performance, unfortunately.

To improve the performance of the control loop, I implemented what I later discovered the literature calls a nonlinear modeled predictive controller (NMPC). I measured the output response of the sensor assembly when subjected to something approximating an impulse heating disturbance. This was then used to generate a piecewise linear function that was used as a model the output of the sensor, given the heating history. The measured response of the device is shown in Figure 5-2, along with the piecewise linear segments that were used as a model for the device's response. For the control of the system, I was driving the voltage supplied to the resistor heaters such that the predicted heat flux fifteen seconds in the future was held at zero. I experimented around with several different controller topologies for implementing a PID control loop to correct for any residual errors from the NMPC, without too much improvement over running the NMPC in an open loop at the temperatures where the device's response was measured. I kept tinkering with the control loop, trying to minimize the range that the measured heat flux wandered in. A diagram of the setup of a typical control loop is shown in Figure 5-3. However, having the controller operate with at least a closed loop gain in addition to the predictive model used was essential for reasonable operation far from the room temperatures where the device was characterized, with the simple and temperature independent response approximation I used.
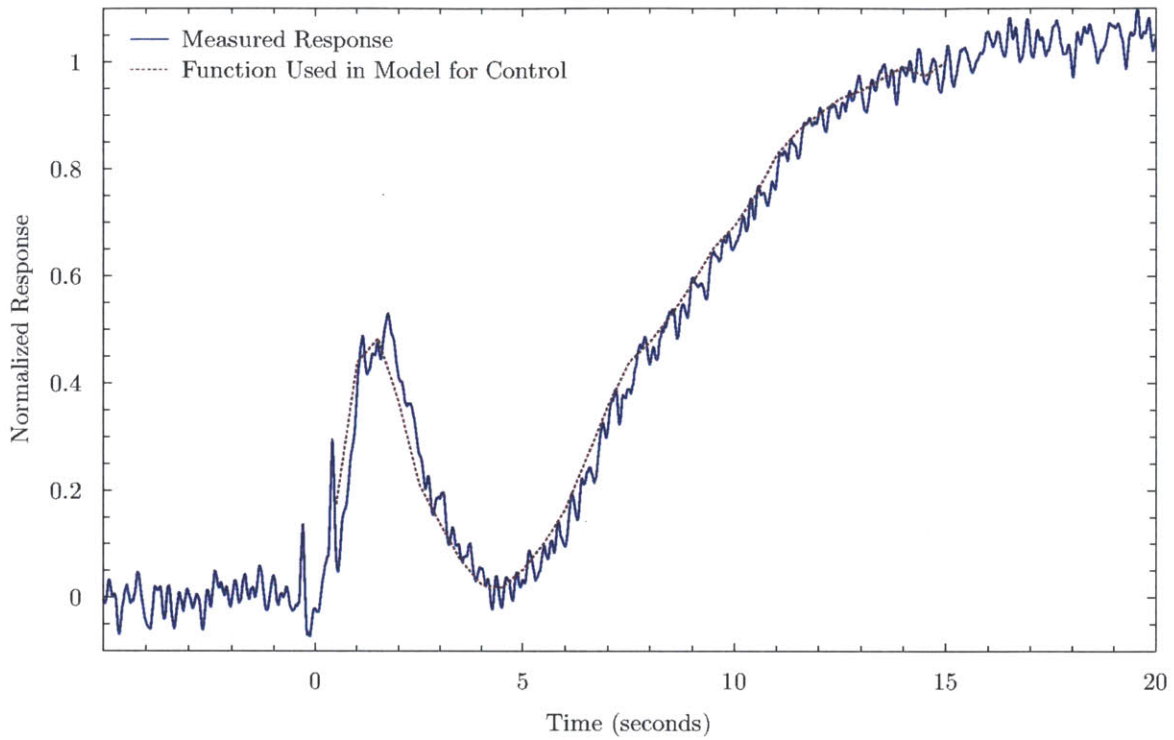
56

Figure 5-2: Measured impulse response of the unit with the device at room temperature.
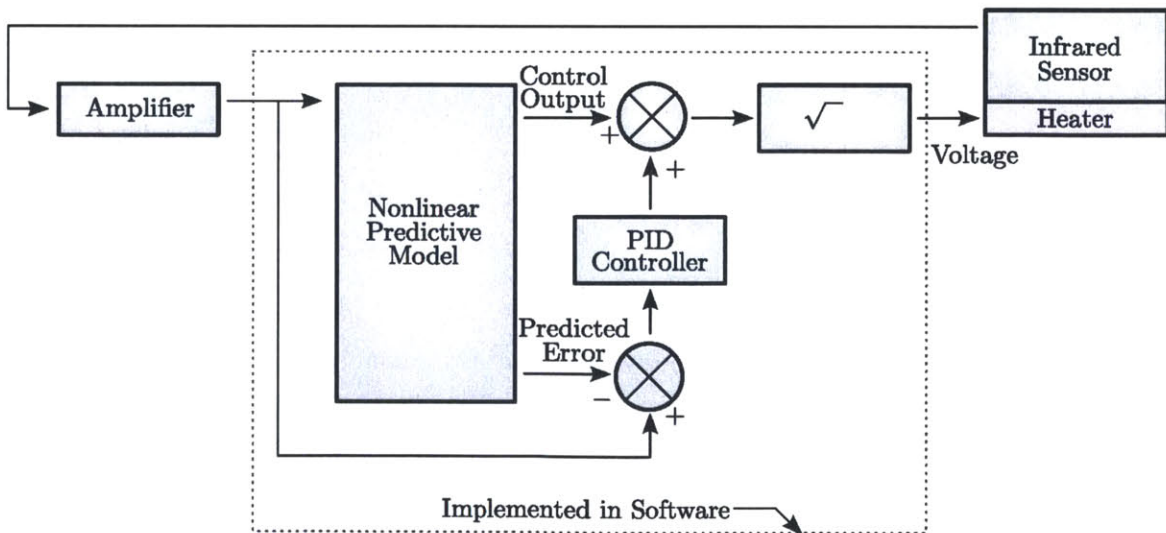

Figure 5-3: A diagram of a typical control loop used for maintaining a null heat flux between the sensor and target.

5.2 Thin Films

The second set of experiments for testing the achieved performance of the sensor was conducted by attempting to measure a moving thin film of "stretch wrap" with the sensor.

I initially envisioned a reel-to-reel setup for driving the film one way, then reversing the direction when all the film was used. The first efforts to implement this did not go well. The film would wander around on the rollers excessively, and often would roll completely off eventually. In an effort to just get something working, a strip of plastic film was spliced into an endless piece, and the film handling grew enough parts to let it run true. Temperature control was of the heating block inside the chamber, and the rest of the box was just assumed to match the temperature shown by the controller after the apparatus had reached steady state. This setup worked – after a fashion, but didn't give results useful for verifying the performance of the new sensor.

In the third effort, the performance of the film handling was further improved. A fairly precise and made-to-be-accurate reference measurement of the film temperature was added. A heat-sink-fan combination was added to the heater to promote tighter thermal control and fewer gradients in the box. The PID controller for the box now takes its feedback measurement from an IRt/c (an infrared sensor whose output is designed to mimic a thermocouple over a fairly narrow range of temperatures) pointed at the outgoing spool of film instead of a thermocouple buried in the heater block. Even though the film handling was now probably usable for reversible reel-to-reel operation, the experimentation needed to all be in a single direction now that all the measurement and thermal control was concentrated on the outgoing film, in an effort to achieve the maximum possible accuracy for testing.

In the final setup for measuring film, an Exergen DX-501 was used as a reference to measure the temperature of the film (and mostly the plastic roller it's on at the measurement location) just before it exits the temperature controlled chamber. The DX sensor cannot accurately measure the film itself, but the film will very quickly take on the temperature of the roller that's guiding it, and that can be accurately measured. The DX-Series thermometer is also completely inside the temperature controlled box, which has circulating air, and was allowed to "soak" at a temperature for about an hour prior to any measurements being made. The film, the roller, and the DX should all be quite near isothermal, therefore the output of the DX thermometer should be as close as possible to the actual temperature

of the film.

The overall design of the final experimental setup used for the testing of the sensor with thin film targets is shown in Figure 5-4.
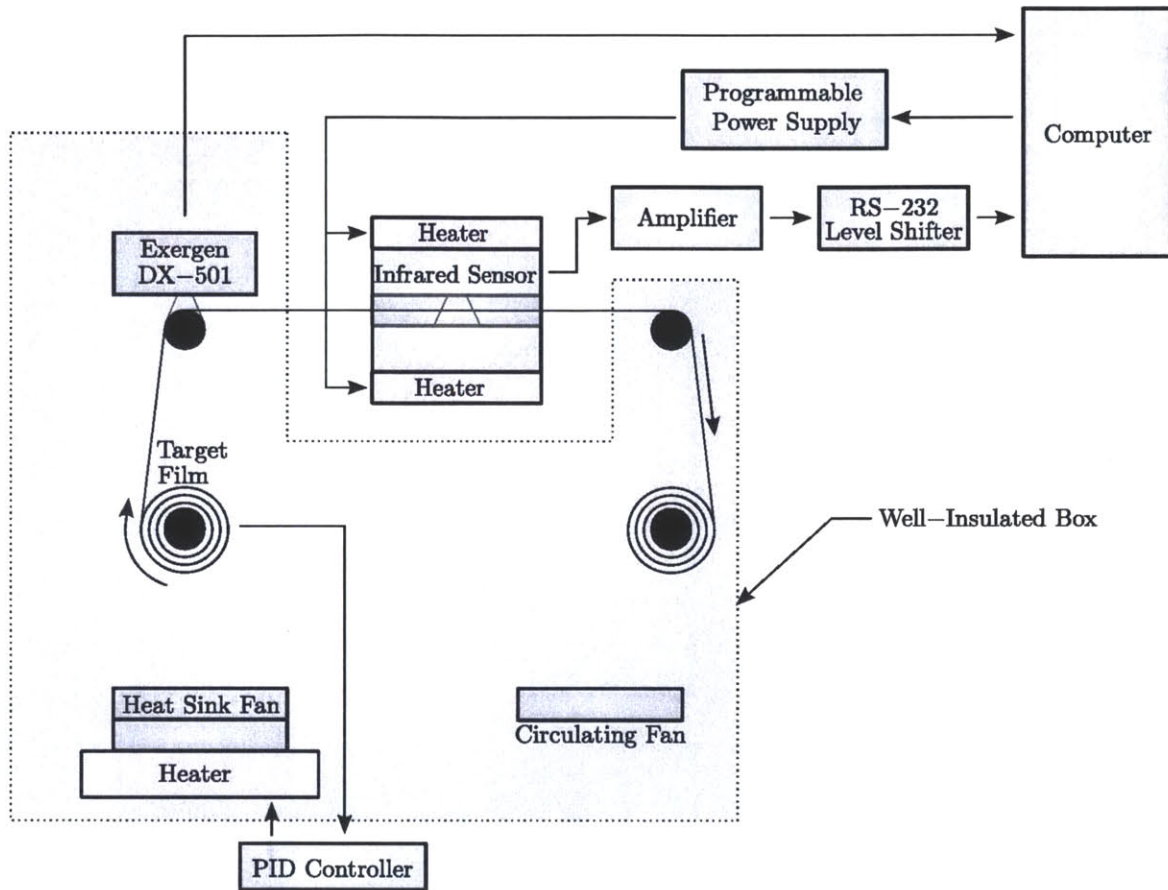

Figure 5-4: Final configuration of the experimental setup built for measurement of thin films.

Chapter 6

Testing of Final Hardware

I'm going to gather some general notes regarding the experimental testing here. During the experiments, the external PID controller stayed at the setpoint, within the display resolution of the controller. I attempted to keep the film speed slow and steady when it was moving, as a noticeable change in the film temperature could be seen if I varied the speed of the hand drill used to the drive the take up reel for the film. Since I was now running the film through the test setup in a "run-once" manner, in order to minimize the amount of film used during testing, I would advance the film in pulses while the sensor was warming up. This shows up in the graphed data as the large irregular square wave in the computed temperature, before the sensor reached steady state operation. The battery on the DX ran low, and needed to be replaced during testing. It is suspected that the extra noise in the output of the DX during the test at a nominal 110° F was due to the low battery.

I used an assumed ε to handle non-steady state readings, in the typical manner of a one color brightness infrared thermometer. In this case, I set the emissivity to ~0.04, to match the actual emissivity of the film. Using an application specific guess for emissivity like this will give the best results, but on the other hand, it defeats the point of wanting this sensor to be a universal cure-all. No matter the emissivity assumed by the sensor, the performance of the sensor in steady state operation should be acceptable for all but the most demanding applications. So the question of an adjustable or a per application setting of emissivity is only about the dynamic response required by the application. Figure 7-4 in the Results section of this paper shows the computed target temperatures using emissivities of 0.9 and 0.04 for the set of experiments on a thin film presented there. Perhaps a fixed value in the

middle of the range could be a good compromise.

The sensor was not carefully calibrated during any of this testing, as the critical issue from a performance perspective is that the null signal level is very near to zero.
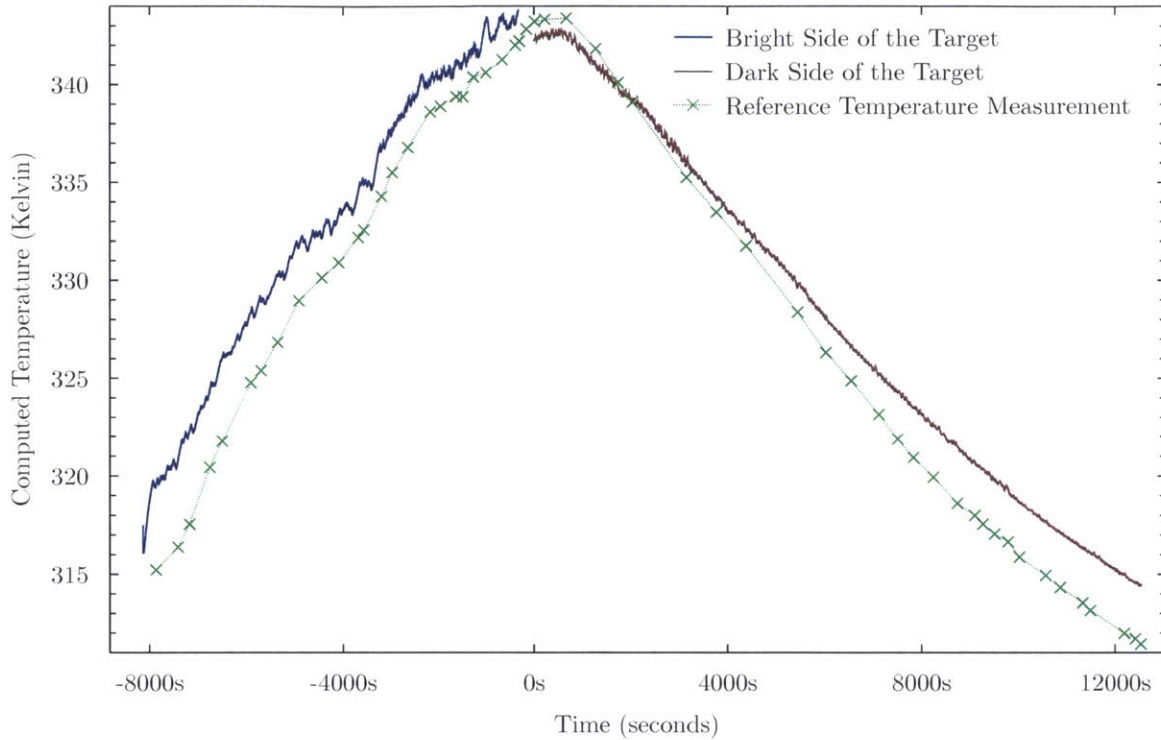

Figure 6-1: The first test of the new sensor with the control loop active.

The initial tests of the device were troublesome, as can be expected from initial tests. The errors were quite large, and different between the two differing emissivities on the same plate. This led to a shotgun approach of improvements intended to improve the performance. Much of this effort was directed at the control loop of temperature, and also to the reduction of noise in the circuit, since the heating/cooling was asymmetric and will cause noise to be rectified. And lo, the performance improved.
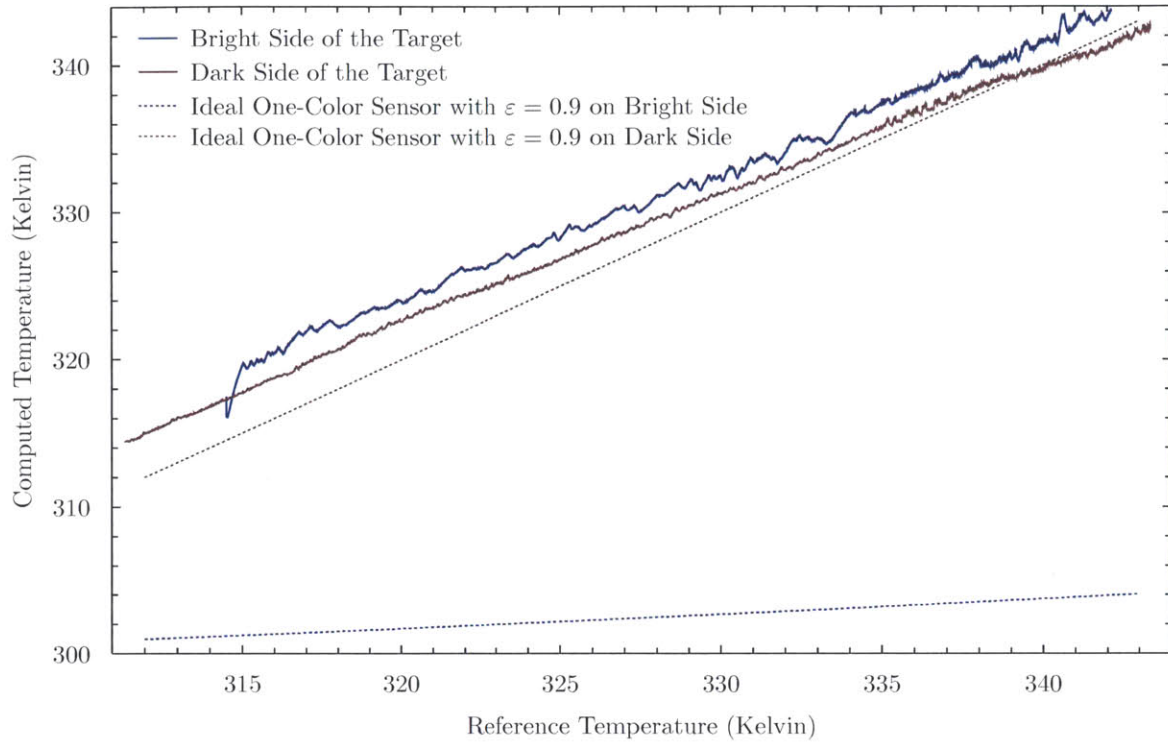
Figure 6-2: The first test of the sensor with the control loop active. Plotted along with the theoretical response given by an ideal one color brightness sensor.
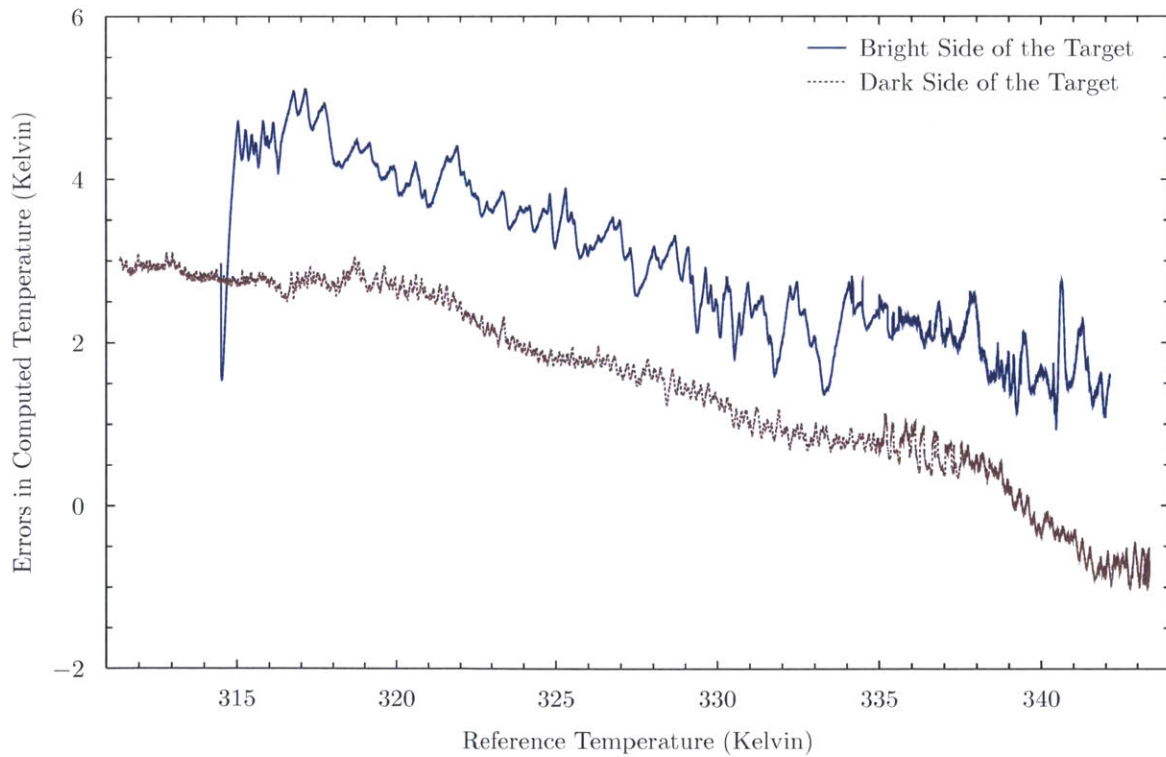

Figure 6-3: Errors from the first test of the sensor with the control loop active.

Chapter 7

Results

The closed loop control system was not yet well tuned during this testing, as can be seen by the large oscillations that developed twice during this test. (See Figure 7-1.)
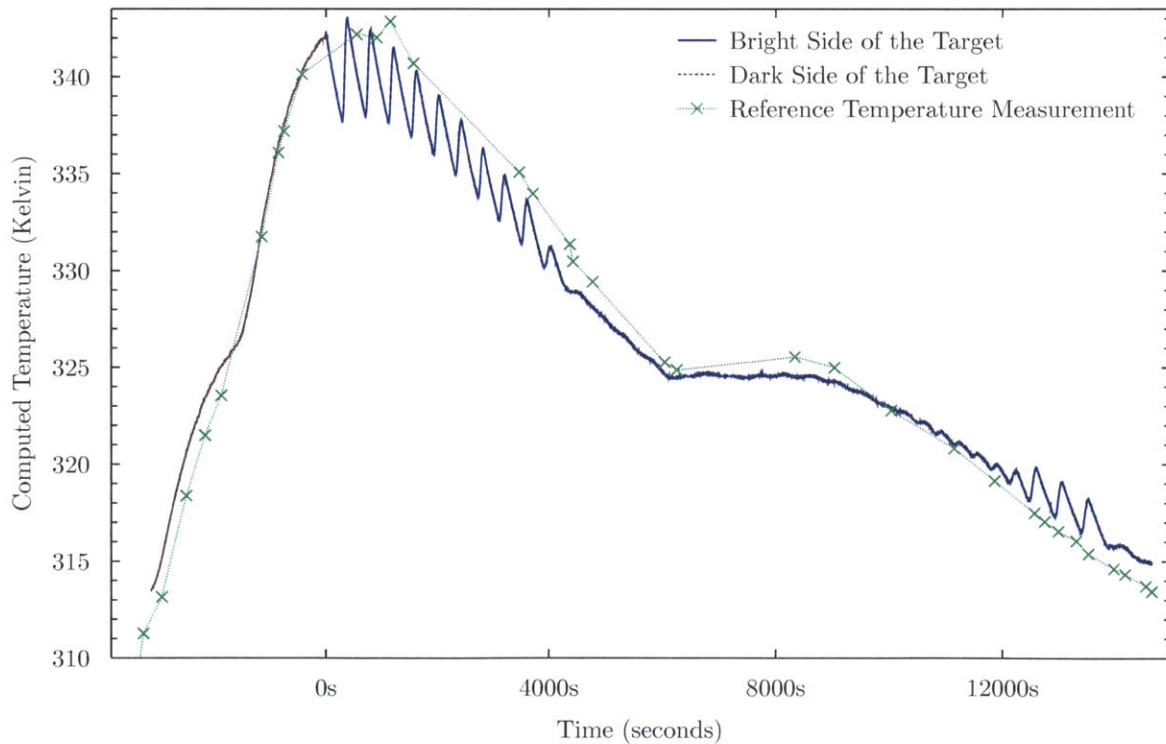

Figure 7-1: Measured output of device versus time, for the duration of one test sweeping over the range of target temperatures.

Part of the difficulty with the control system was the nonlinear response of the system to the control effort. I probably made this worse than it had to be, by my reluctance to do the obvious thing to improve the dynamic response of the thermopile, which was to use a
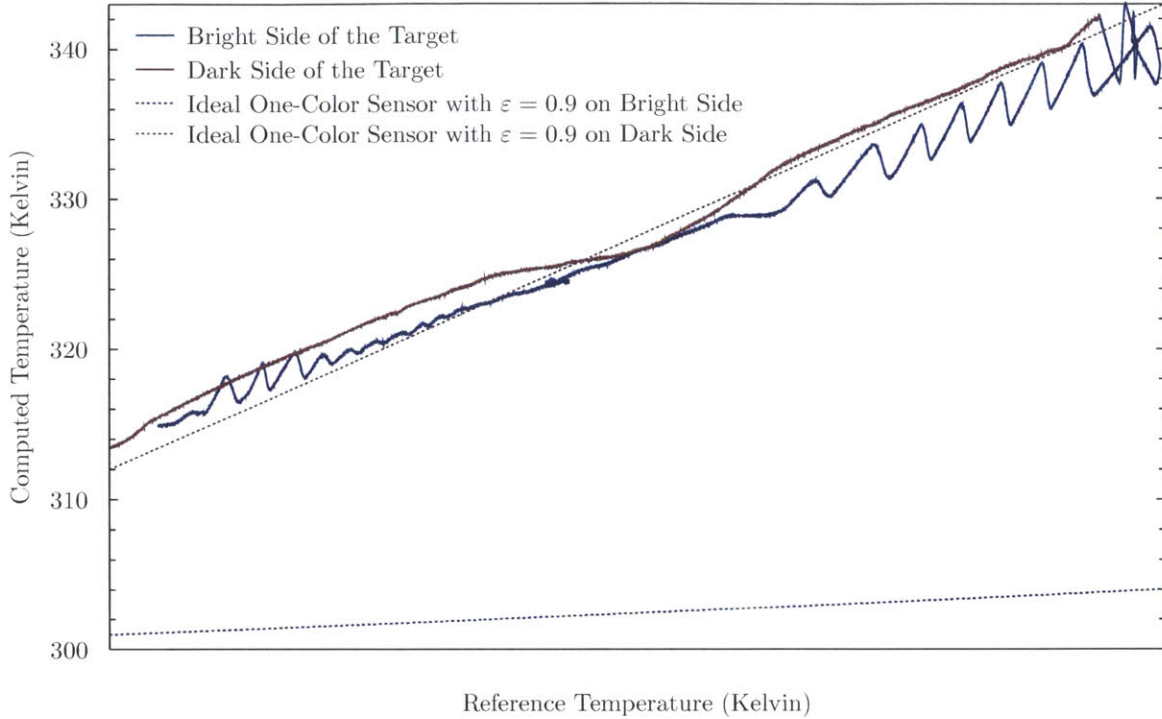
Figure 7-2: Measured output of device, compared to an ideal one color brightness sensor.

metal filled conductive epoxy to permanently pot it in place. The reluctance to do so was rooted in not wanting to chance destroying my one good working sensor assembly.

Another cause of difficulty during this testing, is that the surface of the sensor housing will not actually be uniform, as the design was intended to be used in comparison with a background target plate, with a spacer between the plates to hold them apart. Without the extra material of its half of the the aluminum spacer to heat up, the end of the sensor housing furthest from the thermopile will heat up disproportionately, causing thermal gradients.

I am seeing a much larger than expected (and worse, *variable*) offset on the sensor amplifier. As it was not noticed on the bench prior to device integration, my biggest suspicion is that this is due to a mismatch between the impedances of the thermopile, and the zero reference signal on the amplifier. It may be possible to do a first order compensation for this, if it turns out to be repeatable. Although much better of course, would be to minimize any such effects.

I measured the transmission of the stretch wrap purchased for testing as 0.904. I found a reference for the reflection and index of refraction of LLDPE, which puts the reflectance at ≈ 0.055, and the refractive index at ≈ 1.42.[1] This is at near-IR wavelengths, and not

---

[1]Plass, M., Streck, R., Nieto, J. & Siesler, HW Rheo-optical FT-IR spectroscopy of LLDPE: effect of
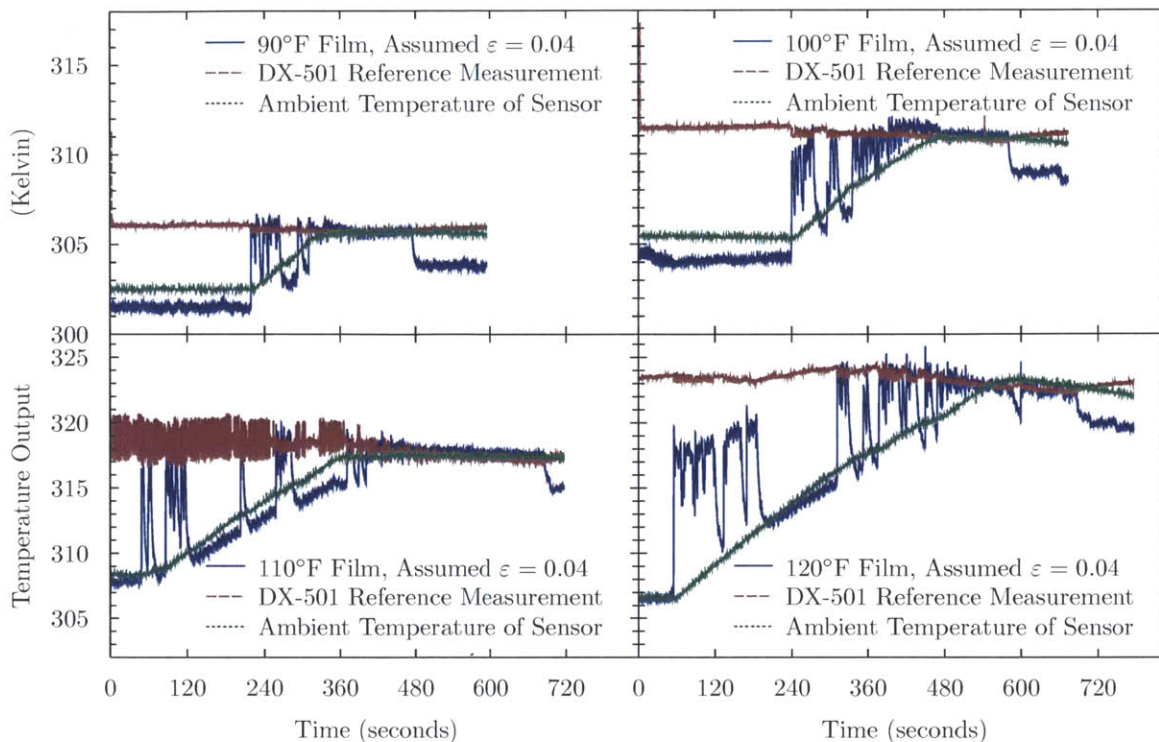
Figure 7-3: Graphs of the sensor servo-ing in vs time at four different film temperatures.

actually the mid-to-far wavelengths I'm using, but beggars cannot be choosers, and I'll use these figures as representative. The values for refraction and reflection agree well for the expected values in air for the normal direction, as given by the equation:

$$\rho = \left(\frac{n_1 - n_2}{n_1 + n_2}\right)^2$$

where ρ is the reflectance and n_x is the index of refraction for the two materials that make up the interface where reflections are occurring. This equation (using a value of $n = 1$ for air) gives a ρ of 0.0301, which if you account for internal reflections off of both faces and assume that there is no adsorption in the film (which is not true, but this calculation will then give an upper bound on the reflection), results in an expected total reflectance of under 0.0584, which agrees well with the $\approx$0.055 value published. Using the published number as the total reflectance of the film, along with the measured transmission of 0.904, gives an emissivity estimate of $\sim$0.04 for the film.

---

comonomer and composite materials. Macromolecular Symposia volume 265 "Modern polymer spectroscopy selected contributions from the conference : "17th European Symposium on Polymer Spectroscopy (ESOP 17)" Seggauberg, Austria, September 9-12, 2007", edited by Peter Wilhelm. Page 130.
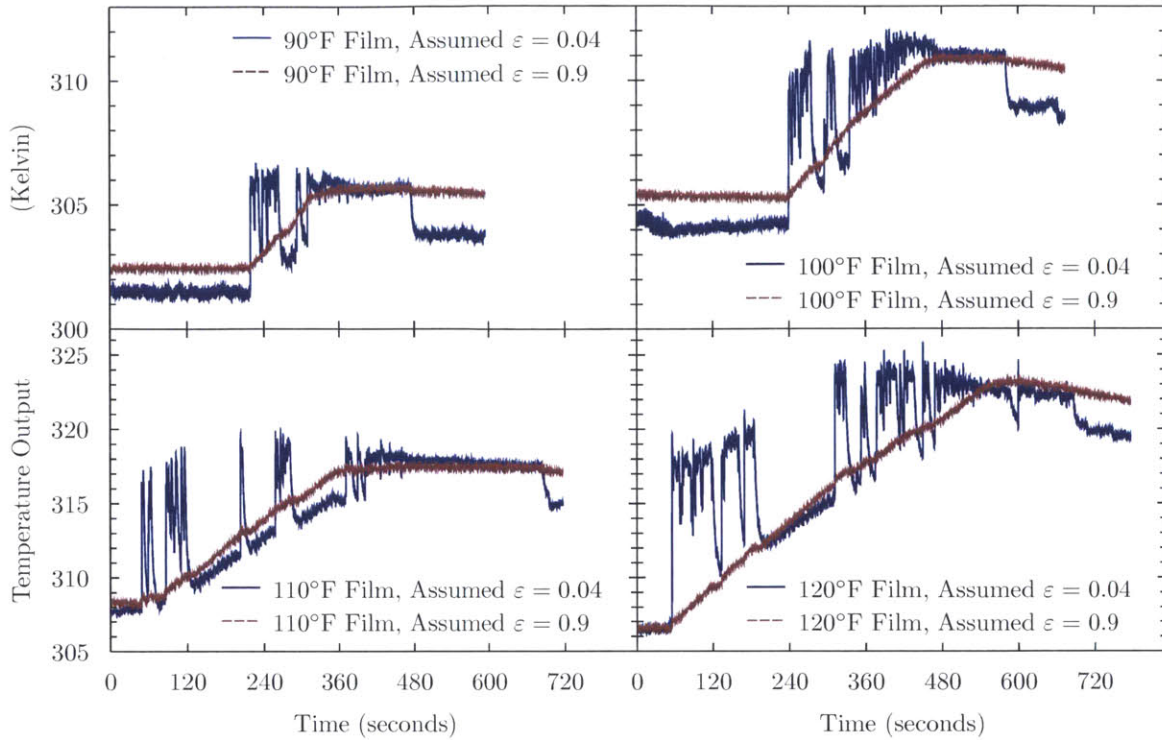
Figure 7-4: Graphs of the sensor outputs over time using two different emissivities to compute temperature as the sensor servos in to a zero heat flux, at four different film temperatures. Note that in all four cases, the computed temperatures agree well when the sensor has achieved steady state.

Figure 7-5 is a graph of the temperature calculated by the sensor for the film under test at the four different temperatures (90°F, 100°F, 110°F, and 120°F) at which the experiments were run. The data for each point is the average temperature computed for the period after the sensor's transient response had settled down, and just prior to the end of the experiment.

In the plot of the temperatures calculated by the new sensor shown in Figure 7-5, the sensor output is in very good agreement with the reference temperature measurement. In fact it is actually kind of hard to see the errors, so I present the graph in Figure 7-6 to make them a little more clear.
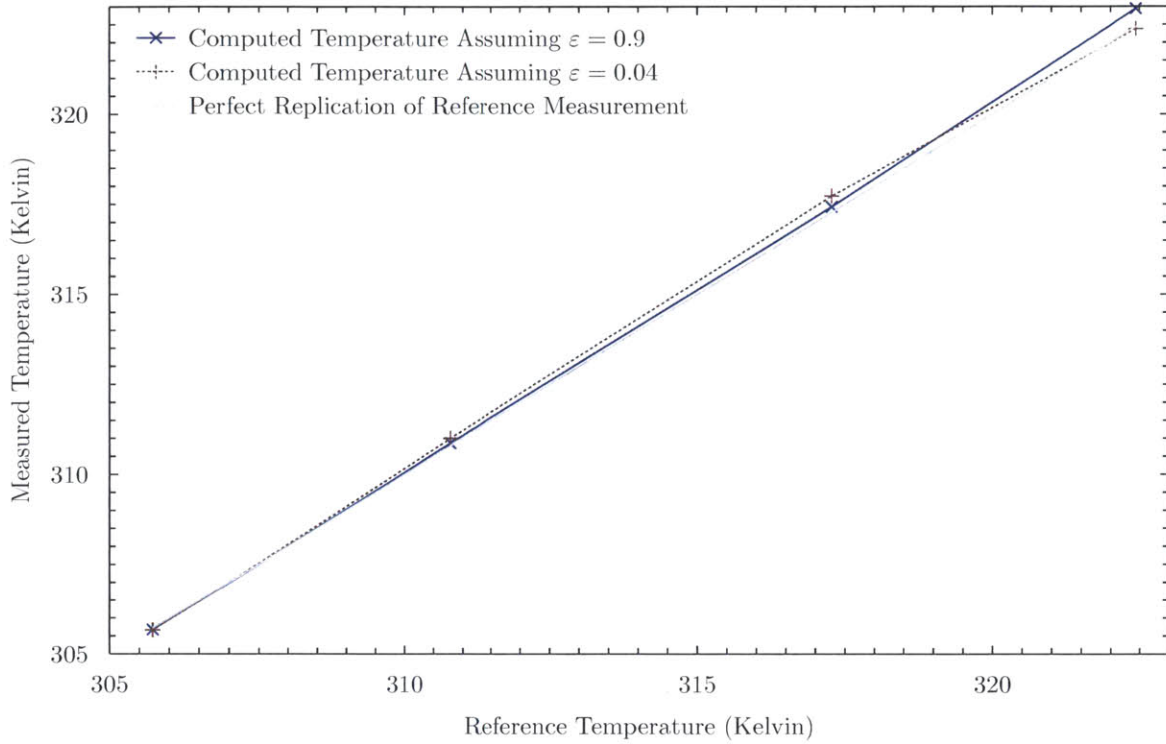
Figure 7-5: Measured output of device, compared to the reference measurements.
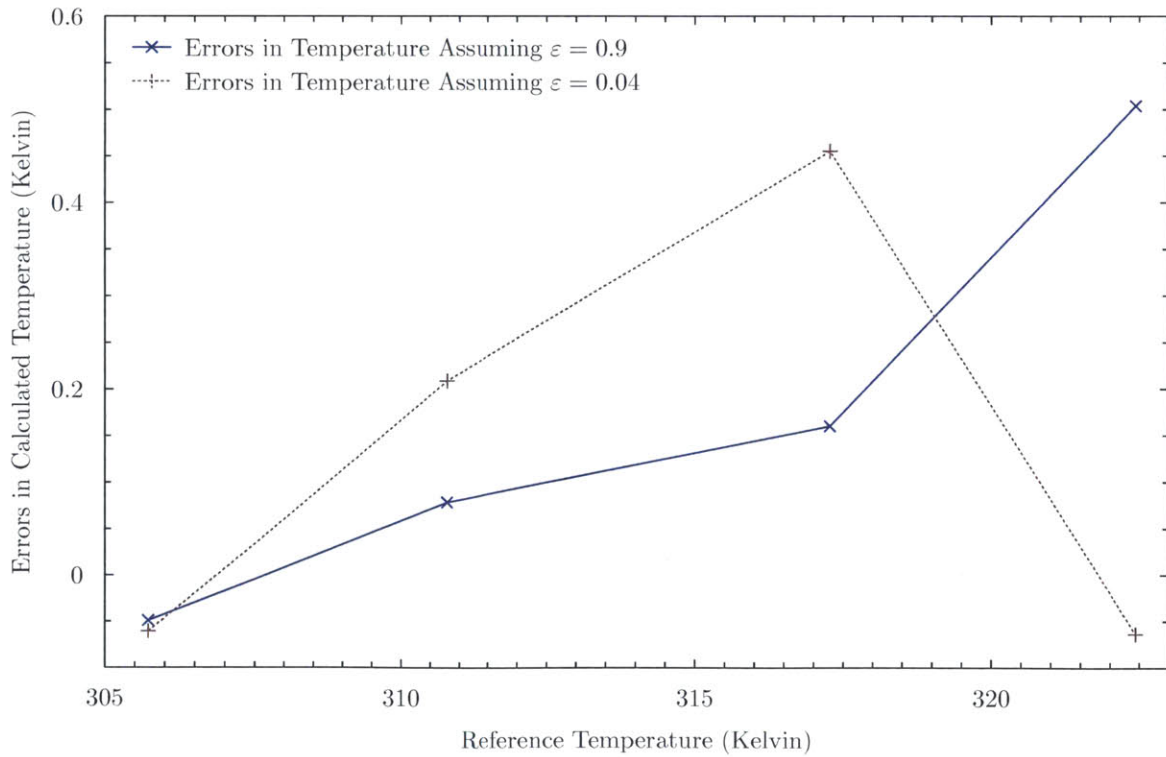

Figure 7-6: Errors in the measured output of device, with respect to the reference measurement.

7.1 Commentary/Conclusions

Please note the magnitude of the error shown in Figure 7-5 and Figure 7-6; Although the reference device used to measure the temperature of the film was given every advantage I could think of to make an accurate measurement, the errors in the temperatures calculated by the disappearing film sensor are too small to characterize by the equipment used in these experiments.

The achieved noise levels were not rigorously quantified, but seem to be lower than my goal of 1°C RMS at an emissivity of 0.01, based on eyeballing the RMS noise to be something like 0.25°C for the film with an emissivity of 0.04, and that's with the actual and very low emissivity in the calculations. Using a higher value will make the readings less noisy, although the resulting output is less accurate when the device is not at steady state.

Further work

I must admit I'm a little disappointed with the performance of the circuit when hooked up to a actual sensor in a real application sort of way. The noise is higher than I'd hoped, and I feel that the quality of the null signal, while good, could be improved. The source of the noise needs to be tracked down, whether it is thermal noise and gradients in the thermopile, thermal EMF effects (Seebeck, Johnson, Kelvin) in the conductors, strain on the amplifier PCB, insufficient EMI rejection/shielding, mismatched impedances between the sensor and the null reference for the autozeroing, or what. The overall performance of the sensor seems pretty good, but I would really like to have seen another factor of three or four improvement. Which I think is probably possible, but will require some more time and effort.

I wonder if the "popcorn noise" could be due to mechanical shifty things, and could be lessened by making the PCB more compliant (or more rigid, or the mounting more compliant or more rigid, or making slots in the board, or making the SMD pads a different shape, or controlling the temperature of the amplifier, or annealing the PCB, or making ...).

Appendix A

Software used

Free Software:

- Emacs

- Latex

- Octave

- pyxplot

- Inkscape

- g3data

- gnumeric

- gerbv

- gmsh

- Elmer

Proprietary Software:

- Eagle (I tried gEDA and PCB, and while usable it wasn't Doing It for me.)

- AutocAD LT

- LabvVIEW

- Excel

- LTspice

- An exceedingly old Bytecraft MPC compiler (And DOSBOX)

- Microchip's MPLAB

- Excel