DEVELOPMENT AND IMPLEMENTATION OF AN FPGA-BASED CONTROL SYSTEM FOR A PASSIVE, DISTRIBUTED APERTURE MILLIMETER-WAVE IMAGING SYSTEM

by

James Lawrence Bonnett

A thesis submitted to the Faculty of the University of Delaware in partial fulfillment of the requirements for the degree of Master of Science in Electrical and Computer Engineering

Spring 2015

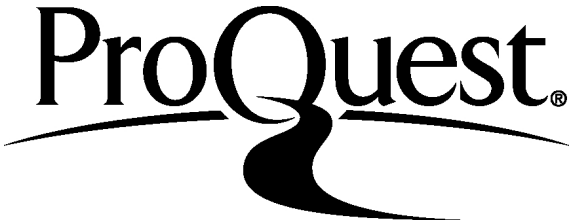

DEVELOPMENT AND IMPLEMENTATION OF AN FPGA-BASED CONTROL SYSTEM FOR A PASSIVE, DISTRIBUTED APERTURE MILLIMETER-WAVE IMAGING SYSTEM

by

James Bonnett

Approved: _____

Dennis Prather, Ph.D.
Professor in charge of thesis on behalf of the Advisory Committee

Approved: _____

Kenneth Barner, Ph.D.
Chair of the Department of Electrical and Computer Engineering

Approved: _____

Babatunde A. Ogunnaike, Ph.D.
Dean of the College of Engineering

Approved: _____

James G. Richards, Ph.D.
Vice Provost for Graduate and Professional Education

ACKNOWLEDGMENTS

I would like to thank my wife, Lindsay Bonnett, my parents, Paul and Wendy Bonnett, and the rest of my family for their constant support and encouragement to push the boundaries of my abilities, as well as the occasional push to help me hold my face to the fire.

I also wish to thank all of my friends and colleagues at the University of Delaware, EM Photonics, and Phase Sensitive Innovations. Special thanks to my adviser, Dennis Prather, and my friend and mentor Petersen Curt, without whose patience and direction this work would not have been possible.

TABLE OF CONTENTS

LIST OF FIGURES

ABSTRACT

Degraded visual environments, particularly those caused by brownout or whiteout, pose one of the most prominent threats to rotary-wing aircraft operating in areas with unimproved landing zones, accounting for nearly half of the losses of rotary-wing aircraft experienced by the Air Force, and providing the leading cause of those suffered by the Army. In an effort to mitigate this threat, Phase Sensitive Innovations has been developing a passive millimeter-wave sensor with the ability to image through the obscurant clouds causing the degraded visual environment. Novel imaging systems such as these likewise require novel control methodologies. In conjunction with PSI, EM Photonics, and the University of Delaware, my research has been focused on developing and implementing such a system.

Chapter 1

INTRODUCTION

1.1 Motivation

Much attention has been focused in recent years on the issue of degraded visual environments (DVE) and safe navigation thereof. Of particular interest is the issue of brownout, or of DVE experienced by aircraft which is caused by clouds of sand, dust, or similar particulates. This is a common obstacle encountered by rotary-wing aircraft when landing in unimproved areas, especially in sandy areas such as the desert. One can easily imagine the challenges such an environment would pose as the downwash of an approaching helicopter transforms a sandy landing zone into a sandstorm; all visual reference quickly disappears, leaving the pilot to fly blind. The problem compounds if we consider that this aircraft may be only the first of a convoy, leaving those that follow without a visual reference not only to the ground but also to any aircraft which have already landed. With this situation in mind, it is not difficult to understand why degraded visual environments, particularly brownout, are currently the single largest contributor to military rotary-wing losses[1]. In addition, the impact of missions that must be canceled due to poor visibility cannot be quantified. Brownout effects operational scenarios and endangers helicopter crews, as well as those who depend on their support.

Current sensor suites employed by the military on these airframes that allow pilots to see in the dark or through fog are unable to compensate for these environments. As such, new imaging techniques and modalities are being investigated

to mitigate this threat. Among these is a system developed by Phase Sensitive Innovations, Inc (PSI) in conjunction with the University of Delaware and under contract with the Navy. This system is of novel design and operates by means of passive detection of millimeter waves (high-frequency RF waves in the range of 10's to 100's of GHz). This program has resulted in several proof-of-concept systems, as well as a fieldable prototype which has now undergone both ground and flight testing. As will be discussed in greater detail later, such a system is inherently sensitive to phase variations induced by vibration, making operation in a harsh environment or onboard an aerial platform unfeasible without the ability to mitigate these variations, necessitating a sophisticated control system. This thesis presents the development of the electronic control systems developed in conjunction with the imaging systems which provide the necessary control functionality that enables the use and deployment of this technology.

1.2 Contribution

During the course of this project, I operated as the second of a two-man electronics development team. The other engineer on this team was Petersen Curt, Senior Engineer at EM Photonics, who was the lead engineer on this project. The design methodology was developed by EM Photonics and has been patented under publication number US 8897656 B2[2]. Our team designed, implemented, and integrated several revisions of the control system based on this design pertaining to developmental prototypes built by PSI. The effort required involved a wide range of disciplines, including analog circuit design, PCB layout, firmware and software generation, and no small amount of troubleshooting.

My contributions have been further detailed in the following publications:

1. Thomas E. Dillon ; Christopher A. Schuetz ; Richard D. Martin ; Daniel G. Mackrides ; Petersen F. Curt ; **James Bonnett** ; Dennis Prather " Nonmechanical beam steering using optical phased arrays ", Proc. SPIE 8184, Unmanned/Unattended Sensors and Sensor Networks VIII, 81840F (October 04, 2011); doi:10.1117/12.898356; http://dx.doi.org/10.1117/12.898356

2. Martin, Richard, Christopher Schuetz, Thomas Dillon, Daniel Mackrides, Peng Yao, Kevin Shreve, Charles Harrity, Alicia Zablocki, Brock Overmiller, Petersen Curt, **James Bonnett**, Andrew Wright, John Wilson, Shouyaun Shi, and Dennis Prather "Optical Up-conversion Enables Capture of Millimeter-wave Video with an IR Camera." SPIE Newsroom. 13 Aug. 2012. Web. <http://spie.org/x89063.xml>.

3. Petersen F. Curt ; **James Bonnett** ; Christopher A. Schuetz and Richard D. Martin " Embedded electronics for a video-rate distributed aperture passive millimeter-wave imager ", Proc. SPIE 8715, Passive and Active Millimeter-Wave Imaging XVI, 871508 (May 31, 2013); doi:10.1117/12.2018519; http://dx.doi.org/10.1117/12.2018519

4. Christopher Schuetz ; Richard Martin ; Thomas Dillon ; Peng Yao ; Daniel Mackrides ; Charles Harrity ; Alicia Zablocki ; Kevin Shreve ; **James Bonnett** ; Petersen Curt ; Dennis Prather " Realization of a video-rate distributed aperture millimeter-wave imaging system using optical upconversion ", Proc. SPIE 8715, Passive and Active Millimeter-Wave Imaging XVI, 87150I (May 31, 2013); doi:10.1117/12.2016138; http://dx.doi.org/10.1117/12.2016138

Chapter 2

THE PSI IMAGING SYSTEM

To understand the requirements of the electronic control system developed in this project, it is necessary to understand the purpose for which it was designed and likewise to understand the basic concepts behind the imaging system developed by Phase Sensitive Innovations (PSI). PSI's imager operates in the millimeter-wave region of the electromagnetic spectrum, so denoted because the physical wavelength of radiation in this area of the spectrum exists on a range of about 1mm to 1cm. By contrast, other common sources of electromagnetic radiation encountered in the course of daily life range from FM radio waves, as long as 3 or 4 meters, to the so-called 'optical' wavelengths that our eyes are sensitive to, measuring mere hundreds of nanometers, a billion times smaller; between these two common examples sits a wide range of other frequency regimes, each suited to different applications. Millimeter waves in particular are of special interest for a variety of reasons, primarily in relation to their properties regarding three particular physical phenomena, namely scattering, blackbody radiation, propagation, and diffraction, each of which will be addressed in limited detail in the following few paragraphs.

2.1 Scattering

In effect, scattering is at the heart of the brownout issue, and thus I address it first. When an electromagnetic wave, or rather a wave of any type, comes into contact with an object, its reaction with that object depends quite strongly on the ratio of the wavelength of the light to the size of the particle. A very long wave that is incident on

a very small particle will hardly scatter at all (e.g. a typical RADAR signal would hardly be affected by a small object like a baseball), while a much shorter wave would scatter strongly (e.g. optical wavelengths scatter specularly off of a baseball, which is what allows us to see it). Here, then, we can see the inherent physical issue taking place during brownout in that the optical wavelengths with which we see are quite short in relation to the size of the sand and dust particles of the obscurant cloud, and are thus scattered entirely. While visible light ranges from approximately 400 to 700 nanometers, the smallest dust and sand particles are merely a few microns in size, and some are even larger. Even the FLIR (forward-looking-infrared) optics already employed aboard the airframes in question are unable to penetrate the dust cloud, despite a slightly longer wavelength. Attempts have been made to apply computational techniques to imagers operating in the visible and infrared regimes with varying amounts of success, but as of yet these wavelengths cannot yield a solution which can see through the dust cloud. Here we encounter the first advantage of millimeter-waves as they are long enough to pass through the obscurant cloud with little scattering, and thus minimal attenuation. Why stop here, then? If longer wavelengths bring better penetration performance, why not utilize even longer wavelengths? The answer to this basic question lies within the issue of diffraction, which will be addressed in a moment. First we must consider the implications of blackbody radiation.

2.2 Blackbody Radiation

Active imagers must first illuminate their object before imaging it, like a RADAR system, whereas a passive imager operates using the ambient illumination already present in the scene. In the context of the PSI imager, this illumination is

created by blackbody radiation. Blackbody radiation is the emission of electromagnetic radiation as a function of heat. This is commonly seen in daily practice as a hot stove will often glow red when it becomes very hot. Cooler objects also emit blackbody radiation, though we may not be able to see it with the naked eye. As the heat of an object increases, so does the energy, and consequently the frequency, of the emitted radiation. Very hot objects, like a hot stove, can emit at visible frequencies, while people, with a typical core temperature of 98.6 degrees Fahrenheit, radiate in the infrared regime. As we consider cooler and cooler objects we encounter the comparatively cool range of terrestrial blackbody radiation.

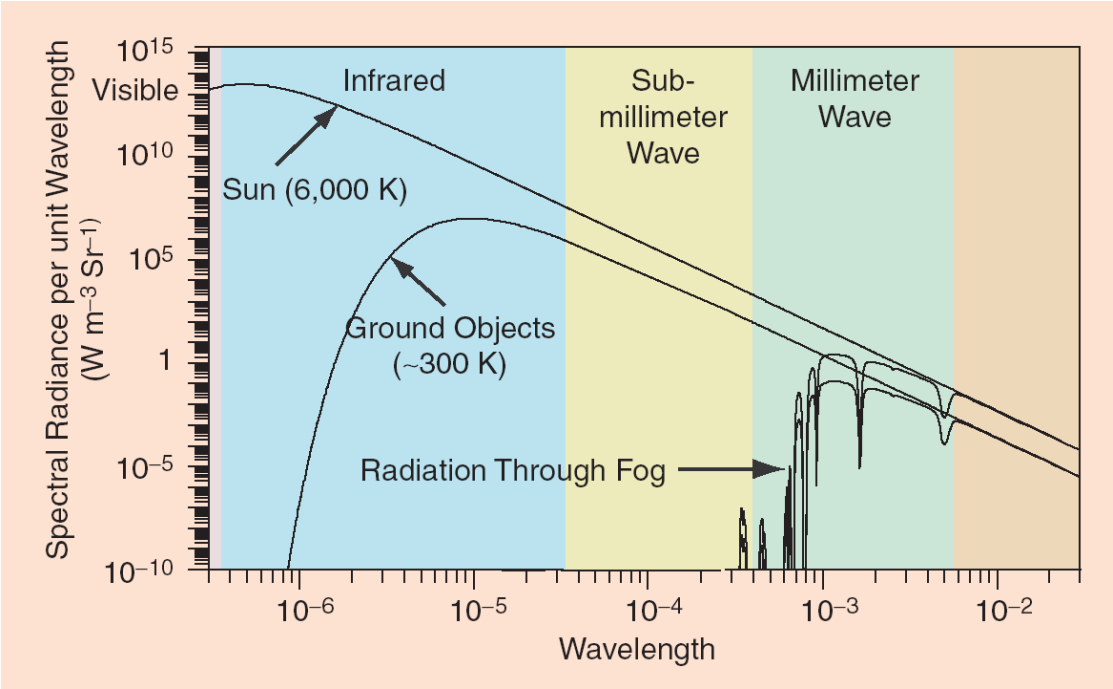

Figure 2.1 Passive blackbody radiance through 1km fog[3]

Now we must consider the combination of blackbody radiation with the discussion of scattering above. As shown in Figure 2.1 above, the radiation spectrum

of ground objects follows that of the sun fairly closely up to the infrared range, reaching a peak. Because of this, a passive imaging system operating around this peak would exhibit better performance, but we can also see that radiation through fog drops off dramatically in the lower millimeter-wave regime, making the infrared portion of the chart useless in the context of degraded visual environments. Here we can see quantitatively why existing imaging modalities that make use of this portion of the spectrum fall short in these situations. Given the combination of these factors, millimeter-waves present an attractive solution in terms of both scattering and blackbody radiation.

2.3 Propagation

In addition to the previously mentioned properties regarding scattering and blackbody radiation, millimeter waves have a significant ability to penetrate the atmosphere. As can be surmised, such a quality is absolutely essential in the development of imaging modalities expected to operate at any appreciable range. Loss in propagation through the atmosphere is highly dependent on frequency. Much as radiation from a variety of frequencies may be scattered by obscurants in the air, attenuation of electromagnetic radiation is highly frequency dependent, as shown in figure 2.2 below. Of particular interest are the local minima shown, which correspond to frequencies which are not attenuated to the extent that others are. As can be seen, several frequencies of interest can be found in both the Q-band (33-50GHz) and W-band (75-110GHz). As such, these are the operation frequencies of the various iterations of the PSI imager.

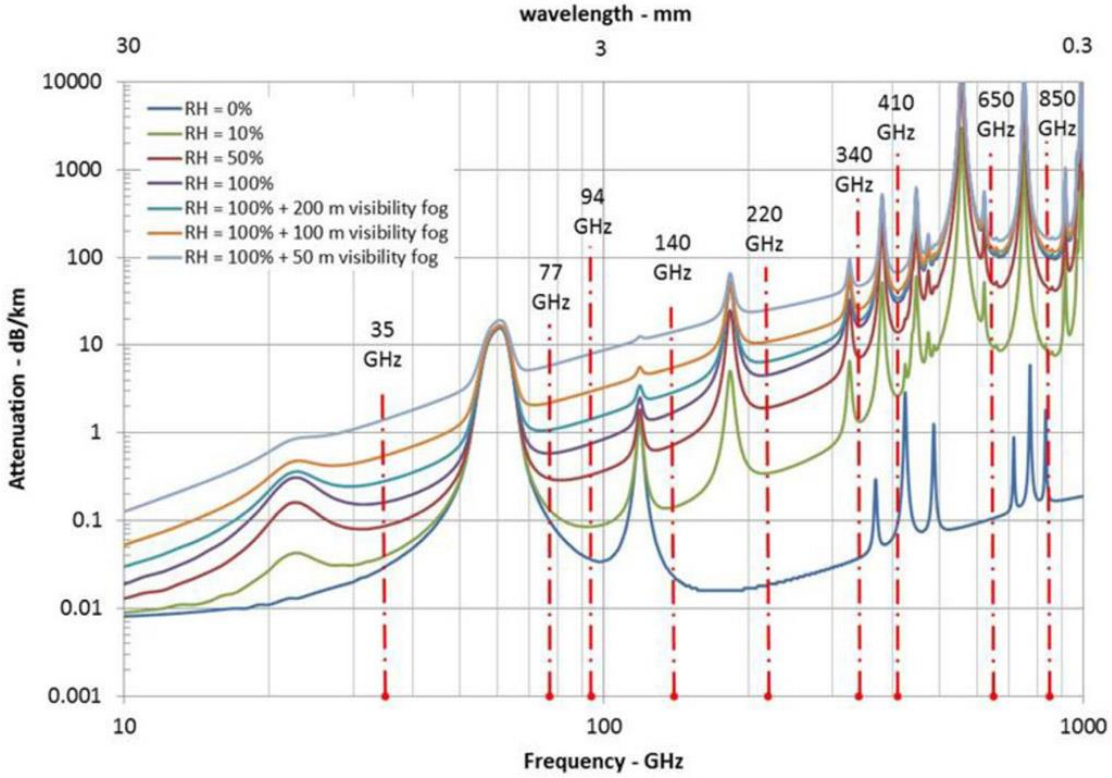

Figure 2.2 Graphical representation of atmospheric attenuation of electromagnetic radiation with relation to frequency. Local minima are noted, and of particular interest for imaging modalities. This plot was generated by PSI using atmospheric codes developed for the North Atlantic Treaty Organization.[4]

2.4 Diffraction

The resolution of any imaging system is limited by diffraction. Even a 'perfect' lens cannot achieve infinite resolution unless it is infinite in extent due to the fundamental physical nature of light. As such, the resolution of an imaging system is typically described by the minimum angle discernable between the focal point of the lens and resolvable points in the image. In the common case of a focal-plane array with a circular lens, this minimum angular resolution is effectively estimated by the Rayleigh Criterion, which defines it as being directly proportional to the wavelength in

question and inversely proportional to the aperture diameter as shown by the following equation:

$$\alpha \approx 1.22 \frac{\lambda}{D}$$

It follows, then, that an imaging system utilizing long wavelengths is at an immediate disadvantage when it comes to resolution, as the aperture must grow in direct proportion to the wavelength to achieve the same level of resolution. This sounds straightforward enough, but only until we consider the sheer extent of the electromagnetic spectrum. Visible light ranging from 400-700nm can be used to capture high-definition imagery with devices small enough to fit inside a cell phone, but a similar device utilizing millimeter-waves would have to be over a thousand times larger to compensate for the difference in wavelength. Such a system would no doubt prove useful for some applications, but would be inappropriate for our purposes of brownout mitigation aboard rotary-wing aircraft.

Much of the diffraction limitation we see here is imposed by the system topology. As mentioned above, most cameras and imaging systems of similar modalities make use of a combination of optical lenses and a focal plane array, but implications of the diffraction limit inherent to such a system topology make it prohibitive to use to image millimeter-waves with an appreciable level of resolution. Systems of this type do indeed exist, but their size and weight make them ill-suited for use on a weight-sensitive aerial platform, particularly rotary-wing aircraft, and as such they are relegated to use on ground platforms or on very large aerial platforms.

In light of this, PSI eschewed the conventional configuration, opting to design their imaging system around a distributed aperture. Where a conventional aperture

would be defined by a single focal-plane array and any relevant lenses, a distributed aperture is comprised by many discrete antennae which are spatially dispersed. A high-profile example of such a system can be seen in use on-board the Lockheed Martin F-35 Lightning II, or Joint Strike Fighter. This system, the Northrop Grumman AN/AAQ-37 DAS (Distributed Aperture System)[5], consists of an array of high-resolution electro-optical infrared sensors, the combined imagery of which affords the pilot an unobstructed 4π steradian view of his surroundings, effectively allowing him to see through the hull of the aircraft.

In the case of the PSI Distributed Aperture Imager (DAI), the aperture is defined by multiple discrete antennae distributed over the front surface of the imager. Each of these distributed channels consists of a horn antenna, a chain of low-noise amplifiers, and a phase modulator. These channels would appear to be analogous to the pixels of a conventional focal-place array, but this is not the case. Rather, they would be better described as components of a wavefront sensor, collecting information relating to the spatial frequencies of incident millimeter waves from which an image can be reconstructed using interferometric imaging techniques. The employment of this approach allows the DAI to achieve the performance of a much larger imaging system without the additional size and weight.

Crucial to the operation of the PSI DAI is the upconversion modulator, which operates by means of frequency modulation in similar fashion to a heterodyne mixer. In typical applications of frequency modulation, ie an FM radio, audio signals (ranging from 10Hz to 14kHz) are modulated onto a carrier of much higher frequency (88 to 107MHz) by varying the instantaneous frequency of the carrier wave. This familiar process is governed by the equation:

$$\cos(a) * \cos(b) = \cos(a + b) + \cos(a - b)$$

If ‘a’ is taken to be the carrier frequency and ‘b’ is taken to be the frequency of the encoded data, we can see that two signals are generated at frequencies of a+b and a-b; these are known as sidebands.

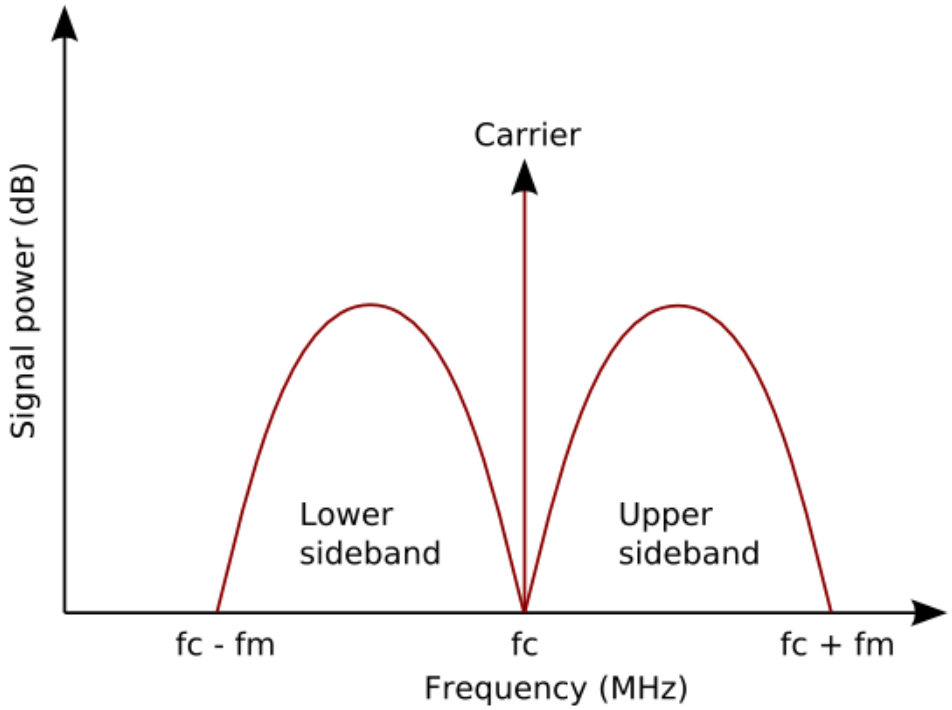

Figure 2.3 A graphical illustration of sidebands and how they appear on the electromagnetic spectrum

In similar fashion, information captured from the millimeter-wave regime, in this case ranging from 35 GHz to 77 GHz, can likewise be encoded on higher-frequency carriers such as short-wave infrared light. The implications of this are numerous, but the primary benefit is that information gathered from millimeter waves can be encoded on infrared light, which can easily be imaged without the debilitating effects of

diffraction seen when attempting to image millimeter waves. This carrier is generated by a ~1550nm laser, a wavelength that is very common in the telecommunication industry, and for which sensor arrays can be easily acquired. Imaging with a distributed aperture requires that magnitude and phase information be saved from each detector in the distributed array. Typically this requires the use of distributed Local Oscillators and mixers to down-convert the relevant data for digital storage, requiring fairly intense computational methods (which themselves call for fairly large and heavy computational equipment) for image reconstruction. This complication is mitigated in PSI's system by the use of the aforementioned optical upconversion and interferometric imaging techniques. A faithful reproduction of the wavefront is created in the infrared regime by coherently recombining the sidebands of the infrared carrier on each channel in free space, allowing the realization of an image by measuring the far field of the wavefront with an IR camera, effectively performing a Fourier Transform. This technique allows the back-end of the system to operate as a common infrared camera, which is an easily-accessible and mature technology, and results in higher resolution than would otherwise be possible with a conventional imaging system with the same name number of sensors. To faithfully reproduce the incident wavefront in the infrared regime, the physical layout of the distributed aperture is mimicked by the output fiber-optics of each respective channel by mounting them in a very small silicon array which was itself developed with a special etching process by PSI. This array is sized so as to be proportional to the distributed aperture, but scaled by wavelength, and so great care must be given to the placement of the fibers and the respective optical path-length (OPL) of each channel. These fibers are then allowed to emit to free space, where the carrier frequency and one of the two sidebands are

eliminated by a narrow-bandwidth optical filter, leaving only the modulated image data. This data can then be recorded by placing a SWIR (Short-Wave IR) camera at the focal plane of the optical system.

The unique combination of a distributed aperture with optical upconversion enables unprecedented millimeter wave imaging performance, but also introduces complication in that each channel must be controlled individually. An added complication of a distributed aperture imager is a phenomenon known as ‘fringe washing’. The formation of an image requires both coherence and phase calibration, which becomes increasingly difficult to maintain as the number of individual channels in the aperture increases. Differences in the path length of each channel correspond directly to differences in phase between these channels, which must be constrained within the associated coherence length relating to the bandwidth of the system in order for an image to be generated. As such, instabilities in these relative phases caused by effects such as vibration pose a considerable threat to system operation when left unchecked. In the context of millimeter waves, coherence is fairly easy to maintain, as the signals are sensitive to perturbations on the order of their wavelengths, which are uncommon. After upconversion, however, the signals are now sensitive to perturbations on the order of their new wavelengths, which are much smaller and thus the threshold for degradation is much lower. As such, very sensitive phase stabilization is required to maintain the fidelity of the infrared signals.

The unique architecture of the imager developed by PSI presents a collection of unique challenges for a prospective control system. Given that such an imaging system is very novel, it follows that a novel control system must be developed for it so that it may be tailored to its needs.

Chapter 3

ELECTRONIC PHASE CONTROL METHODOLOGY

In order to normalize the fluctuations in phase produced by each channel, a control methodology had to be developed which was capable of detecting, calculating, and compensating for the perturbations on any given channel in an appreciable amount of time to allow real-time video display without aberration. This methodology was pioneered at EM Photonics and has been published in a patent held by EM Photonics, Inc. Work completed by EM Photonics on this project was likewise performed under contract with the Navy, who funded its development.

The primary function of the electronic control system is to stabilize the phase of each individual channel so as to enable coherent imaging. All modulators in the system are fed by the same laser, causing them to modulate coherently, at least in theory. Also important to the generation of an image is the relative phase of each channel. We know that a phase shift in the frequency domain corresponds to a shift in the time domain, and so variations of relative phase between sensor channels can have a variety of effects on the final image, from shifting focal distance to completely scrambling the image. In practice, however, this phase relation between channels is frequently disturbed by vibration and movement of the system. While the front end of the system operates in the millimeter-wave regime, and is thus sensitive to vibrations proportional to long wavelengths (a few millimeters), the upconverted signals that can be found in the optical processor are sensitive to perturbations on the order of their own wavelengths, about 1550nm (a common telecommunication wavelength in the

infrared (IR) regime). In this context, and fluctuation in the path of the optical fiber on the order of a micron can cause phase fluctuations on the image contribution of that particular channel, and the phase-coherence is broken.

Since all of the modules are intended to be coherent by virtue of being fed by a single laser source, that being coherent by definition, these perturbations can be detected and measured by comparing a given channel with a reference. As such, an additional reference module is added to the system which receives no RF input. The output of this modulator is then interfered optically with those of the signal carrying modulators, creating an interference pattern from which the phase aberrations on the input channel can be inferred. The mechanism of a phase modulator is to convert an input voltage signal into a variation in phase on a given carrier signal. As more voltage is applied, the variation in phase on the carrier signal becomes more pronounced, until it eventually reaches the 2π radians (360°) point where it wraps back to its original state. Given this repetitive nature of phase, the range of voltages needed to produce a phase shift from 0-2π (generally referred to as $V_{2\pi}$) can accomplish any adjustment required. Thus, assuming a linear modulator, a signal of increasing linear amplitude would produce the same results as a signal which increases to $V_{2\pi}$, returns immediately to 0, and repeats – effectively forming a sawtooth with inflection points at every multiple of $V_{2\pi}$. When such a signal is applied to a phase modulator, the output of which is then interfered with another coherent modulator with what should be a constant phase input, the output interference pattern will present as a classic sine wave. Since the phase inputs on the signal modulators are ostensibly constant, varying only by the phase aberrations inflicted on them by vibration or discrepancies in path length, each individual interference pattern will bear the shape of a shifted sine wave.

This shift is then directly proportional to the phase error present on that particular channel, allowing it to be calculated. This relation is shown in Figure 3.1 below.

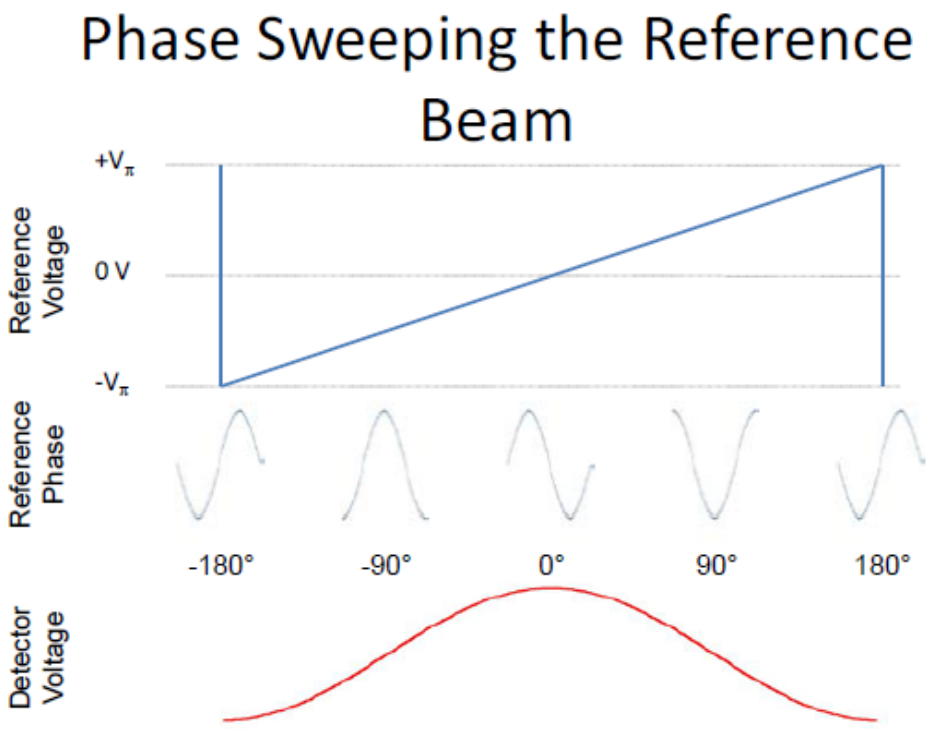

Figure 3.1 Sweeping a reference with a DC voltage allows the creation of an interference pattern which varies at a frequency more easily manageable by the control system

The next challenge is to devise a way in which this phase error can be automatically detected by a digital system. These signals are easily captured due to the abundance of photodetectors developed to function in this particular wavelength range (~1550nm). However, these signals are analog by nature and cannot be understood or acted upon appreciably by a digital system. The classical solution would be to add an

analog-to-digital converter circuit to the input of each channel, but this design is not always feasible or optimal. Analog-to-digital converters are expensive and would in this case add quite a bit of unnecessary complication to system operation. As discussed earlier, channel signals present in the form of shifted sinusoids. In the context of a system input then, a digital representation of the complete signal would be superfluous information, as only the relative phase offset of the given signal with respect to the system reference is of consequence. As such, the methodology developed consists of analog circuitry which comprises what is essentially an edge detector or 1-bit analog-to-digital converter in practice. The original concept of this circuit is shown below.

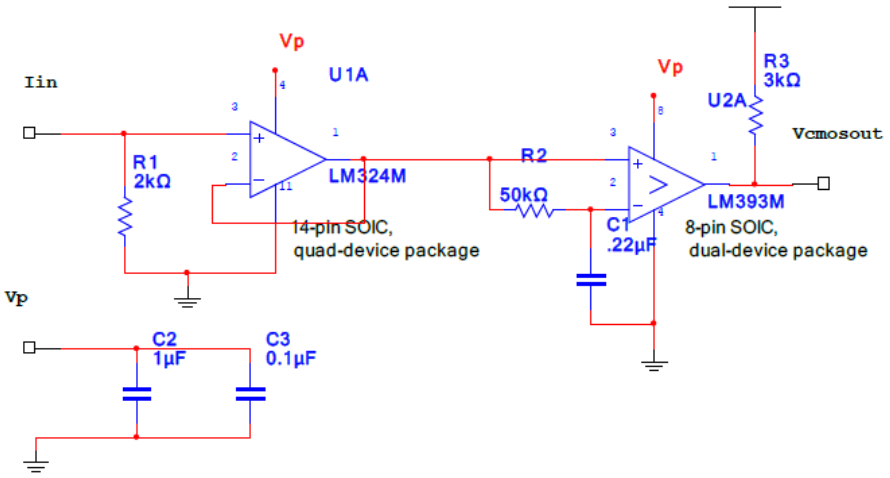

Figure 3.2 Schematic of original interference readout circuit for phase detection

The particulars regarding the operation of the various parts of this circuit will be discussed in more detail in the following sections, but in general the left side of the circuit comprises a trans-impedance amplifier, converting the input current signal from

a photodiode to a voltage signal and amplifying it, and the right side of the circuit comprises the comparator circuit. As can be seen, the output of the trans-impedance amplifier is split and one leg is subjected to a low-pass filter, the cutoff frequency of which is tuned to be lower than the expected frequency of the sinusoidal signal. As such, this leg effectively represents a DC average of the original signal, and is connected to the negative terminal of the comparator. The unfiltered signal is then passed without alteration to the positive terminal of the comparator. When the signal value is less than its average, the output is driven low. When the signal is greater than its average, the output is allowed to float high due to the pull-up resistor attached to the output. In this way, the incident sinusoidal waveform is effectively transformed into a square wave, which can more easily be understood by the FPGA control system. Also, the inflection points of the incident signal are preserved in the generated waveform, allowing for calculation of relative phase offset via measurement of temporal delay of these edges.

As the reference sawtooth signal is likewise generated by the FPGA, the inflection points of each incident signal, indicated by digitized edges generated by the method previously explained, could easily be compared to the corresponding inflection points of the reference signal to calculate a phase offset value for each channel. This offset will then correspond to a given error value, in that each channel will have a desired relative phase offset which the system is programmed to correct to. Phase correction is then performed in a two-step process wherein the system processor computes a compensation factor and this factor is applied to the array by means of specialized circuitry. This compensation factor is calculated by means of a ‘PID’ control technique which is purposed to minimize steady-state error, overshoot, and

settling time. The three letters ‘PID’ correspond to three computed values, namely the Proportional, Integral, and Derivative parameters, corresponding to present, past, and predicted error values, respectively. These three parameters are combined by means of a weighted sum to determine an appropriate error-correction value. A graphical representation of this process is shown below.

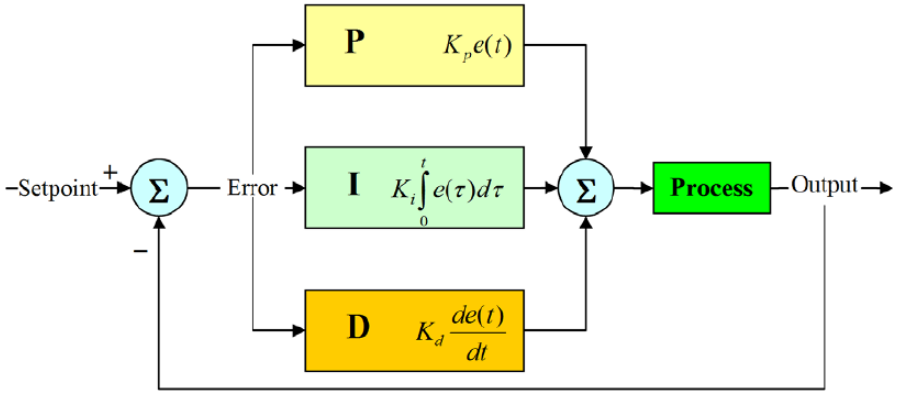

Figure 3.3 Phase error computation process, in which three feedback branches are summed to minimize steady-state error, overshoot, and settling time

Once calculated, the phase compensation signal is output to the phase modulator. The signal must therefore be acted on by a DAC, or digital-to-analog converter, as the modulator is naturally analog and cannot accept digital signals. This process is subsequently repeated multiple times per second for each channel in the system, constituting a closed-loop feedback system in which the proper phase correlation for each channel is constantly maintained.

The various design iterations of this system will be discussed in more detail in the following sections.

Chapter 4

ITERATIONS OF SYSTEM DESIGN

4.1 4-Channel Prototype

I joined this project directly prior to the PSI Imager's Critical Design Review in May of 2010. At this point in time, EM Photonics had already begun development of the control system using the previously detailed phase control methodology. As with most circuitry projects, this particular project began on a solderless breadboard with a rough prototype of the phase detection circuitry; this prototype is shown in the following figure.

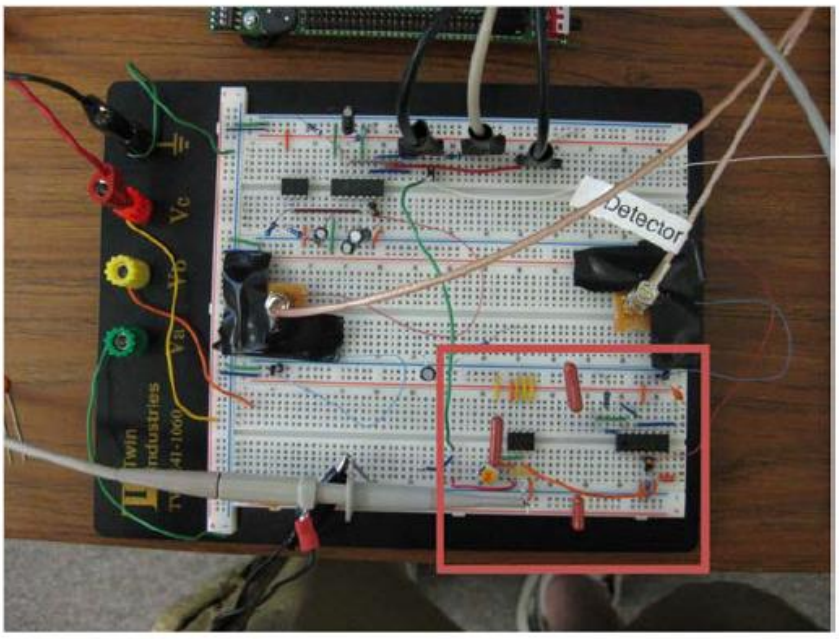

Figure 4.1 Original solderless breadboard prototype of the phase detection circuitry, indicated by the red square. The feedback distribution prototype circuitry can also be seen to the upper left of the figure.

After so proving the capacity for operational functionality, EM Photonics developed a proof-of-concept system of custom design which had the ability to synchronize the phases of four channels. This system consisted of a custom printed circuit board (PCB) and a COTS Xilinx ML-605 FPGA development board. The custom PCB contained the circuitry necessary to digitize each control channel for phase detection, and to generate the reference sawtooth signal. This PCB could be attached to an FPGA development board by means of an FMC (Field Mezzanine Connector) interface. The FPGA 'parent' board could then perform the necessary phase error calculation, determine appropriate feedback compensation, and output it back to the PCB via the same interface, and the onboard circuitry would in turn apply the feedback voltage to the modulator. The development test setup likewise incorporated two phase modulators, one for reference and the other for feedback, with which we were able to replicate the conditions of a channel working in the proposed PSI Imager, demonstrating successful closed-loop feedback control of a phase modulator.

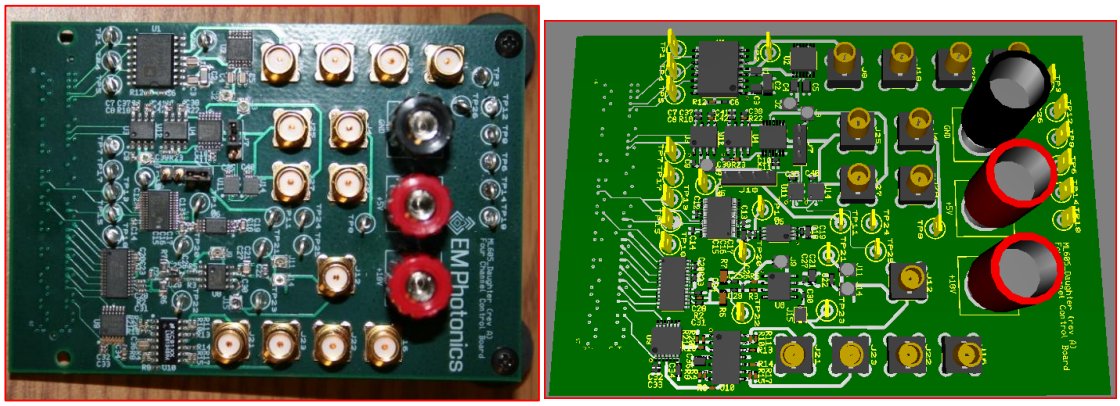

Figure 4.2 A picture of the 4-channel prototype PCB together with a software-generated 3D rendering of the same

This system sampled inputs at 100kHz, or half the expected operational frequency of the subsequent iterations of the control system. Nonetheless this system was able to satisfactorily exhibit close-loop phase control of a phase-control modulator under vibration, maintaining a standard deviation of only 2.2° under no vibration and 6.8° under simulated AH-53E vibration. As this prototype was not intended for system integration, it included no electro-optic detectors for receiving optical signals. Instead, a discrete photodetector was used and its output was injected onto the PCB via test connectors.

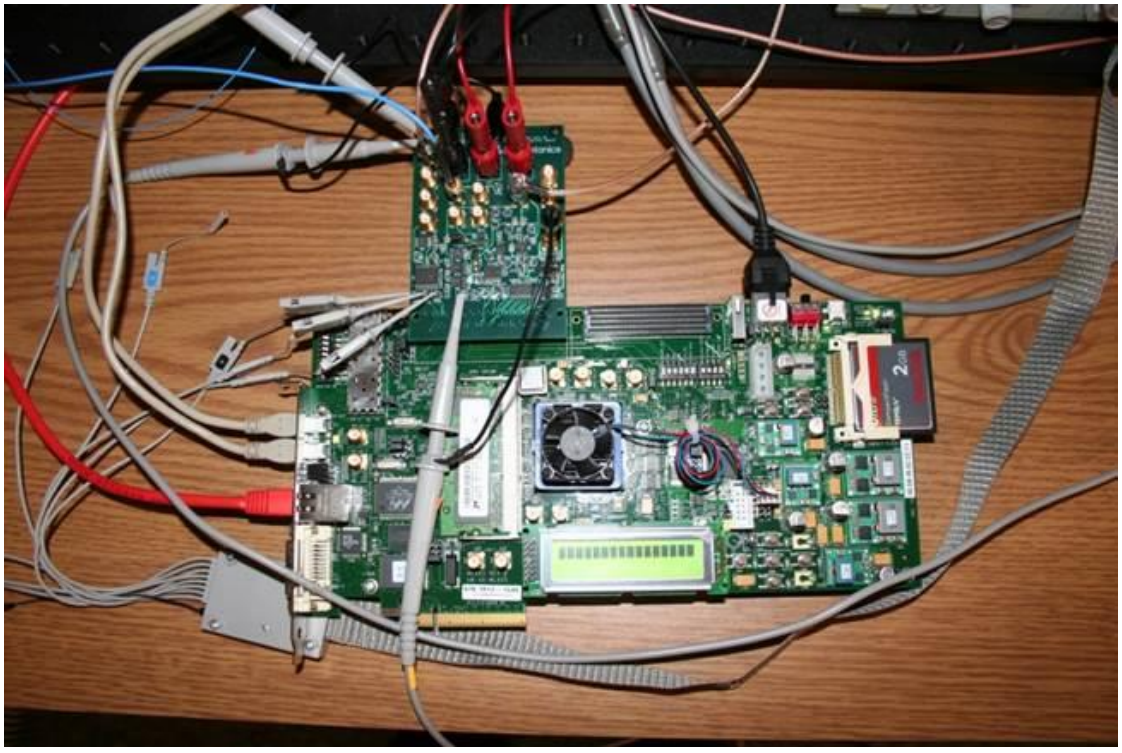

Figure 4.3 The 4-Channel Prototype PCB mated to a Xilinx ML-605 Motherboard under test

4.2 Pathfinder IIB

4.2.1 Overview

PSIs first DAI imager was a thirty-channel array using 35GHz modules designated internally (and subsequently in relevant design literature) as ‘Pathfinder’. The 35GHz operational frequency was chosen due to the availability of low-cost phase modulators developed for use in high-speed communication applications. This system was implemented on a standard optics table within the proper enclosures given the necessary inclusion of an IR laser.

Though known commonly to us as simply ‘Pathfinder’, this moniker actually describes several systems. The first, eventually known as the Pathfinder I, had been implemented with a periodic hexagonal array. The spatial periodicity of the distributed aperture, however, was found to be detrimental due to excessive aliasing and imaging artifacts. As such, PSI developed the next iteration of the system, the Pathfinder IIA, using the same RF parts as the original Pathfinder, but utilizing an aperiodic detector array in the shape of a 5-arm spiral. This new aperiodic sensor placement helped to eliminate aliases and artifacts in the images that had been seen in the previous iteration of the system. This change likewise required PSI to develop new control electronics, as the optical detectors for the phase control system had to follow the same layout as the sensors on the array face. As will soon be described in greater detail, the system designation was again changed, now the Pathfinder IIB, following the integration of our phase control electronics, the development of which is described in the following sections.

4.2.2 System Topology

Following the paradigm of the four-channel prototype board, we developed a system composed of several custom PCBs and a COTS development board which provided an FPGA and several peripherals. Each RF channel assembly contained Low Noise Amplifiers (LNAs), whose power was delivered by RJ-45 connections, and phase feedback was sent over SMA, requiring two connectors to interface with each channel.

As discussed above, the 4-Channel prototype control system was implemented with the use of a Xilinx ML-605 development board. Development boards, like the one in question, are typically used in proof-of-concept or prototypical designs as they provide a considerable variety of pre-implemented interfaces and features, thereby simplifying incorporating designs (and likewise reducing associated costs). As physical space was in this instance a non-issue, we utilized the same model development board as before, replacing our original 4-channel custom PCB with an array of newly-designed custom PCBs, again designed to interface with the development board's FPGA via its mezzanine connector.

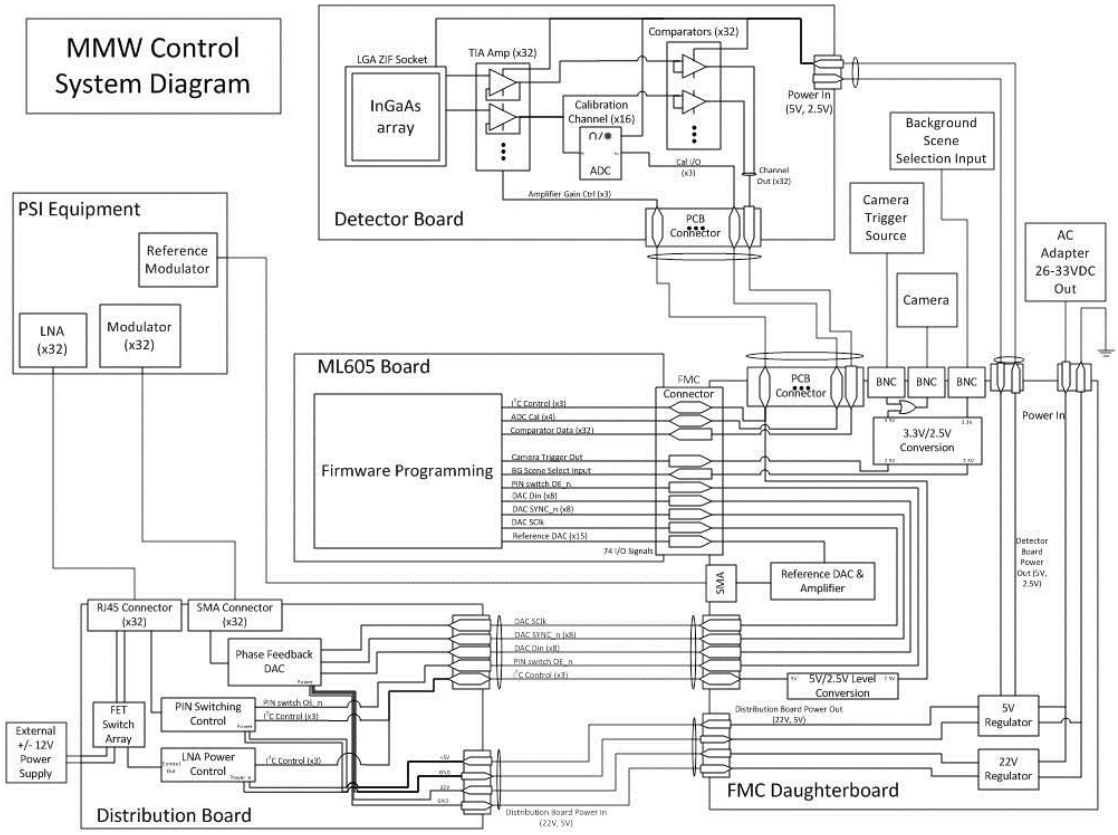

Figure 4.4 Pathfinder Electronics System Diagram.

4.2.3 FMC30

One of my first design tasks on this project was the FMC30 PCB. This PCB, so designated due to its connectivity to the FPGA Mezzanine Connector (FMC) and support of the 30-channel Pathfinder system, closely followed the design of the 4-Channel Prototype FMC board, but was designed for integration into the Pathfinder imaging system. Originally designated the FMC32, the PCB indeed supports 32 channels, two of which are unused as the Pathfinder does not utilize channels 10 and 29. The FMC30 is the primary interface of the FPGA processor to the rest of the system, supporting the interfaces to both peripheral boards, the Receiver30 and the

Distribution30 PCBs, including interfaces for both communication and power. The FMC30 accepts a 26-33VDC input which mimics the AH-53 onboard power format which is then regulated so as to supply 3.3VDC and 2.5VDC to the Receiver30 PCB and 5VDC and 22VDC to the Distribution30 PCB. In addition to power, there is also a 51-pin flat-flex cable interface to both the Receiver30 and Distribution30 PCBs. These cables carried comparator outputs from the Receiver30 PCB and phase correction voltages to the Distribution30, as well as various communication signals such as IIC and SPI busses. The FMC30 also holds the necessary circuitry to generate the phase reference sawtooth signal, which was output to the system's reference modulator via an SMA connector.

The FMC30 likewise handles the trigger for the SWIR camera. The board includes two ports for the trigger, one of which functions as a synchronization input as it may occasionally be desirable to trigger the camera from an external source. Support for various triggering modes was implemented in firmware, selectable from the LabVIEW GUI. A third BNC was also included to monitor the status of PIN switches in place on the Pathfinder II's phase modulators. These PIN switches essentially constituted a 'load' or background setting, which would allow the capture of noise induced by the module. While most noise is non-deterministic and can thus be mitigated by means of averaging, some noise induced by the modules is deterministic, and thus cannot be dealt with so easily. By alternating successive frames between the imaging scene and the 'load' scene, one could easily subtract the background image from the scene image, thus eliminating additive noise and improving the fidelity of the image.

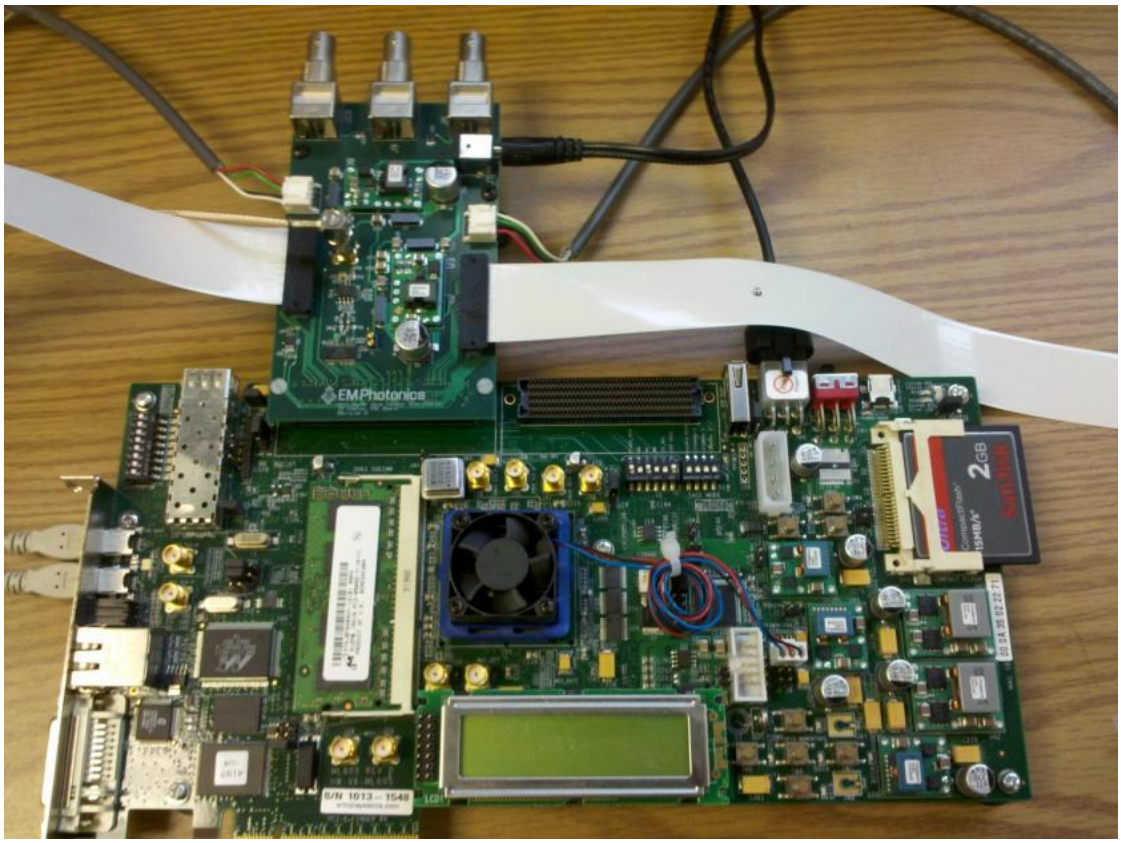

Figure 4.5 The FMC30 PCB attached to its ML605 carrier

4.2.4 Receiver30_RevA and Receiver30_RevB

My largest responsibility in the design of the phase control electronics for the Pathfinder imaging system was the design of the Receiver PCB. While the design of the FMC30 was relatively simple and the layout was a derivative of the previous design, the Receiver board was largely new territory, posing a new challenge as a 'ground-up' design. All subsequent control systems were likewise derivatives of this design. As such, the design of the Receiver30_RevA and the subsequent RevB will be discussed in a higher degree of detail.

4.2.4.1 Optical Interface

The purpose of the Receiver board is to extract phase information from each channel and pass it to the Control board. Since the signals are carried optically, this necessitates splitting the optical signals (such that one branch can travel to the camera and the other to the Receiver board) and converting the optical signals to electrical signals with photodetectors. Since the optical channel configuration follows that of the antenna layout, these photodetectors must likewise mimic the antenna pattern. As a result of this the alignment of the Receiver board and its array of detectors is very sensitive with respect to lateral translation, magnification, and optical focus. The imaging system inside the optical processor is designed such that the short-wave infrared (SWIR) camera images the Fourier plane of the fiber array, necessitating the use of a lens to inflict the requisite far-field diffraction on the fiber array image. By contrast, the Receiver board must be placed in the image plane of the fiber array so that each channel can be individually sampled. The original Pathfinder phase control electronics developed by PSI made use of discrete photodetectors. This allows some simplification in layout, but also required magnification of the imaged fiber array for proper alignment to allow for part-clearance tolerances. This likewise necessitated the addition of more optical components which were not desired. To eliminate these additional optical components, the physical layout on the PCB must be made considerably smaller, requiring tighter design tolerances and making the use of discrete photodetectors prohibitively difficult. The solution determined was to use a focal plane array (FPA), which would be bump-bonded to a copper land pattern on the PCB surface which replicates the array layout. For a variety of reasons, detailed in the following section, we decided to modularize this portion of the design, separating the

FPA interface onto a specially-purposed Detector Interposer PCB, which would be connected to the Receiver PCB by means of a zero-insertion force (ZIF) socket.

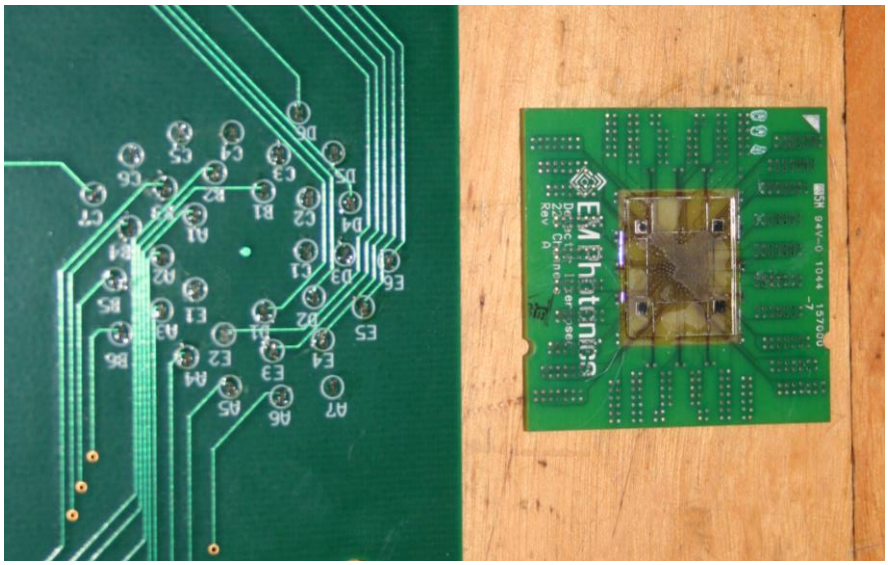

Figure 4.6 Detector Interposer PCB is shown at right, original PSI receiver layout is shown at left. The addition of the interposer removes the need for 40x optical magnification and occupies significantly less space.

4.2.4.2 Input Conditioning

As was discussed in the previous chapter, the input signals from the FPA are relatively weak current signals; the responsivity of the particular FPA in use was approximately 1mA of current per 1mW of optical power. As such, it is necessary to amplify each channel and perform conversion to a voltage signal before attempting to extract phase information. This circuit is shown below in Figure 4.7.

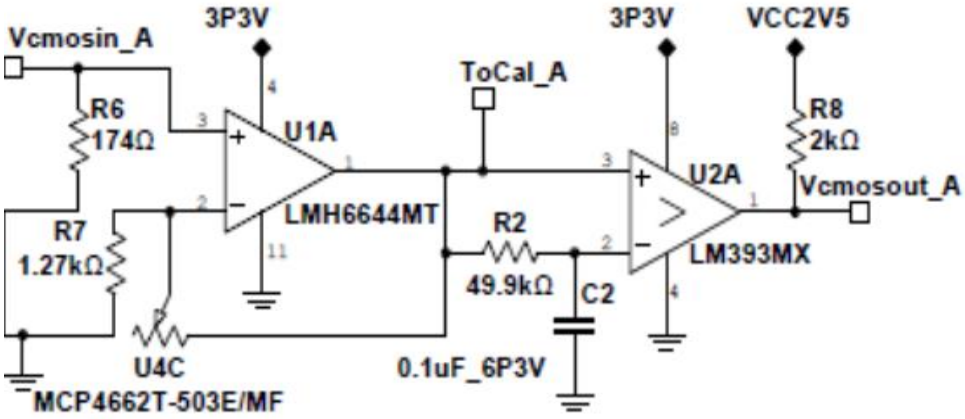

Figure 4.7 Receiver30 Trans-impedance amplifier unit cell, showing input amplification and translation, as well as output digitization

4.2.4.2.1 Current Signal Translation

Op-amps are by nature differential voltage amplifiers and, as such, the input signals carried as variations in electric current must be converted to voltage signals. Within the properties of op-amp design, it is assumed that no current will pass through the inputs of an operational amplifier. The current signal can thus be effectively converted to a voltage signal by the simple relation of Ohm's law (V = IR) through a parallel resistor to ground; as no current can pass through the input of the op-amp, the effective resistance experienced by the input source is determined entirely by this value. As shown by the relation of Ohm's law, the input voltage to the op-amp is directly proportional to this resistance. A consequence of this relation is that this resistor value may be tuned to provide a desired operating range of input voltages in that the resistor value 'R' directly corresponds to resultant voltage 'V'. As such, an expected range of voltage signals will correspond directly with a given range of current signals, and so the system can be adapted to varying input levels with the substitution of only this initial resistor.

4.2.4.2.2 Input Amplification

The input circuit shown in the figure above is largely the same as that used in the proof-of-concept, with a minor change. Due to a variety of parameters relating to everything from the RF modules, the optical processor, and the receiver FPA we anticipated that was likely to be a considerable amount of variance in the optical power from channel to channel, which will relate directly to the signal current from each channel. In order to support this and broaden the range of supportable input signal powers the feedback resistor of the amplifier was replaced with a potentiometer so that the gain on each channel of the system could be individually adjustable. Since such fine adjustments on so many channels would hardly be feasible by hand, our solution was to use an array of digital potentiometers with electrically-controllable wipers with a control methodology built around an I2C bus, a short explanation of which follows. The advantage of the I2C bus is that it is a well-established standard with a wide base of support in COTS devices. This method allows fine-grained user control via a terminal or GUI (graphical user interface), but also lends itself easily to automation in that it can be independently controlled by the processor without requiring interaction from the user.

4.2.4.3 I2C Bus Operation

The I2C standard denotes a single-ended (or non-differential) serial communication interface consisting of a clock signal, designated SCL, and a data signal, designated SDA. A third signal, an active-negative reset is often employed and was included in our design, but is not strictly necessary. A serial bus, such as this, meets our needs here perfectly as the gain control for the amplifiers need not be high speed, and the serial nature of the bus requiring only two conductors adds minimal

complication to the design. The communication on the bus takes place under the assumption that one device is acting as the ‘master’, the FPGA in our case, broadcasting messages to the rest of the devices on the bus, which act as ‘slaves’. Some bus topologies include multiple masters, but it remains true in such cases that only one master can use the bus at a time. Generally, the communication follows the standard that the master will send out the specific address of the slave chip it wishes to address with an additional bit denoting its intention to ‘read’ or to ‘write’, meaning to ascertain the current state of the slave device or to overwrite its current state with new data, which will be met with an acknowledge from the device if connected. In the case of a write, the master will then transmit the memory address of the specific register it wishes to write (generally the address of the wiper value to set in our case) followed by the data to be stored there, each again met with an acknowledge from the targeted device. Otherwise, in the event of a read, the slave will respond with the contents of whatever register was last accessed, thus a read operation generally begins with a write command to the desired register, but without actually overwriting the data contained therein. In either case, assuming the transaction is successful on all points, the device will acknowledge again and the exchange is complete.

4.2.4.4 Bus Multiplexing

Another complication added to our gain control methodology was bus multiplexing. As implied above, devices on the bus are differentiated by means of addressing; in general, a chip operating on an I2C bus has an address which is mostly hard-wired with the last bits determined by physical input pins. As devices of the same model have the same hard-wired address component, the number of devices that can be used on a single bus is determined by how many unique addresses can be made

with the physical address pins. In the case of the digital potentiometers we specified for this design, there were two pins meaning that only four devices of the same model could exist on the same bus, addressed ‘00’, ’01’, ‘10’, and ‘11’, respectively. In order to support all the digital potentiometers required in the design it was necessary to multiplex multiple busses with an I2C switch, each individual bus then supporting four uniquely-addressed potentiometers. This adds a prerequisite step to the bus communication, in that the switch must be addressed first and told which bus is needed before communication can be passed through. Though an added complication, we were able to abstract this in software such that merely indicating a channel index would automatically address the appropriate bus as part of the routine.

This bus multiplexing also achieved a secondary purpose in the way of automatic level conversion. An I2C is connected to a particular source voltage by means of pull-up resistors, such that it need only be driven ‘low’ to transmit data. As such, an I2C bus in a resting state is always sitting ‘high’, where it remains until the master device pulls it low. The FPGA we employed included ports specifically purposed to drive I2C buses at 2.5V, but the digital potentiometers we specified were designed to receive control messages on a 3.3V bus. This brings into account a principle known as ‘hysteresis’, or the necessary conditions for a digital device to recognize a change in state from ‘high’ to ‘low’ or vice-versa. In general, hysteresis requirements dictate the specific percentage of VCC that a signal must achieve to constitute a rising edge condition and the corresponding percentage it must drop below to constitute a falling edge condition. This concept is important because it allows for a certain level of fluctuation or noise on a digital signal without recognizing this as a change of state, however it also means that a 2.5V digital signal cannot be used to

control a 3.3V digital signal if the ‘high’ requirement is greater than ~75% (2.5/3.3), thus necessitating the use of an intermediary level-conversion circuit.

4.2.4.5 Receiver30_RevB Upgrades

After test and verification of the Receiver30_RevA PCB, we identified additional needs and began to formulate a series of upgrades that would be beneficial to integration with the Pathfinder imaging system.

Of particular concern was optical alignment; we had expected this process to be difficult, and found isolation of individual channels to be problematic during testing. As such, we devised a system to alleviate this. Five additional channels were added to the fiber array to assist in the alignment process, these located in each of the four corners of the array and at its center. These channels were fed directly from a laser, and were not subjected to phase modulation. A new block of input-conditioning amplifiers was introduced in the PCB layout, but with no comparators as there was no need for phase extraction. The amplifier outputs, essentially just DC input as no modulation was present, were then fed individually to 8-bit ADC (analog-to-digital converters), the eight bits of which were fed to eight LEDs. As such, during the alignment procedure, the eight LEDs pertaining to any of the four alignment channels would show an 8-bit binary representation of input power, the intent being to provide visual feedback such that the board could be effectively aligned by moving the board until all four channels were maximized. These modifications to the Receiver design are shown below in Figure 4.8.

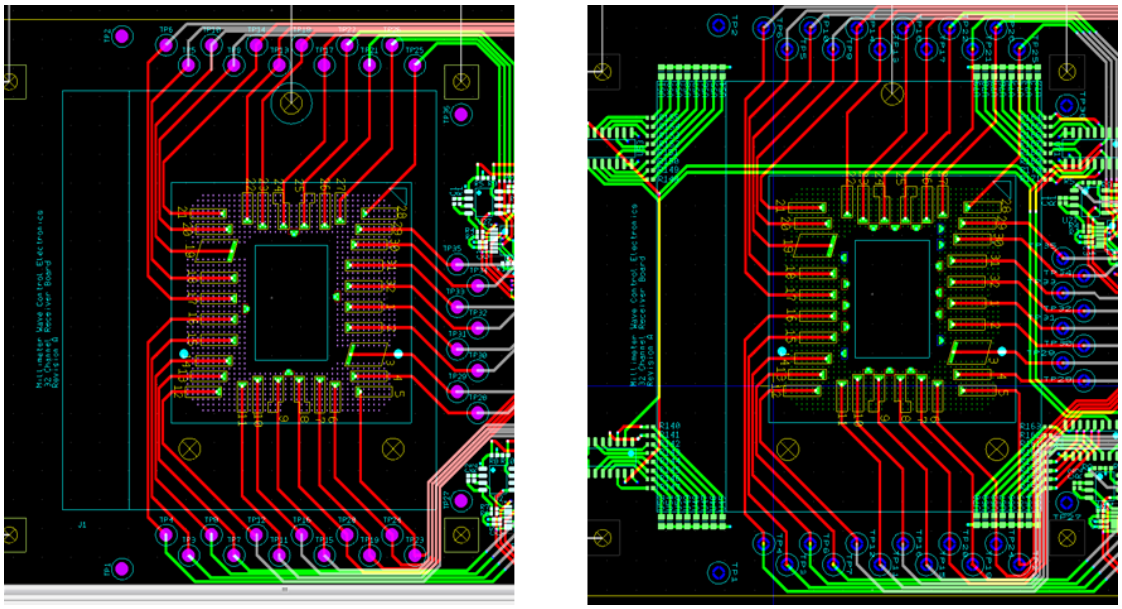

Figure 4.8 A comparison of the layouts of the Receiver30_RevA (left) and Receiver30_RevB. The addition of alignment LEDs and the necessary driving circuitry is showing surrounding the Detector Interposer socket. Additional ADC channels are also present.

This iteration of the Receiver board also saw some changes to the input circuit topology with the addition of a series filtering capacitor placed at each signal input to set the channel bandwidth. Together with the pre-existing parallel resistor to ground, this circuit now comprises a classic low-pass filter with a cutoff determined by the following relation:

$$Fc = \frac{1}{2\pi RC}$$

This is necessary as the gain-bandwidth-product of the operational amplifier may allow for some high-frequency noise to be amplified along with the signal. This value was set such that there would be no distortion imposed on the input signals, but that the high-frequency noise would be attenuated.

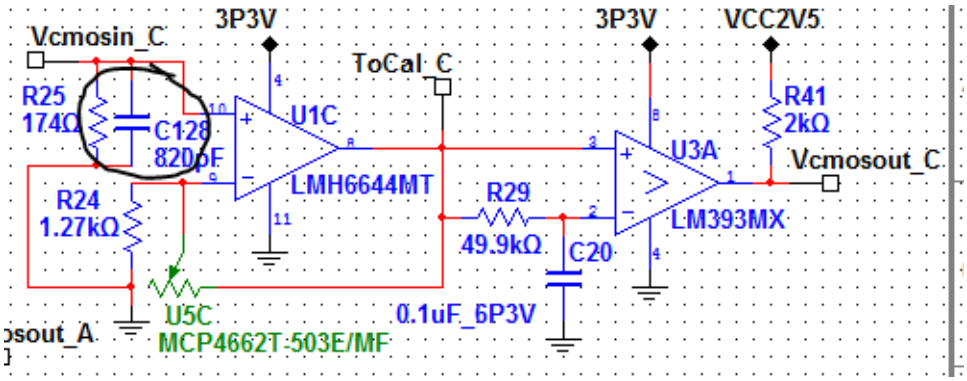

Figure 4.9 Modifications to the input circuitry of the Receiver board include the addition of a low-pass filter

The final major modification to the design was the inclusion of 16 more ADC sampling channels. This was necessary, in part, to support the addition of the alignment channels, but was also beneficial for system operation in general. The RevA design had included amplitude sampling support for only half of the input channels, but we found this feature particularly useful in testing and opted to expand it to the remaining channels. As each ADC sampling IC can support 16 inputs, the RevB design can support 32 channels in this manner. The addition of the four alignment channels necessitated that two of the thirty operational channels would not be sampled, but this was deemed an appropriate risk.

Schematic design and circuit board layout were accomplished with software using National Instruments Circuit Design Suite.

4.2.5 Detector Interposer

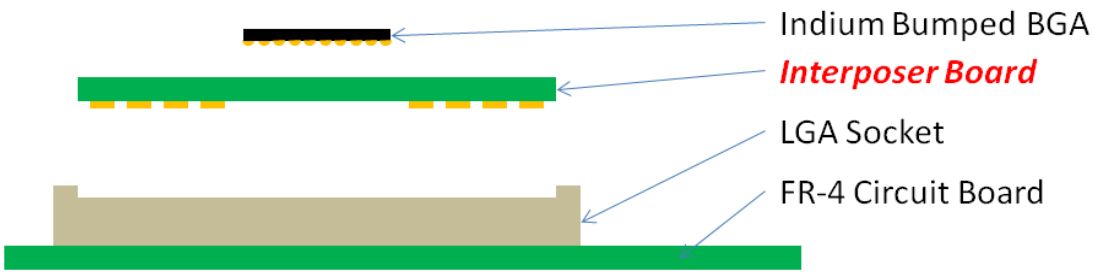

Figure 4.10 Conceptual rendering of the Detector Interposer PCB and its interfaces

The design of this PCB and subsequent FPA integration was a considerable challenge in the development of this control system. The PCB layout involved much smaller design features and tighter tolerances than any of the other PCBs we had yet designed. While the general design of this PCB is very simple in theory (simply routing one interface to another), the execution of the design was a fairly meticulous and time-intensive process.

4.2.5.1 Overview

The purpose of the Detector Interposer is to further modularize the design of the Receiver in the phase control electronics system. The Detector Interposer consists of an Indium-Gallium-Arsenide (InGaAs) focal-plane array (FPA) mounted to a custom printed circuit board which routes the electrical signals from the FPA to a land pattern which mimics the layout of common microprocessor chips. This design allows the Detector Interposer, and consequently the delicate and expensive InGaAs array to be easily installed and removed using COTS hardware that is readily available without any disturbance to the rest of the electronics. Unique to the Interposer was the intent of static design; as the rest of the system was expected to evolve, the Detector Interposer

was expected to remain the same for all iterations of the control system, including the final 220-Channel EC Prototype. In this regard, the design of the Detector Interposer was more stringent in that errors would be more costly, both temporally and monetarily given that the necessary design tolerances required more expensive fabrication processes, making subsequent design re-spins very expensive. As such, this design in particular was meant to be done once, and to be done properly in the first iteration.

Though the modularization of the Detector Interposer adds some complexity to the design of the Receiver, it was deemed necessary for several reasons. Firstly, the replication of the spiral antenna layout at the appropriate scale called for very tight tolerances, on the order of three mils (or three thousands of an inch). As one might expect, the processes necessary to fabricate PCBs at these tolerances are very costly. Further, these tolerances were more than twice as strict as those required by the outstanding circuitry of the Receiver board so it made little sense to fabricate a large PCB at an unnecessary level of precision to the majority of the relevant layout, this also at an exorbitant cost. Secondly, as the spiral layout was to be used in all planned future iterations of the system, the separation of the detector layout into a discrete element meant that the design material and each manufactured element could be modularly reused in all future iterations of the system, reducing design time. Thirdly, the FPA itself is both very delicate and very costly. A dropped or mishandled FPA is easily cracked or shattered, and the process of bonding the array to the PCB (discussed later) is not trivial. As such, a broken FPA could necessitate a complete replacement of the Receiver board; conversely, a damaged Receiver board could render a working FPA useless if it could not be removed.

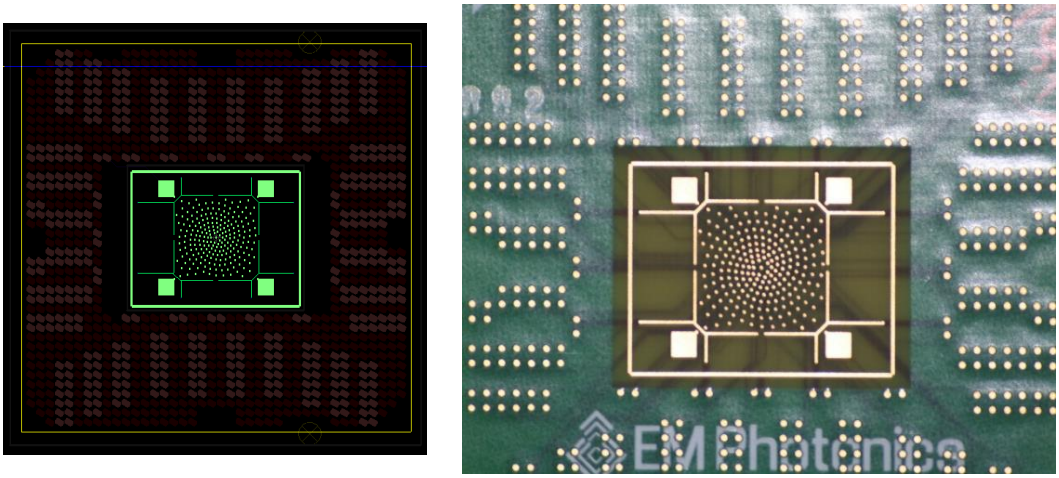

Figure 4.11 A design rendering of the Detector Interposer PCB (left) next to a photo of the physical PCB (right). Input channels are routed from the spiral at the center of the PCB to exposed pads on the PCB underside which are also exposed for probing on the top

As has been mentioned previously, we designed the interposer board to support all planned future iterations of the system, which meant it had to support 225 channels (220x signal channels and 5x auxiliary channels for optical alignment). These were separated into banks of channels, denoted by the letters A to G, and individual channels were denoted by alpha-numeric combination of its bank and channel number (E.G. B-5 or C27). Banks A-F were comprised of 32 channels each, with G bank containing the final 28. There were also five auxiliary channels added for alignment purposes, one at each extreme corner of the array and the fifth at the precise center of the array (boresight).

As an anecdote, the center alignment channel was incorporated into the design at the specific request of Dr. Thomas Dillon, one of the key PSI engineers and a fellow UD graduate, whose responsibility it was to define and fabricate the optical processor of each iteration of the PSI Imager. We were happy to oblige, of course, and I added

the requested channel to the Interposer design, but having no official designation to give the channel, we referred to it internally as ‘tom’. The nickname stuck, and the ‘tom’ channel can be seen in all subsequent design iterations, the only channel to eschew the typical alpha-numeric naming scheme.

4.2.5.2 FPA Integration

The contact side of the of the photodetector array consists of a ball-grid array (BGA) composed of Indium, and is designed to be attached to a carrier by means of Indium Bump Bonding, wherein the array is carefully positioned over the contacts of the carrier and compressed into place. The Indium bumps then deform around the contacts of whatever medium the array is being attached to and form an electrical bond. This interface can then be underfilled with epoxy to promote durability and maintain bond integrity under stress. The thickness of the Indium balls comprising this array is precisely 10µm, allowing for fairly little compression, and the balls are placed at a pitch of 25µm. By comparison, the contact area of the channel receptor pads on the top surface of the interposer are 6mil diameter circular copper pads, translating to approximately 150µm, so each channel is expected to make contact with several pixels of the FPA to form a so-called 'superpixel'. One unanticipated issue we encountered in the development of the bonding process was the issue of planarity. The photodetector array is manufactured to very stringent standards and thus requires very little preparation, but we realized after some research that the manufacturing standard of PCBs composed of FR4, a composite material comprised of a sheet of woven fiberglass suspended in epoxy resin, did not guarantee planarity to within our requirements to be able to bond it to the photodetector array without preparation. This created a serious issue for two reasons. As mentioned before, the optical alignment of

the FPA surface within the system is very delicate, and thus any angular displacement of the FPA face with respect to the optics and the Receiver PCB would distort the alignment of the control channels, and thus would undoubtedly have deleterious effects on the performance of the system. Secondly, as mentioned above the thickness of the Indium bumps is a mere 10μm, meaning that all contacts of the miniaturized array layout on the Interposer PCB face must be within a maximum deviation of 10μm in height to make contact with the Indium bumps. Ideally, this figure would be closer to 5μm to ensure reliable contact and bonding of all of the channels so we set this as the maximum height deviation requirement in efforts to maximize the yield of our bonding process. As such, it was necessary to ascertain the planar quality of the PCB samples before bonding. This was accomplished using a Dektak Profilometer, a stylus-based profilometer with approximately 5Å of vertical resolution, owned by the University of Delaware. PCB samples were binned according to their planarity and their propinquity to the requirements. Samples that were deemed to be the closest to the planarity requirements were then polished using lens paper and a calibrated flat test plate to ensure planarity, being checked periodically under a microscope to track progress as one could recognize a change in the finish of the copper pads to determine when a few microns of material had been successfully removed, and retested with the Dektak. Sample traces from this process are shown in Figure 4.12 below.

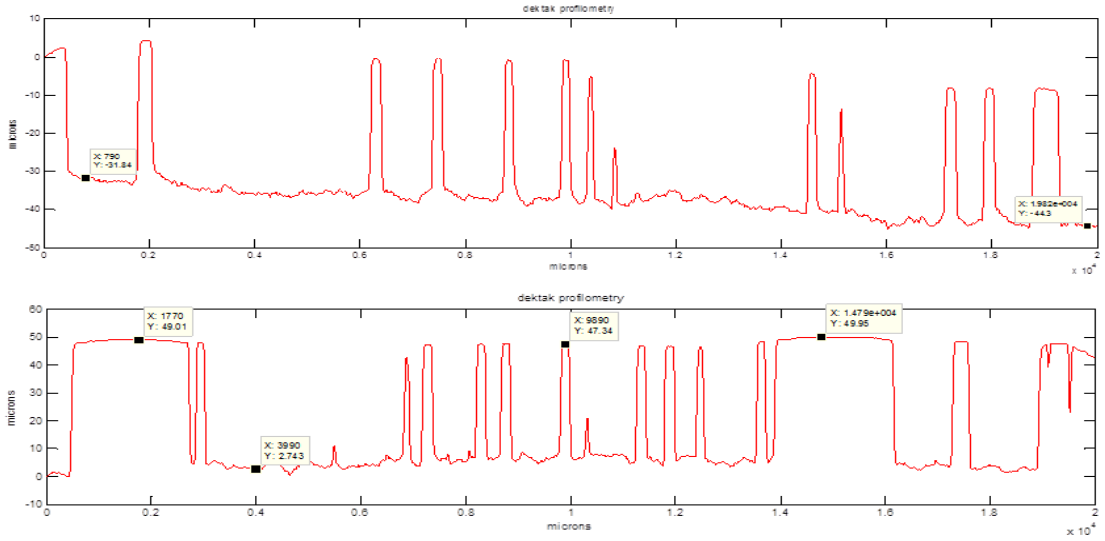

Figure 4.12 Sample planarity traces taken of Detector Interposer PCBs. The peaks indicate areas of copper with valleys of FR4 in between. While the top graph shows considerable variance in pad height, the bottom figure is much more planar after polishing.

Samples that required planarity correction to the extent that the removal of FR-4 material and solder mask was necessary were not used. With the planarizing of the PCB complete, it could then be cleaned, first with Acetone then with Isopropanol, and bonded to the Photodetector array.

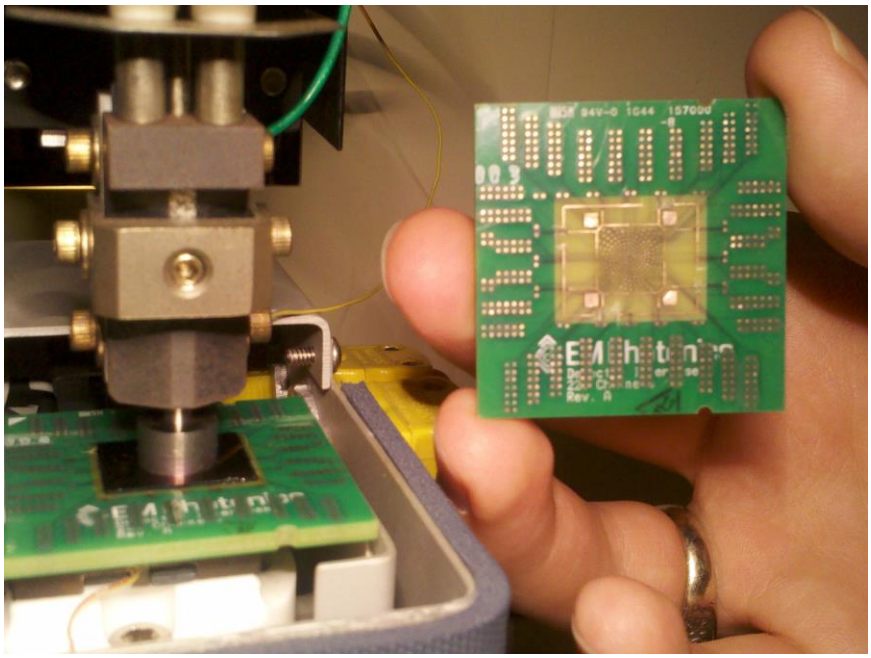

Figure 4.13 A bare Detector Interposer PCB is shown next another of the same, which is undergoing the process of having an FPA flip-chip bonded to its upper surface

In the end, we were able to create four working detector packages, one for each of the planned control systems (Pathfinder IIB, EC Prototype 30, and EC Prototype 220) and a spare. After testing we found that three of these had achieved 100% channel yield, with the fourth lacking only a single channel (this not among the inner 32, making it appropriate for use in either of the two thirty-channel systems).

4.2.6 Distribution30

The distribution30 PCB was designed by Petersen Curt, the lead engineer in this effort. His design incorporated the two connections necessary to interface with the RF phase modulators: an SMA phase-feedback interface and an RJ45 LNA power interface.

4.2.7 NI LabVIEW Interface

National Instruments LabVIEW (Laboratory Virtual Instrument Engineering Workbench) is a visual programming development environment[6]. Rather than textual code, programs are formed by graphically forming connections between VI's (virtual instruments), which can perform functions ranging from standard math operations to customized user routines. This method easily integrates control functions and can generate user-interactive switches and indicators to determine and illustrate system status. We utilized this program to create a GUI (Graphical User Interface) to the system. PSI had likewise done so with their previous system iterations, allowing us to incorporate some of this existing functionality, and allowing them to easily incorporate our phase control tools. This interface communicates with the system processor on the FPGA by means of TCP/IP and Ethernet. This was accomplished primarily through the use of LabVIEW's 'library call' functions, which allowed us to create our own VI's which called specific functions that were included in a DLL file. Within the DLL, a packet would be constructed and sent to indicate the desired action of the FPGA, which would receive the packet and send an 'acknowledge' signal after performing the requested action.

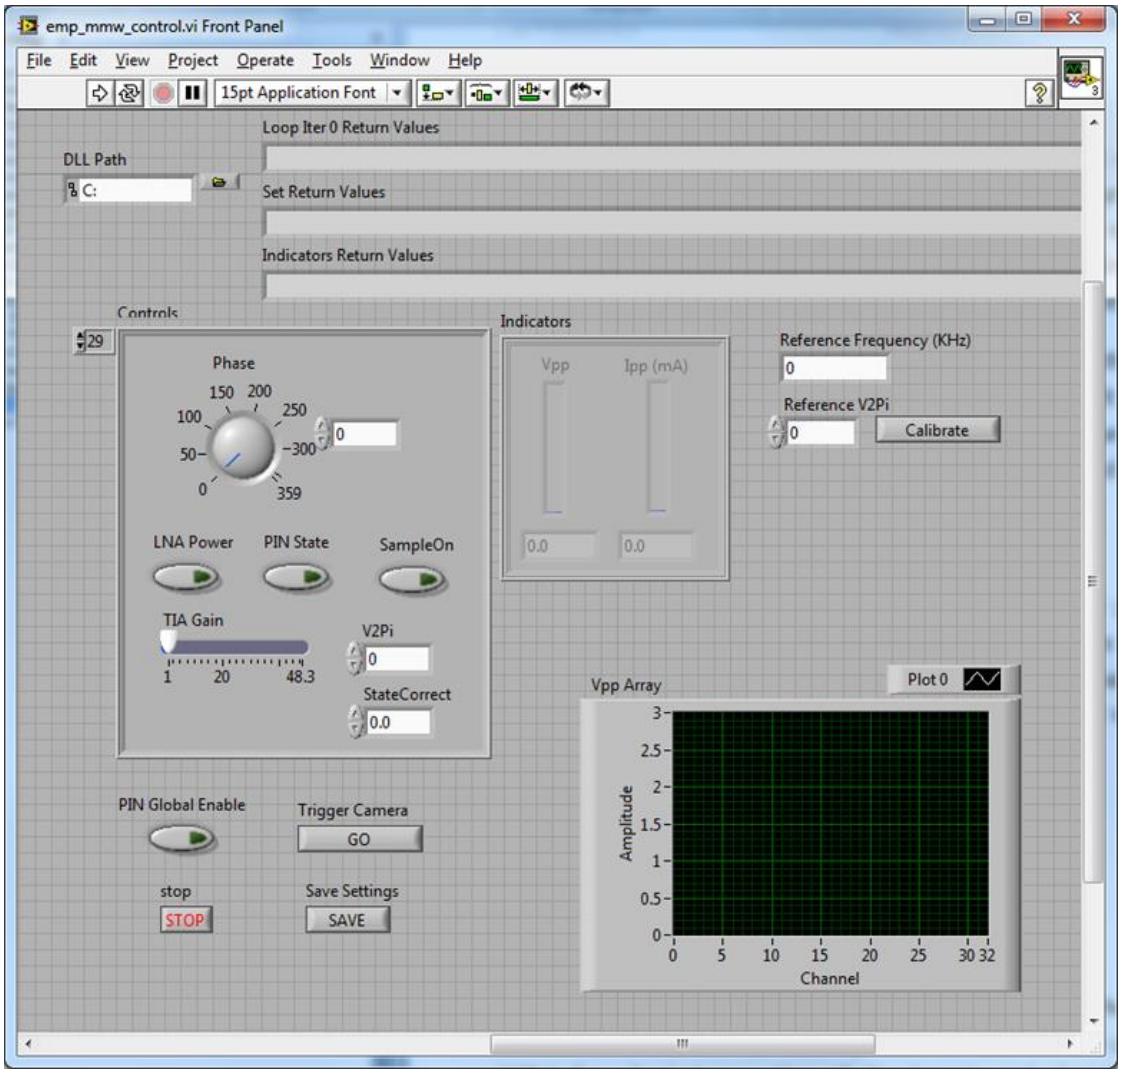

Figure 4.14 Front panel of the LabVIEW VI we developed for interfacing with the Pathfinder control system

4.2.8 Image Acquisition System

Generally referred to as the 'IAS', the Image Acquisition System is a software application that we developed to handle and process frame data, and to interface with the camera. The camera in question is a Flir SC2500 SWIR camera, and as such much of the interfacing is accomplished with tools provided by Flir. Camera capture is

accomplished using Flir's VirtualCam server, which scans for available cameras on the network and creates a 'virtual camera' object which passes frames by means of shared memory. In addition to interfacing with the camera, the IAS would perform rudimentary image processing routines using the OpenCV image processing library[7], including operations such as windowing, averaging, and background subtraction (for use with PIN Switching or Coherence Switching later on). This application is able to execute independently on any locally-connected host machine, as all communication is accomplished via the LAN interface. The final IAS application is purposed to run on a dedicated embedded system running Windows 7 Embedded. Development and Pathfinder operation, however, were accomplished using a variety of laboratory workstations, requiring that the application be easily portable and compatible across common versions of Windows, including XP, Vista, and 7.

4.2.9 Pathfinder Integration

Having completed our designs and preliminary testing of our PCB's, we integrated our control system with the Pathfinder Imaging System.

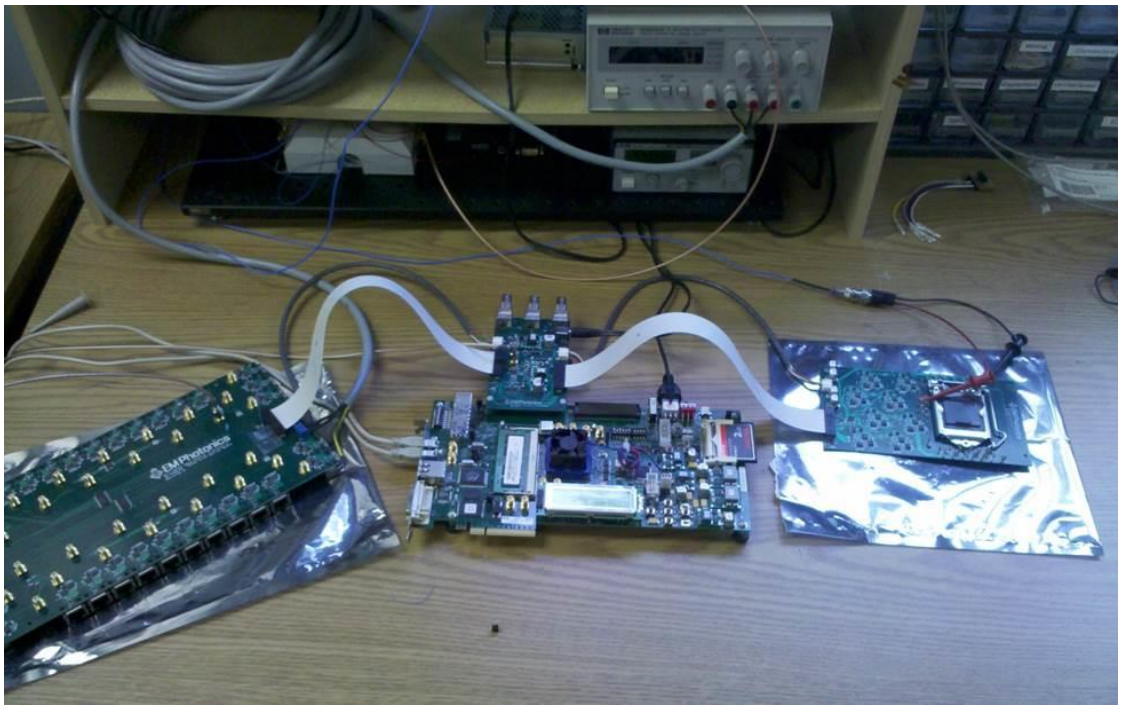

Figure 4.15 Pathfinder phase control electronics are shown under test prior to integration. Left to right: Distribution30, Xilinx ML-605 with custom FMC30 PCB, and Receiver30_RevA (sans Detector Interposer - empty socket). Input is provided by the test clips shown

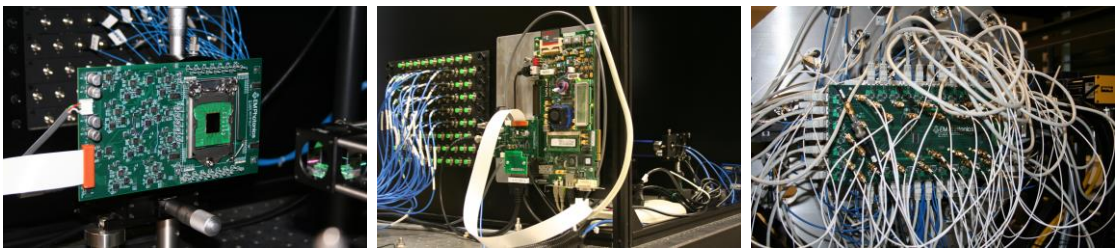

Figure 4.16 Phase control electronics are shown as integrated with the Pathfinder Imaging system. PCBs include the Receiver (left), the Xilinx ML-605 with custom FMC-30 PCB (center), and the Distribution30 (right)

As can be expected, we encountered several challenges during the integration of the Phase Control System with the Pathfinder imaging system, the most notable of which, along with their respective solutions, will be discussed in more detail now.

4.2.9.1 Phase-Feedback Dampening

During integration with the Pathfinder system, we found that the distribution circuitry was overdriving the LNA circuitry, resulting in signal overshoot and ringing, as can be seen below in Figure 4.17. While this did not drastically impact the contrast of resultant images, we determined this may have a detrimental effect on the system. Given the waveform shown, we determined that the power-distribution circuitry was underdamped, an unforeseen consequence of driving such a high-capacitance load. This issue effectively illustrates the sensitivity of a PCB design, as even a seemingly innocuous problem such as this is difficult to solve without a redesign of the PCB in question. In this instance, however, we determined that the problem may be solvable by imposing a larger series resistance within the power-delivery circuit. As such, I fabricated a custom connector which consisted of two female SMA connectors joined by a surface-mount resistor package. After a few iterations of varying resistor values, I demonstrated that this fix could render each distribution output critically-damped as per our desired design. This fix was noted for future design iterations, and I fabricated an identical series-interconnection for each of the 30 output channels on the board.

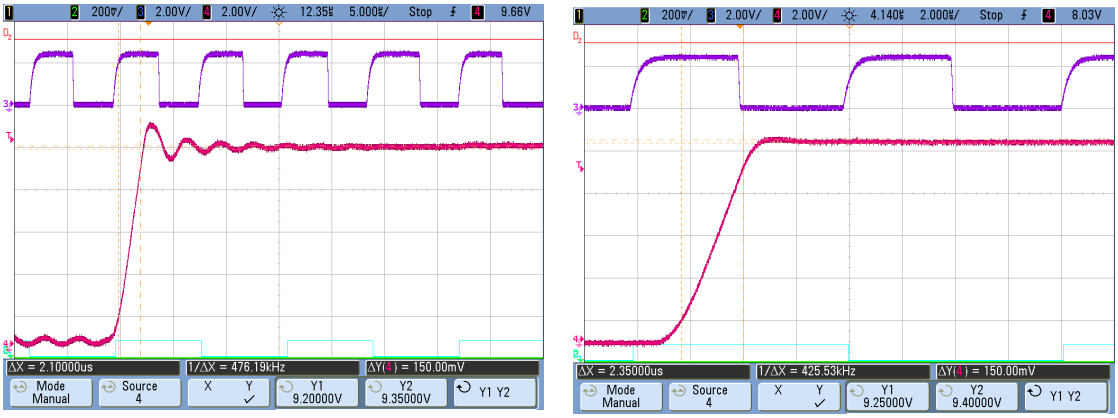

Figure 4.17 Underdamped LNA control output is seen to the left. Critically-damped LNA control output is seen to the right.

4.2.9.2 Flat-Flex Crosstalk Mitigation

Another challenge arose during Pathfinder integration in the form of high-frequency signal crosstalk. While the input signal frequency is relatively low (~200kHz), these are pulled high by weak pull-up resistors, resulting in a rounded rising-edge, and driven low, resulting in an abrupt falling edge. An unforeseen consequence of this was that these falling edges dissipated power very quickly, on the order of 25V/µS, which likewise dissipated current very quickly and generated a corresponding magnetic field in accordance with Ampere's law, which was occasionally strong enough to affect the weakly-pulled-up rising edge of neighboring signals within flat-flex connectors. While this phenomenon was not common or strong enough to reliably interrupt phase-locking specifically, it did have a very detrimental effect on the system's IIC control busses, which likewise rely on weak pull-up resistors for operation. As such, we quickly found the operation of the control bus to be unreliable. The anomalies in question are shown below in Figure 4.18.

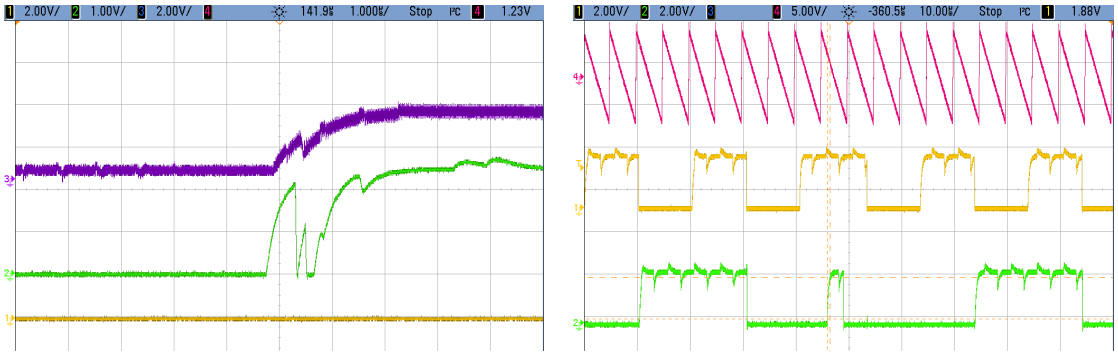

Figure 4.18 Signal errors caused by crosstalk of high-frequency signals are shown, measured from Pathfinder control electronics

The left image shows the rising edge of a control signal, which is reset after reaching hysteresis due to induced EMF from a neighboring comparator signal in the flat-flex cable. The result of signal interactions such as this is shown in the right image, in which the clock (SCL) and data (SDA) signals of an IIC bus, indicated by the yellow and green traces, respectively, display a momentary loss of synchronization which corrupts the entire communication. Incidentally, the periodicity of neighboring comparator edges can likewise be seen in the right image as noise in the variations of the signals when pulled high by the weak pull-up resistors. It is likewise notable that this noise is not seen when the signal is driven low by the comparator circuitry. These signal anomalies were rare and difficult to isolate, likely due to the statistical improbability that an instance of adequately strong noise would occur precisely in concert with the rising edge of a control signal – as the two were not synchronized in any fashion it was inevitable that this would occur at some point, but these instances were impossible to predict or characterize. The fix for this issue was considerably simpler than finding the issue which caused it; our temporary solution was to implement an edge filter in firmware which effectively eliminated adverse effects to

the extent that system operation was acceptable. In order to prevent similar issues in the future, greater care was given to flat-flex cable pinout, with more sensitive control signals isolated from comparator signals by common grounds. Though we did not perceive a detrimental effect to phase registration and locking, we determined that this effect was likely to have some effect on comparator signals as well. As such, we began to consider altering the comparator circuit on the Receiver board such that outputs would be driven high as well as low. This design was later implemented on the Receiver30_revC PCB.

4.2.9.3 SPI Update Level Transition

The physical interfacing of the Distribution board to the RF modules required that the board be placed very close to them, and so it was mounted to the back of the RF array, which can be seen in the integration pictures shown above in Figure 4.16**Error! Reference source not found.**. An unforeseen consequence of this was hat the Distribution board was therefore much farther from the Control board as had been anticipated, too far to use our intended flat-flex cables. As a workaround we quickly created some small adapter boards to convert the flat-flex signal interface to a twisted-pair VHDCI cable to help protect signal integrity. Unfortunately, the parasitic inductance of the long cable run, coupled with the capacitance of the twisted pair, had a deleterious effect on our edge rates which prevented the system from operating at its full 200kHz capability.

We were able to mitigate this problem adequately by compensating with changes to the firmware, but we noted that future designs would need to incorporate buffers to make sure these signals are driven adequately.

4.2.10 System Performance

The addition of our control system with the Pathfinder, now designated the Pathfinder IIB, resulted in an improvement in overall system performance. In particular, the increased update rate of our electronics resulted in more stable phase fronts and a flatter background, which likewise improved the contrast of the final imagery. Our electronics also provided new functionality that was not available before, including the ability to selectively power only channels indicated by the user so as to test the performance of smaller array configurations. The point-spread function of the Pathfinder IIB, recorded by imaging a single active point-source, is shown in the following figure.

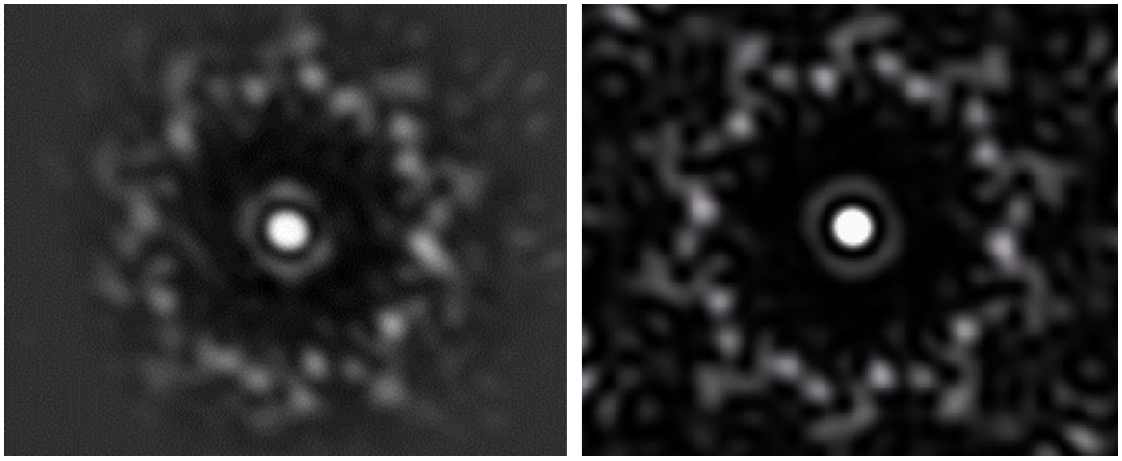

Figure 4.19 Pathfinder IIB PSF comparison of measured results (left) to simulated (right)

Passive imaging performance was found to be similar to that of the Pathfinder IIA, which was expected. A sample of passive images taken at a distance of 3m is shown in Figure 4.20.

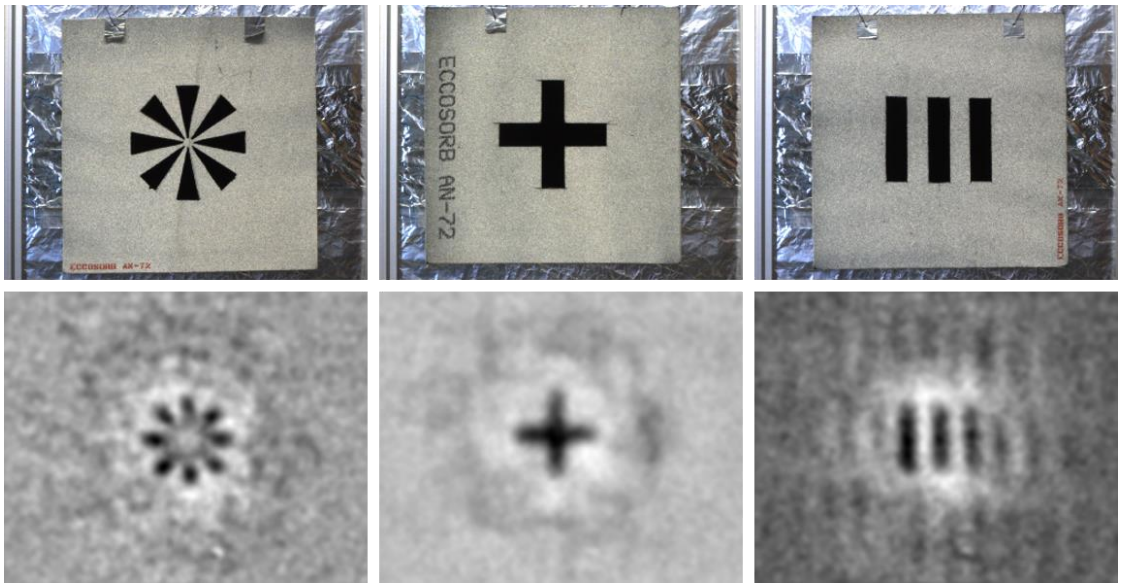

Figure 4.20 A sample of passive images taken with the Pathfinder IIB

4.3 30 Channel EC Prototype

4.3.1 Overview

Following the success of the Pathfinder system, we refocused our efforts on the next phase of the project, which was the 30-channel manifestation of the EC Prototype. This system was to make use of PSI's in-house RF modules, which would likewise be used in the final system, but would consist of only 30 channels rather than the 220 channels which were planned for the final system. This would allow more of a direct development path to move from the 35GHz system to this new 77GHz system, as well as an element of risk reduction in that integration issues could be identified and alleviated on this system rather than the larger, more complicated system. Later in development, this system likewise became a testbed for the full system, with modules being vetted by use in this system before integration into the final system. This shift in frequency from 35GHz RF elements to 77GHz RF elements brought with it several

benefits, including increased resolution (due to a lower diffraction limit) and smaller RF components, which reduced overall system size and weight. The size of the 35GHz modules would have made it impossible to construct an array of so many channels in the form factor required, but the drop in wavelength meant that the antennae and waveguides could be made to be less than half the size that they were before.

From the perspective of our control system, this system was largely the same as the Pathfinder as far as topology, and so our control system remained largely the same. The substitution of PSI's new modules meant that the parameters of the Distribution board would necessarily change, but the array layout and channel count would remain the same. Our control system therefore followed the same methodology as before, utilizing a Xilinx ML605 development board with an FMC card, as well as a single Receiver and Distribution board. Each of these design pieces was either upgraded or redesigned, as will be discussed in the following sections, but each also follows essentially the same design methodology as was employed in the development of the Pathfinder IIB system, and as a result will not need to be discussed in as much detail.

4.3.2 Receiver30RevC

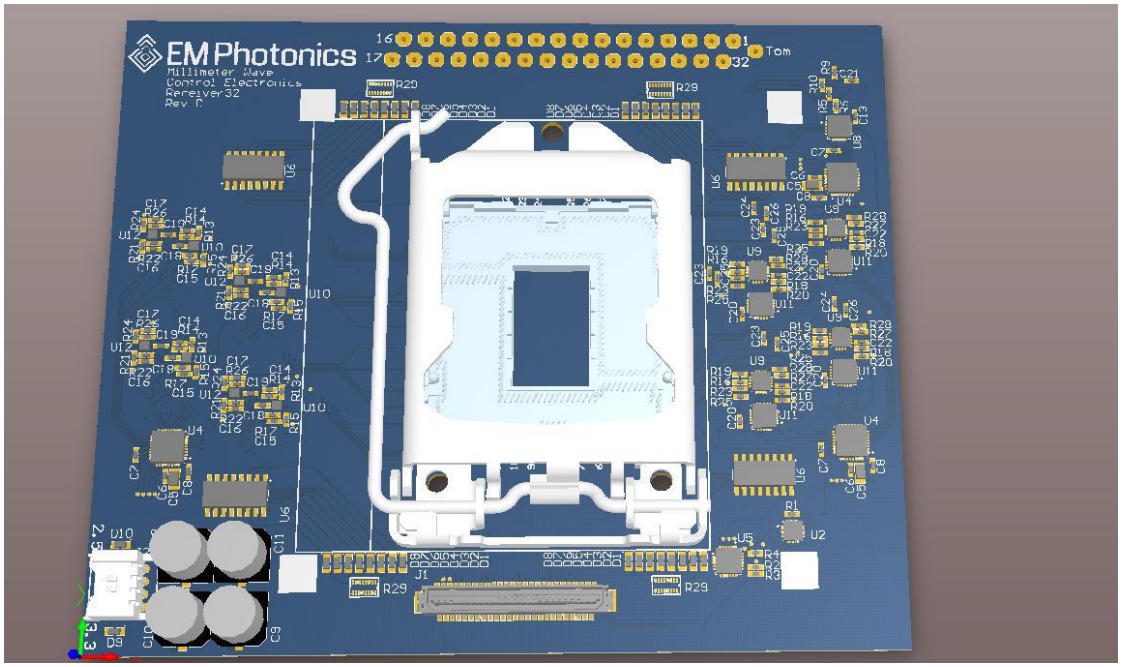

Figure 4.21 A computer-generated rendering of the completed Receiver30_RevC PCB

As discussed previously, several issues were encountered during integration with the Pathfinder imaging system. These were generally patchable in firmware, but also informed our design process going forward. Though not critical so as to necessitate an immediate respin of the hardware, we began working on new designs immediately. We determined that integration with the next iteration of the system, the 30-Channel EC Prototype, presented a fitting opportunity for another revision in design. To further facilitate risk-reduction, we also incorporated more compact design methods, including the substitution of various parts with smaller alternate packages. While care was taken in earlier revisions to limit the size of the design, these efforts

were amplified for this revision. The challenge was to reduce the area requirement of the electronics without losing any functionality (and in fact gaining new functionality).

4.3.2.1 Connector Pinout Modifications

Of particular concern on the earlier Receiver board design iterations was the issue of signal integrity on the I2C control busses, particularly when longer cables were used. We determined that the weakly-pulled-up signals were susceptible to spurious edges caused by crosstalk, this presumably the result of the electromagnetic fields generated by the quick dissipation of current of the falling edge of low-driven comparator signals. We had expected this bus to be resilient to crosstalk effects, owing to its relatively low operational frequency, but closer inspection revealed that signal glitches were occasionally potent enough to cause the control signals to drop below hysteresis levels, effectively sending a false bit and thereby desynchronizing any concurrent communication processes and occasionally causing the devices to become confused and lock up. This issue was patched by reconfiguring the FPGA to drive outgoing communications both high and low, rather than depending on the weak pull-up resistors. While this alleviated the communication issue, the signal crosstalk issue could not be effectively addressed without a hardware redesign. In particular, we determined that the signals causing the crosstalk would benefit from a more tightly coupled signal return path, and so we redefined the flat-flex connector pinout such that grounds were more evenly distributed among comparator signals. Another signal integrity issue we experienced previously came as a result of the 50MHz ADC_SCLK signal being placed directly next to one of the comparator outputs (Vcmosout_12), resulting in many spurious edges being imposed on the signal, which were frequently misinterpreted by the FPGA as being zero-crossings, and thus invalidating phase

measurements and causing that particular channel to lose phase-lock. This issue was alleviated in practice by simply disabling the ADC clock, which restored normal operation, but at the cost of a slight loss of functionality. In the redesign, this signal was moved to the edge of the connector to prevent this issue from recurring. These pinout changes are reflected in the figure below. The difference in appearance between the two designs also illustrates our software transition at this point in the project from the NI Circuit Design Suite to Altium Designer for our PCB schematic and layout designs.

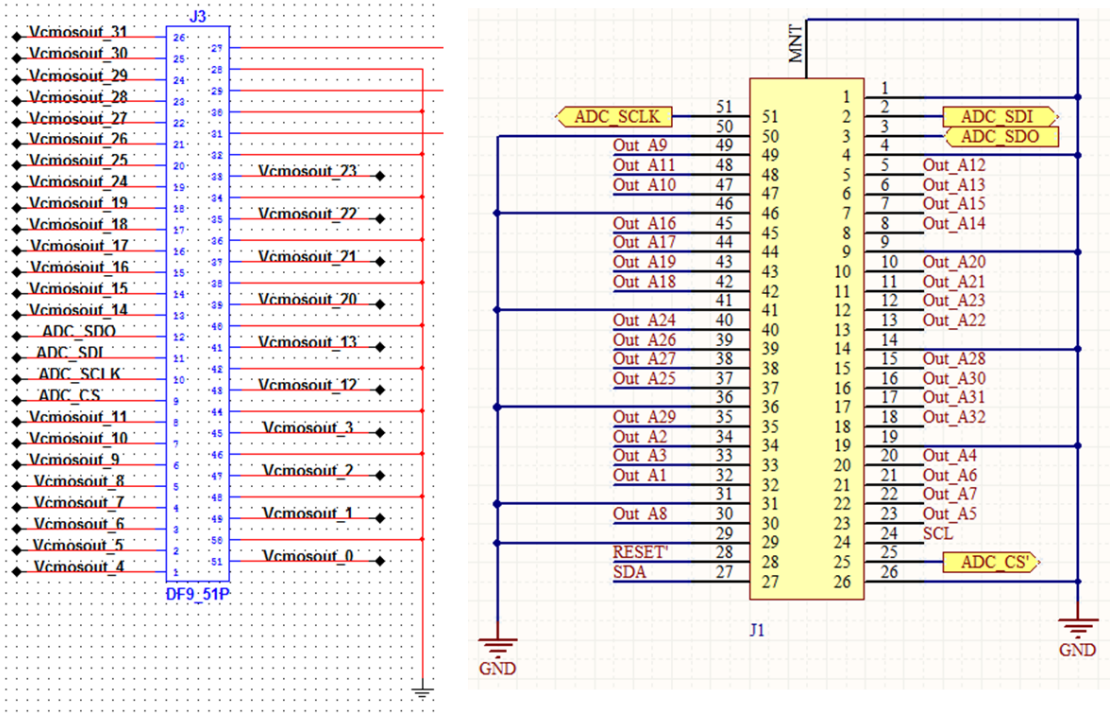

Figure 4.22 A comparison of connector pinouts from the Receiver30_RevB (left) and the Receiver30_RevC (right). The rearrangement of the signals appears subtle, but had a marked effect on our crosstalk issues.

4.3.2.2 Comparator Circuit Modifications

The comparator circuits utilized in earlier iterations of the Receiver board operated by means of selectively either driving a signal low or allowing it to float high by being pulled up by a resistor to VCC. As noted in the discussion above, however, such signals were found to be susceptible to crosstalk. While we did not witness any direct effect on the performance of the system due to any crosstalk affecting rising edges of comparator outputs, it was certainly feasible that occasional false edges would register, but that the high update rate of the system would correct the error before it was noticeable. As such, we modified the comparator circuit such that the output would now be driven both high and low, eliminating any opportunity for signal degradation due to crosstalk.

The design of this board resulted in yet another "lesson learned" in the field of PCB design: the layout schematics for the new comparator part were sourced from layout software provider, and did not follow the same design style as the layout schematics used in previous design iterations. In particular, the comparator symbol is essentially identical to that of a typical op-amp, with a connection directly above and below the part for VCC and GND connections, respectively, which were unlabeled. While the original schematic symbols had placed the positive input above the negative input, these new symbols were reversed and so the new symbol was inverted to correspond directly to the existing layout. As you may already have guessed, this resulted in the chip being reverse-biased. While this did require a respin of the board, it was quite fortunate that we made this board as a risk reduction for the 220-channel board due to the design changes; had we proceeded directly to the 220-channel board this would have constituted a much more costly refabrication. Still, a valuable lesson

was learned not to rely on pin placement or schematic methodology to be consistent between manufacturers or designs.

4.3.2.3 Space-Reduction Measures

As these design modifications posed an opportune occasion for testing new designs, we deemed it prudent to test various size-reduction measures that were likely to prove necessary in the next iteration of the system, the 220-Channel EC Prototype. New chip-scale parts were selected to reduce component space, replacing the DIP and SOIC footprint parts that had been used before with much smaller QFN and BGA footprint parts. Parts were also arranged on both surfaces of the PCB, allowing for more efficient use of space. These parts required more stringent cooling measures, in that conventional parts are effectively cooled by conduction of heat through their contact leads. As these packages incorporate no such leads, large copper pads with vias providing direct connection to ground planes are required to maintain adequate cooling. Care must also be taken to use tighter drill tolerances with thermal vias, as larger holes are more apt to wick solder when heated, pulling it away from the pad and preventing a solid connection to the copper.

These space-reduction measures proved quite effective. Where previously a unit cell supporting four channels occupied approximately 630mm^2, the redesigned layout required less than 170mm^2, a nearly 75% decrease in 2D space.

4.3.3 FMC30_RevB

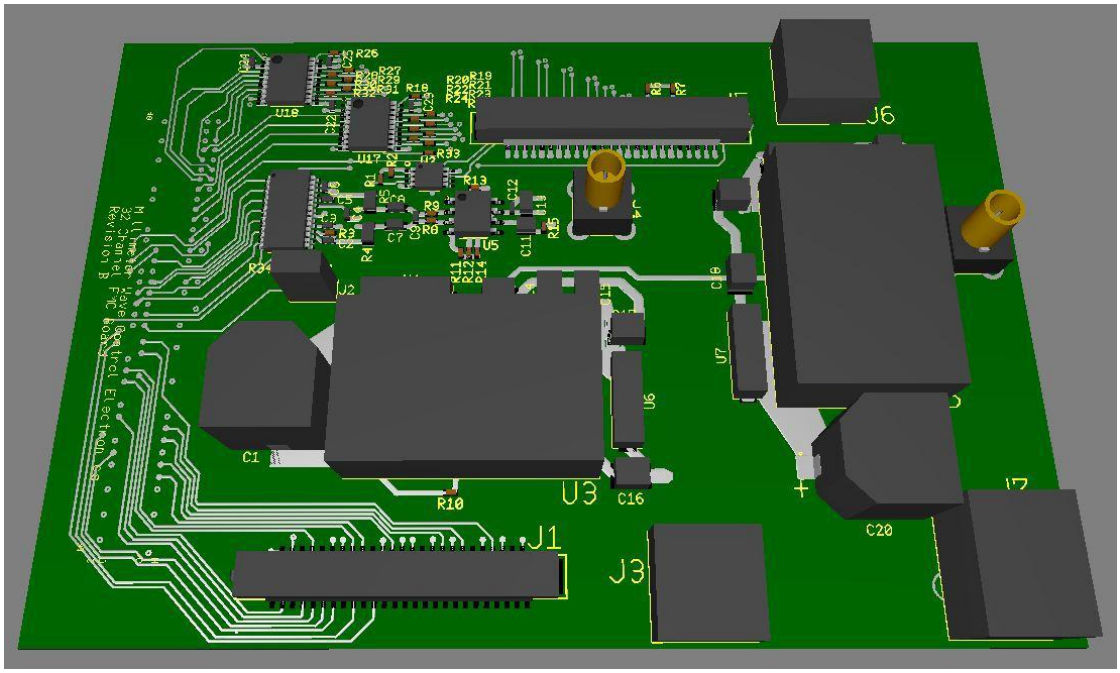

Figure 4.23 A computer-generated rendering of the FMC_30_RevB

As the redesign of the Receiver board included a redefinition of the I/O connector pinout, these changes were necessarily mirrored on the FMC board. As a new hardware design was necessitated, we likewise took this opportunity to make improvements to the PCB design.

After Pathfinder integration necessitated the use of unexpectedly long cable runs to connect the FMC board to the Distribution board, we found that a lack of buffering and/or impedance matching had a deleterious effect on our control signals. As such, we took action to ensure proper level-conversion was in place for these signals and added series termination resistors for impedance matching

We also simplified the board design a bit by removing unused circuits. As the new modules no longer employed PIN switches, the PIN status input was no longer required. We also found the selective trigger input unnecessary, opting to remove it as well. Finally, we replaced the last remaining BNC connector for the camera trigger output with a much smaller SMA connector to conserve space.

4.3.4 Distribution55_RevA and RevB

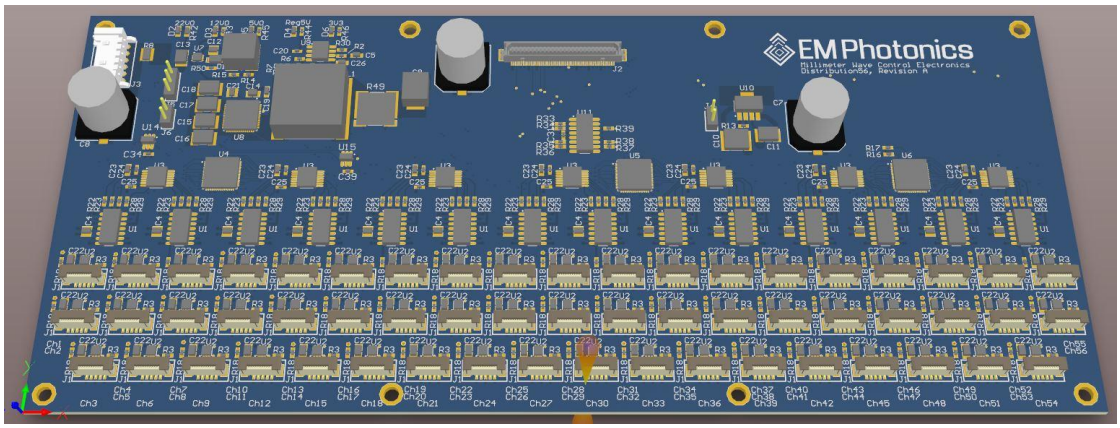

Figure 4.24 A computer-generated rendering of the Distribution55_RevA PCB

In particular, the distribution board had to be completely redesigned. Several improvements could now be made, not least of all that the removal of the bulky connectors that had been in place before meant that a more efficient design could be achieved. We also determined it would be more prudent to design this board to the requirements of the 220-channel system as well as those of the 30-channel system currently underway, thus the specification of 55 channels rather than 32 as before to allow support of the next system with four distribution boards; as with the naming conventions of the previous boards, this PCB actually supports 56 channels, but one is

unused (in that four of these boards support 220 channels, the channel count of the final EC Prototype system). As such, only a single Distribution board will be required for this system. As with previous iterations of the control system, the Distribution board design was completed by Petersen Curt.

The use of new RF modules necessitated the implementation of new interfaces, requiring a redesign of all the relevant interfaces. While size constraints were not paramount at this stage of development, they soon would be as the program progressed to the development of the 200-channel EC Prototype, and since the purpose of this design was to test the methodologies intended for use in its construction we found it prudent to take steps toward scaling the design to the appropriate size. As such, more care was taken beyond merely the scale of the antennae to reduce the size of the electronics and streamline the design. Where before power and phase feedback had been carried to the modules separately via RJ45 and SMA cables, respectively, a new interface comprised of a single, small-footprint, flat-flex cable (FFC) was substituted. The new modules also required higher modulation voltages, originally expected to be about 8V, whereas the original Distribution30 was purposed to support the 35GHz RF components, which had a Vpi of 5.5V.

The FPGA processor is able to communicate with this board by means of the flat flex connector seen at the top of the design, with each module connecting to the board through the 56 discrete connectors that occupy the majority of the lower portion of the board's face. As these boards are intended to be mounted around the perimeter of the EC Prototype 220 enclosure, these connectors are populated only on one side. These flat-flex ZIF connectors accommodate 5mm width cables of lengths up to

500mm and 10 conductors, carrying signals to the modules such as 5VDC RF amplifier power, 0-22V analog phase-feedback, an input to indicate module presence.

The previous board utilized power conditioned by the FMC30 PCB as well as an external rack-mounted power supply to power the RF amplifiers. This new board incorporates the functionality of both in a much smaller form-factor, now accepting a single power input connection at 22VDC, supporting up to 20A of current. This supply is then filtered and regulated to provide 22VDC, 12VDC, 5VDC, and 3.3VDC supply rails.

4.3.4.1 Distribution55_RevB Updates

Initial testing and integration of the Distribution55 PCB uncovered a series of problems that we determined would be best addressed by a revision in hardware design. In particular, the 12VDC supply was not able to reliably maintain a sufficient voltage level, generally providing only about 9VDC. As the 5V supply was subsequently regulated from this supply, it was found to be inoperable as well. In order to alleviate this issue in testing we substituted an external 12VDC supply which allowed us to continue testing and substituted a new 12VDC regulator circuit in our subsequent designs.

We also found that the 5V and 22V power supplies exhibited some appreciable ripple noise when supporting larger loads. This had little effect on our testing procedures, as this was only observed when the 5V display was under stress attempting to power several modules simultaneously. We attempted to filter this ripple by adding a larger capacitor to the input of the switching supply, and by redefining the regulator circuit's layout so as to reduce the ESR of the capacitors and improve the mutual inductance of the return current paths. To further isolate the analog

components from regulator switching noise should this issue persist, we instantiated an LC decoupling network to isolate the analog supply from the main supply with a low-pass filter.

Another major difference in the RevB design was the addition of individual I2C busses for each module to support a new technique used by PSI to calibrate module gain. While I2C busses already existed in the RevA design, each module is not individually addressable and so each had to be given its own bus network.

4.3.4.2 Distribution Board Connector Testing

The ZIF connectors to be employed on the new Distribution55 PCB were largely unknown to us, and as such we took it upon ourselves to evaluate their performance under vibration conditions. After identifying three primary candidates, we designed a very simple PCB which included two of each. A logic ‘high’ would be transmitted from one to the other using the same flat-flex cabling intended for system operation. The signal from each pin of the cable was aggregated by means of a logical AND, which would exhibit a falling edge should any of the conductors of that cable disconnect. This layout was replicated for all three candidates, and the aggregated output of each was monitored by individual channels of our oscilloscope. We then developed a data-logging instrument in LabVIEW which would establish a connection to the oscilloscope and monitor its output, recording an event for every disconnection. This board was then mounted to our vibration-testing setup, allowing frequency-of-failure evaluation of our candidate connectors under harsh vibration conditions.

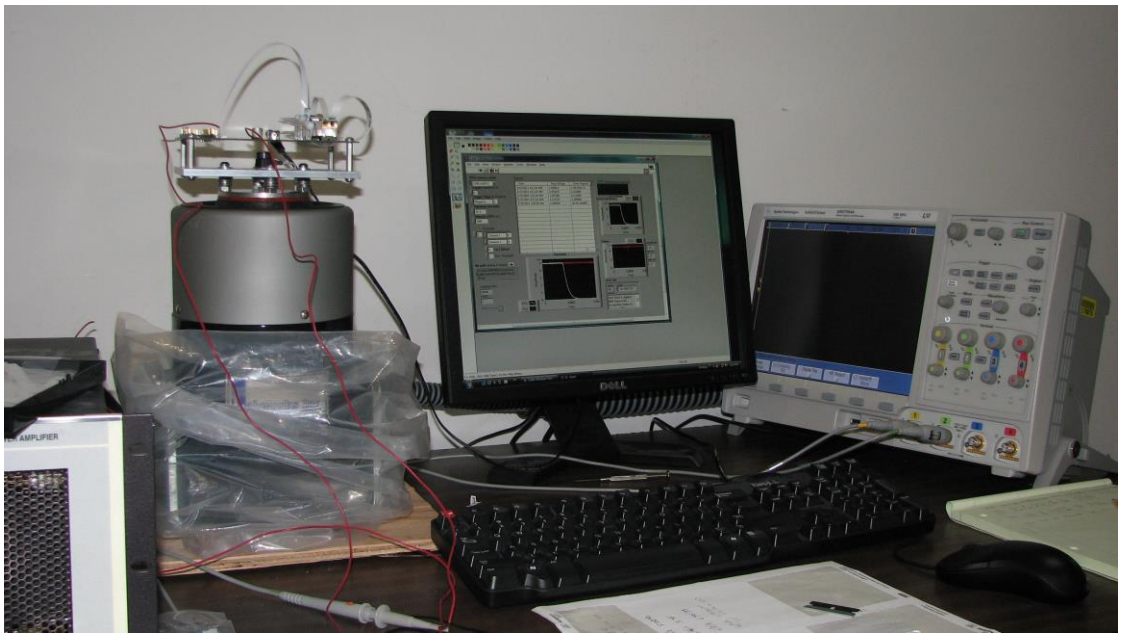

Figure 4.25 Our flat-flex ZIF connector evaluation setup, comprised by a custom test PCB affixed to a vibration piston. Connection integrity is monitored by the oscilloscope, with any events recorded by a LabVIEW VI running on the control PC.

4.3.5 Firmware and Software Improvements

Several modifications and upgrades were instrumented for the system firmware and software during this phase of development, as will be discussed in the following sections.

4.3.5.1 Operational Software Modularization

The various hardware changes in this next iteration of the control system likewise necessitated updates to the system firmware. Some of these changes were fairly simplistic, like updating addressing to accommodate new components, but still carried farther-reaching implications. Rather than simply branch the design, we thought it important to attempt to maintain backwards compatibility with the

Pathfinder system with our development software. In this way, future changes to the EC Prototype can be tested on the Pathfinder and vice-versa. We accomplished this by first streamlining the code base, refactoring the majority of our component-specific code to object-oriented classes. We then parameterized the operational code such that executable software for either system can be compiled from the same source code going forward.

We also added firmware and software support for new hardware features. The system now has the ability to sense when a channel has been connected or disconnected from the Distribution board, which could automatically trigger an alert or an auto-calibration routine. New health monitoring circuitry on the PCBs also supports measurements of power draw and protection from overcurrent events. This monitoring can initiate an emergency shutdown of a module's power source in the event that it starts drawing more power than expected.

4.3.5.2 Camera Capture Improvements

Prior to integration with the EC Prototype system we evaluated improvements to the system software and firmware using the Pathfinder IIB 35GHz system, as the two are operationally similar.

As the new 77GHz RF modules produce more sideband power than those used in the Pathfinder system, we found that we were able to decrease camera integration times and consequently increase framerate. We encountered an issue in that our triggering technique added considerable latency, effectively limiting our maximum framerate to around thirty frames-per-second. While this had little effect on the performance of the Pathfinder system, it would greatly hinder the performance of the EC Prototype, as the performance gains by utilizing shorter integration times would be

effectively moot. The problem manifested in that, as has been mentioned, our system captures frames in “scene” and “load” pairs; subtracting the load image from the scene image greatly improves the signal-to-noise ratio of the scene data by removing deterministic noise. Control of this process was implemented in LabVIEW, which while highly versatile is not optimized for performance. Essentially, the LabVIEW VI would signal the FPGA processor to toggle the switch state, waiting for a reply signaling that this operation was complete before initiating capture of the next image from the camera. As such, the overhead of these network communication routines caused a great deal of subsequent latency in camera trigger operations.

Our solution to this issue was to implement the foreground/background switching logic in firmware. As the FPGA was determining which image to take at any given point in time, it was no longer necessary to send time-intensive network packets to trigger a change. As many solutions often do, this uncovered yet another problem, namely that the software would occasionally drop a frame, causing a desynchronization of the background subtraction process and effectively inverting the contrast of the resultant image. We initially tried to solve this problem with image-processing methods by detecting a change in the contrast of the image, but these efforts proved fruitless. The final solution was found in the metadata that was bundled with every frame captured from the camera; this metadata included a data structure of frame information, including a global “frame_cnt” parameter. It was not necessary to track this number in any detail, but merely to ensure that successive frames were not both even-numbered or both odd-numbered. This is, perhaps ironically, much simpler than our original image-processing approach, and has proven very robust in preventing these errors.

With these changes in place in software and firmware, we were able to increase our capture framerate from 30Hz to over 250Hz – a considerable improvement in performance.

4.3.6 System Integration

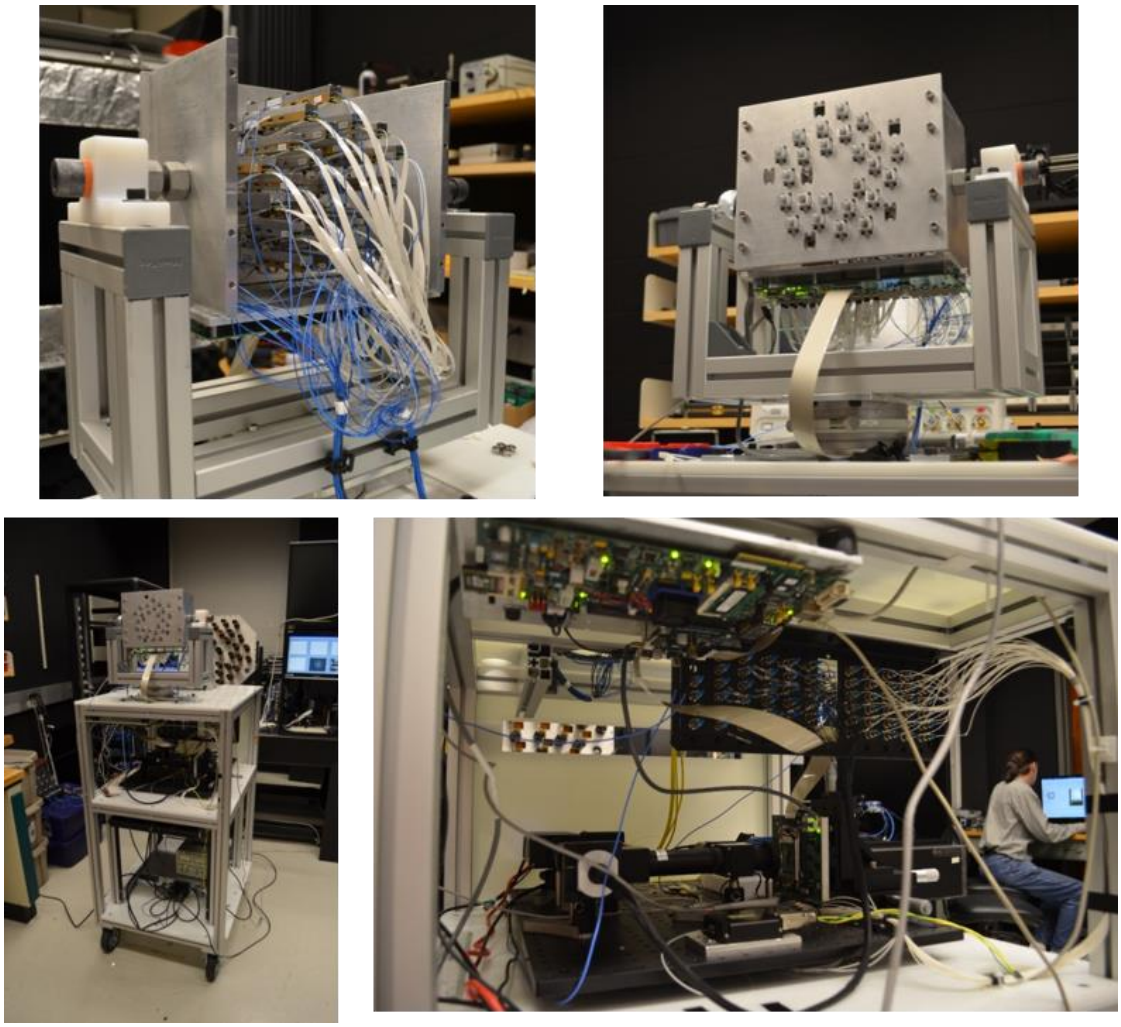

Figure 4.26 Pictures showing our phase control electronics integrated with the 30-Channel EC Prototype imaging system

Where the Pathfinder system was built on an optics table, the 30-Channel EC Prototype system was built into a cart for portability. This was facilitated by the smaller array size allowed by the use of smaller modules (due to the increase in wavelength). The RF modules were arranged in a faceplate mounted to a swiveling gimbal, allowing easy adjustment for imaging targets. The Distribution board was attached below this assembly so as to be as close as possible to the module components. The remainder of the electronics, as well as the optics, were mounted on a lower level of the cart.

As part of the integration process we also performed stress-testing on PSI's upconversion modules, including both vibration and temperature testing. In addition to proving PSI's module designs, this also presented the opportunity for us to test the operation of our control electronics under harsh conditions.

4.3.6.1 Vibration Testing

Using our vibration-testing setup, we repurposed our connector test PCB, which we discussed a moment ago, to serve as a mounting point connecting the module to the vibration solenoid. We first performed survivability test, subjecting a test module to continuous vibrations for increasing periods of time, testing operation intermittently to ascertain the continued functionality of the module. After no failures were detected during this testing, we performed system operation tests while subjecting the module under test to harsh vibration. This testing setup is shown below in Figure 4.27. An oscilloscope is used to visually monitor the interference pattern as sensed by the Receiver board and also to track statistics over time including signal phase average and standard deviation. Following 15 minutes of operation under vibration conditions, the oscilloscope reported a standard deviation of 4.87° in channel

phase. While this provided aggregate monitoring of phase consistency over time, we also configured the FPGA to output instantaneous samplings of measured phase error and calculated compensation to a control PC. This data was later analyzed with MATLAB to extract information regarding instantaneous phase errors.

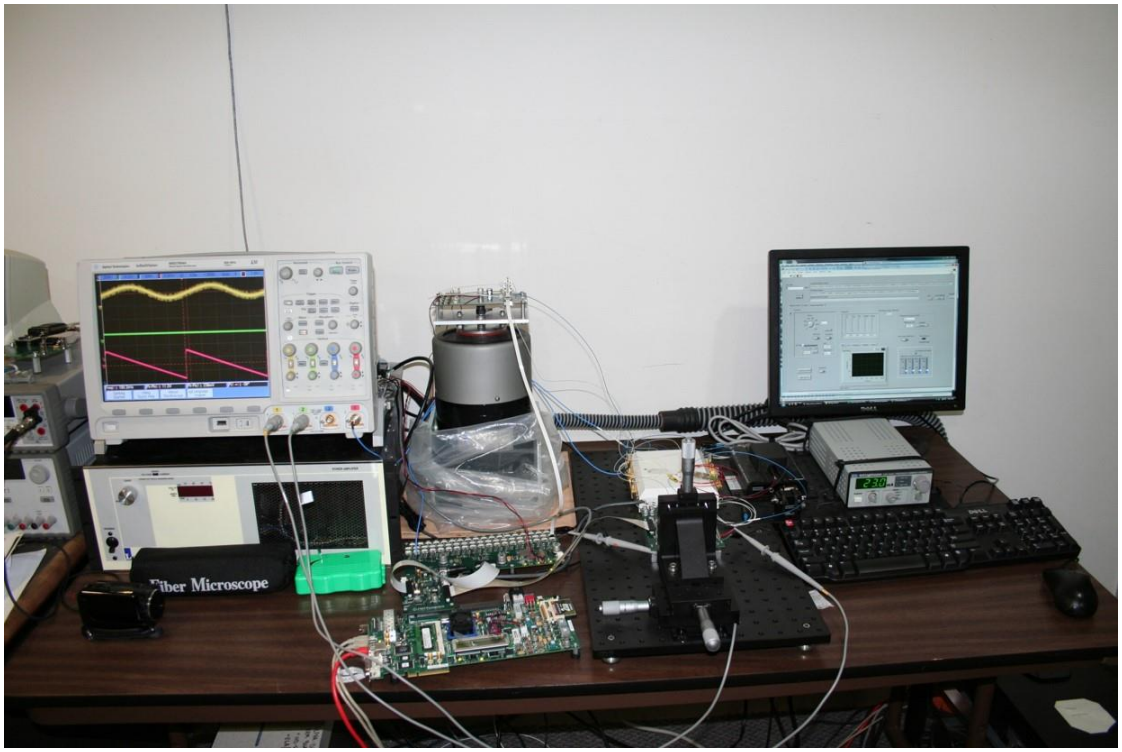

Figure 4.27 Vibration test setup. Instrumentation shown includes FPGA control electronics PCB, Distribution electronics PCB, phase control electronics power supply, vibration solenoid, vibration solenoid power amplifier, PC for vibration control, PC for phase control and data capture, and digital oscilloscope.

The last time we had performed system performance characterization under vibration conditions had been during the testing of the 4-Channel prototype hardware. As appreciable changes had been made to the phase control methodology, from the

software and firmware to the physical circuitry, and given that entirely new phase modulators were now in use, we compared the performance of these systems to quantify whether the performance of our system had improved. As we can see from the data presented in the following figure, performance was slightly worse in steady-state measurements. We determined that this was likely due primarily to noise from our custom Detector Interposer FPA. The performance under vibration, however, was seen to be considerably improved, owing most likely to the faster update rates and higher-quality comparator signals generated by the new hardware. In either case, these systems showed excellent performance under testing, far exceeding our operational requirement of 15° of phase error.

		Test Conditions	Measured Error (3 Std. Dev.)	
Date	**Hardware**	**DUT**	**No Vibe**	**Random**
5/14/2010	35 GHz 4 Channel	35GHz (Fibers Only)	2.2°	5.1°
8/29/2011	77GHz EC	77GHz Upconversion Module	2.7°	3.2°

Figure 4.28 Table showing comparison of the EC Prototype phase control electronics performance with and without vibration to that of the original 4-Channel prototype hardware

4.3.6.2 Temperature Testing

Our temperature testing was performed in three stages. First, a non-functional module (mechanical sample) was held at increasing temperatures to gauge how the temperatures would affect the module housing itself. Given the expense, both monetarily and temporal, involved in fabricating one of these modules, we were eager to identify any risk-reduction steps. After the housings were seen to survive extreme temperatures with no apparent visible defect, we repeated the same test with an operational module, testing performance between test iterations to ascertain continued

module functionality. As a final testing phase, we then subjected this module to high temperatures while under power and observed whether the operation of the module was appreciably affected by temperature over time.

The target of these tests was to confirm that the module could endure 50°C during operation and 70°C ambient temperatures in storage. These tests were accomplished with the use of a temperature control chamber with two temperature-monitoring probes: one to measure ambient temperature and another to measure the precise temperature of the device under test.

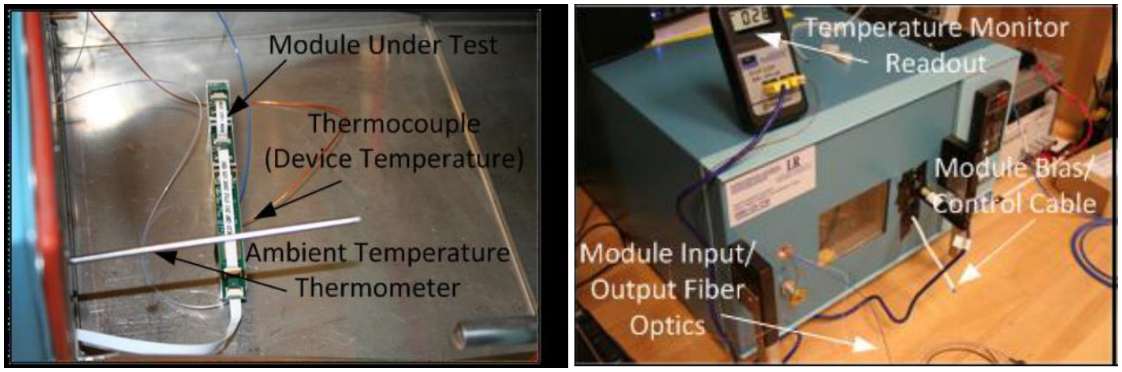

Figure 4.29 (left) PSI Upconversion Module undergoing temperature testing inside our temperature-control chamber. (right) The temperature control chamber is shown with some relevant test equipment.

Module H9	Baseline	After 2hrs @ 50°C & 3hrs @ 70°C
Sideband with noise diode	-33.57 dBm	-34.07 dBm
Sideband without noise diode	-38.93 dBm	-39.38 dBm

Figure 4.30 PSI Upconversion Module performance is shown to be reasonably steady under high-temperature conditions

4.3.6.3 System Performance

We were able to demonstrate full operation of the 30-Channel EC Prototype imaging system. Though utilizing the same number of optical channels as the Pathfinder system, the improvement in performance can be illustrated with a comparison of the point-spread functions of these two systems, shown in the following figure.

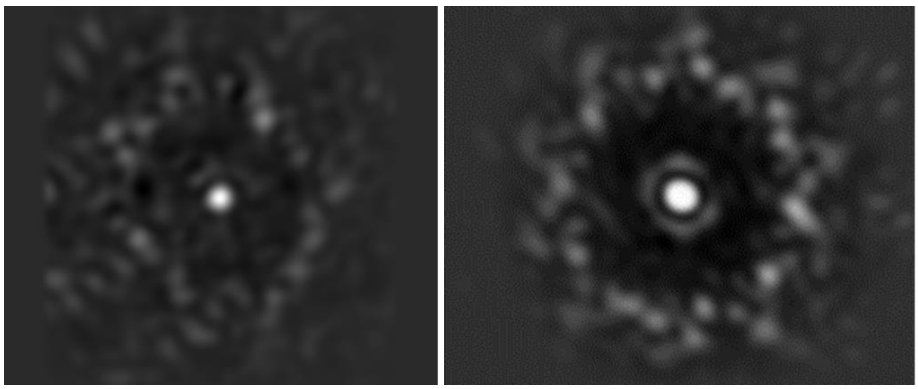

Figure 4.31 A comparison between the phase-locked point spread functions of the 77GHz 30-Channel EC Prototype (left) and the 35GHz Pathfinder IIB (right)

Owing largely to the difference in operational frequency, the new system is able to produce a better-defined point than was possible with the Pathfinder. Though side lobes are still readily apparent, they are considerably fainter and likewise less obtrusive to the image than before.

4.4 220-Channel EC Prototype

The final iteration of our phase control electronic system was designed to support PSI's 220-Channel EC Prototype imaging system. This system represents the culmination of all previous work, incorporating the technology developed in the

previous system iterations, but with improved performance and in a fieldable form-factor.

4.4.1 Overview

In theory, the 220-Channel EC Prototype is functionally equivalent to the 30-Channel EC Prototype system with the addition of many more channels. Upgrading the system to 220 channels meant quite a bit more than just adding an additional 188 channels, however, as this system had the most stringent design requirements. In effect, the challenge was to build a 'larger' and more complicated system that was, in actuality, smaller than some of the prototypes that preceded it. From the perspective of the control electronics, it meant that each PCB would have to be updated to support the new channels, and that our COTS development board would need to be completely redesigned as a custom PCB, as the space required by the original board could no longer be accommodated. Most difficult to redesign were the new Control and Receiver boards, the Control Board due to the much-increased complexity that came with adding an FPGA to a custom design, and the Receiver Board due to the size constraints imposed by its placement within the system's optical processor. Our risk reduction efforts during the 30-Channel system integration provided a head start, but further improvements would have to be made to fit phase detection circuitry for 220 channels into a space not much bigger than that used to fit only 32 of each.

As with the earlier iterations of the system, I was primarily responsible for the design of the Receiver board. My colleague, Petersen, performed the bulk of the Control and Distribution board designs.

4.4.2 Receiver220_RevA

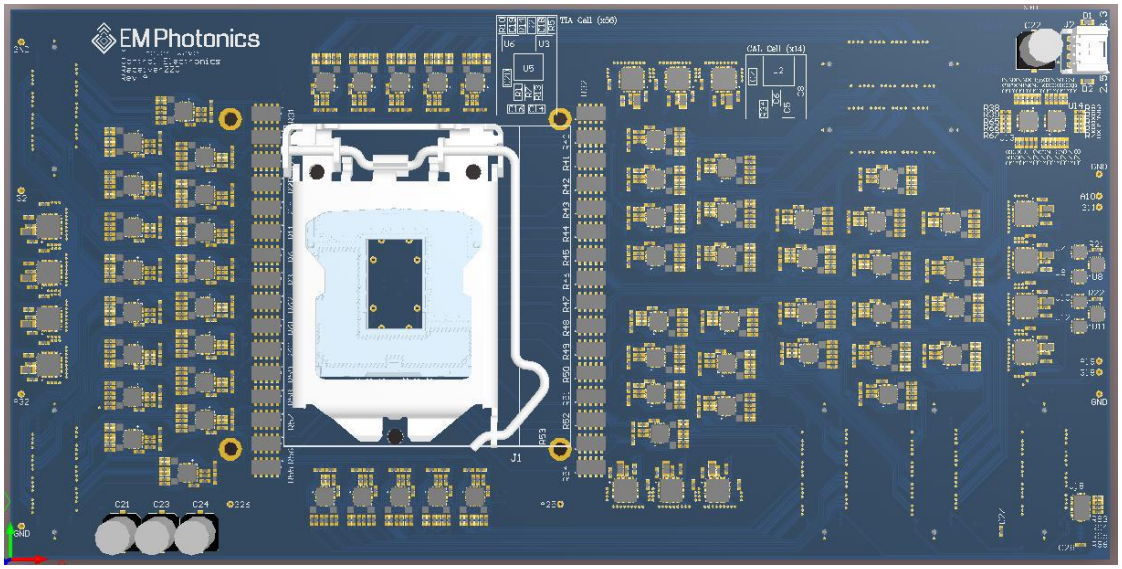

Figure 4.32 Computer-generated 3D rendering of the front face of the Receiver220_RevA PCB

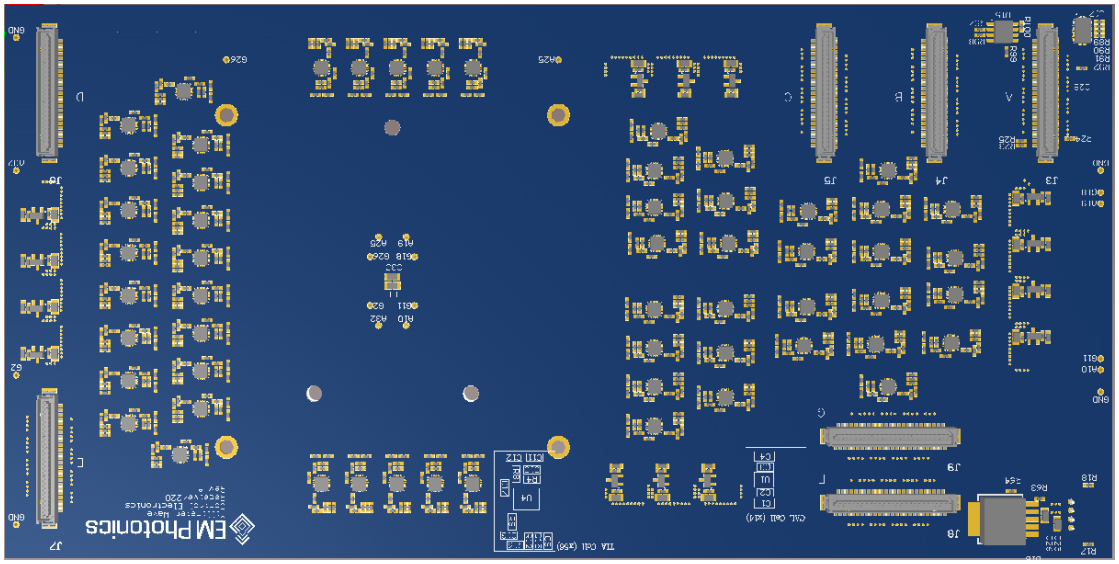

Figure 4.33 Computer-generated 3D rendering of the back face of the Receiver220_RevA PCB

The primary constraint of the Receiver220 PCB was size. As the Receiver board is necessarily located inside the optical processor so as to access the optical signals, these constraints were not negotiable limits. We had tested various space-saving measures in the implementation of the Receiver30_RevC, these were likewise put to use in this design. Additional measures were included to augment space savings, including the addition of more signal routing layers. This change allowed me to keep most of the electrical connections between discrete parts below the surface of the board, which in turn allowed me to devote more of the top and bottom surfaces to component layouts. Another space-saving technique was the elimination of the silkscreen for the majority of the design. This allowed tighter part clearances, whereas earlier designs had to allow for additional space between components for the placement of reference designators in silkscreen. This would add an element of difficulty in testing and troubleshooting exercises, as specific parts would be more difficult to identify, but the reduction in footprint was worth the price. During the 30-Channel system design, we were able to reduce the size of the original unit cells by nearly 75%. With these design modifications we were again able to reduce the unit cell area by more than 50%.

As before, the methodology followed to create the design in repeated unit cells of four channels each, as each operational amplifier IC includes four circuits. The new, smaller, BGA comparators could support only two channels each, so each unit cell comprised a single op-amp chip, two comparator chips, the I2C-controlled digital potentiometer for gain adjustments, and the passive resistors, capacitors, and inductors required by each respective circuit. As such, these components could be arranged in a tightly-optimized design, which the software could then automatically replicate for

any identical circuit configurations. Due to the sheer volume of parts required to support 220 channels, this feature was much appreciated.

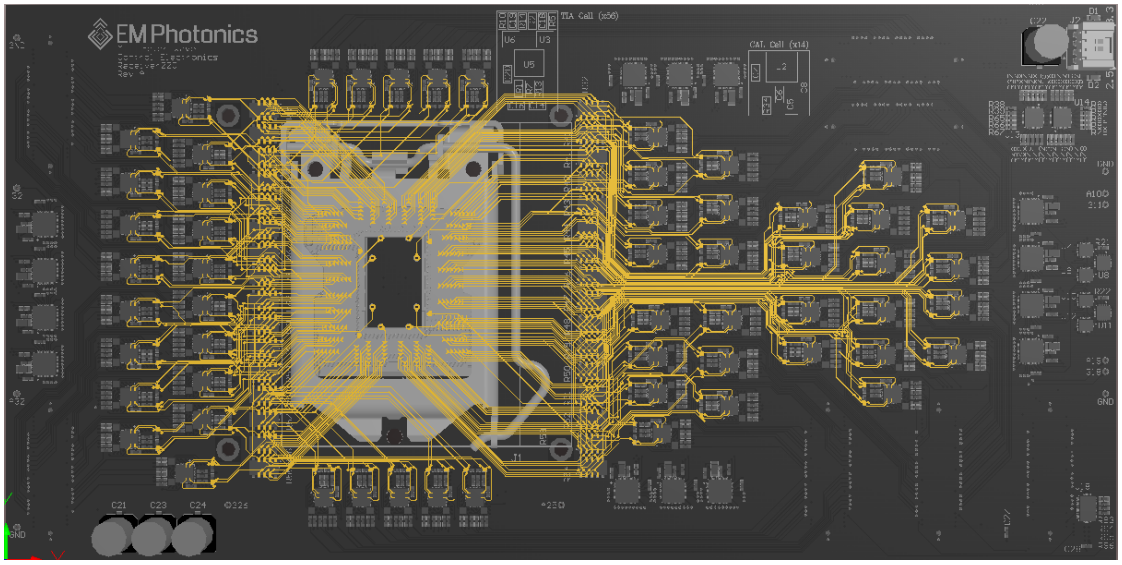

Figure 4.34 Composite image of multiple signal layers showing routing of input signals from the Interposer socket to their respective phase-detection unit cells

One particular challenge in the design of this PCB was the formulation of a design for which the requirements are not completely determined. Due to the complexity of the design, we could not afford to delay its development, lest we fail to meet our deadlines. The size constraints of the Receiver board are largely defined by the space allowances of the optical processor, however, the final design of which had not yet been finalized. As such, the Receiver board design existed in a perpetual state of flux for some time. The primary driver was a determination of what axis the board would be oriented to with respect to the imaging system within the enclosure. Rather than guess which orientation was most likely, I attempted to formulate a design that

could be easily adapted from one constraint to the other in the event of a design change.

Another design complication was uncovered in that it soon became clear that a technician would be able to access only one face of the PCB when installed in the optical processor. As such, it was necessary to place any test points and any connectors that may need to be accessed during operation on the rear face of the board (with respect to the placement of the FPA). Where only one 51-pin FFC had been required for the 30-Channel system, seven would be required now to support the additional channels. In efforts to aid testing, however, I added a constraint to the design that the pinout of the first FFC connector be identical to that of the Receiver30RevC, thus the new receiver board could be used with the electronics of the 30-channel cart if only that connector is used, and the smaller Receiver30RevC could be used to control the A-Bank of channels in the EC Prototype should the need arise.

4.4.3 Distribution55_RevC

Earlier revisions of the Distribution board were designed with an expectation of a 14V modulator $V_{2\pi}$. In actuality, however, this value was found to be higher than expected, approaching the 24V power rail of the Distribution board. The current boards were still able to support this, but later firmware modifications purposed to mitigate noise necessitated still higher voltage requirements. We simulated and then prototyped some potential upgrades utilizing an existing Distribution55_RevB PCB. We disconnected the onboard analog power supplies by removing the LC connection to the main board supplies, substituting a connection to an external tunable power supply. This necessitated the modification of the amplifier circuits with the substitution of IC's designed for use with high-voltage supplies and the replacement of

feedback resistors to more appropriately set channel gains. Our prototype proved successful, and we incorporated these changes into the Distribution55 design, extending the phase feedback voltage rail to a maximum of 41VDC and thereby accommodating a much greater range of $V_{2\pi}$ values.

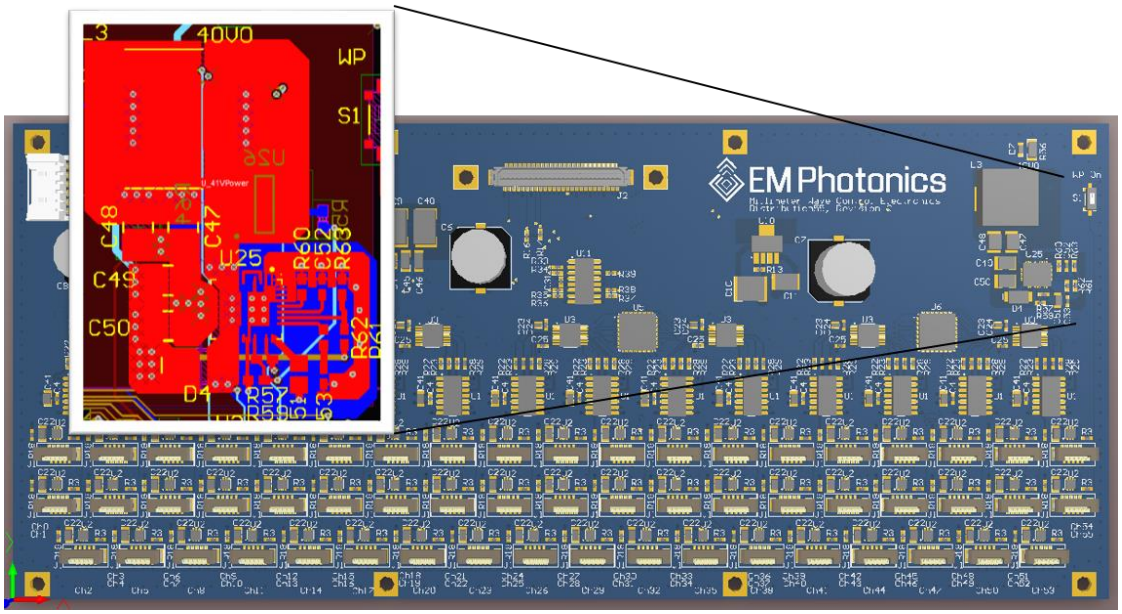

Figure 4.35 Computer-generated 3D rendering of the Distribution55_RevC PCB, detailing the addition of the 41V boost circuitry

4.4.4 Control220 PCB

Central to the operation of the EC Prototype, and a considerable engineering effort was the design of the Control220_RevA PCB. This board effectively replaces the COTS development boards that were used in previous system iterations, housing the control FPGA and forming the central hub of the control system. In addition to providing phase control, this board controls and interfaces with other peripherals of the imaging system and supports the Embedded IAS Platform, enabling the interface to the camera for image acquisition, as well as powering and interfacing with solid-

state storage peripherals for data acquisition. Primary design work for this PCB was completed by Petersen Curt.

This design utilizes a Virtex 6 FPGA, which is flip-chip bonded to the PCB as part of the fabrication process. This FPGA supports up to 600 individual general-purpose I/O connections, of which we utilize 597. While our previously-used COTS development boards included a wide variety of interfaces and features, most of which were unnecessary, this design was tailored to the specific needs of the system, including only those interfaces and features which were either required for operation or deemed potential avenues for expanding functionality in the future. Due to the complexity of the design, this board comprised 16 copper layers, including seven ground planes, two power planes, and 5 inner planes for signal routing, in addition to the two outer surfaces of the PCB.

This design was the first of our custom PCB designs that had to place priority on the routing of high-frequency signals. Some of our other boards make use of 50GHz SPI clocks, which have occasionally posed a mild integration challenge, but the high-frequency signals utilized by the FPGA support systems required more stringent design considerations, including precise trace-length matching of differential signals and precise impedance matching of transmission traces.

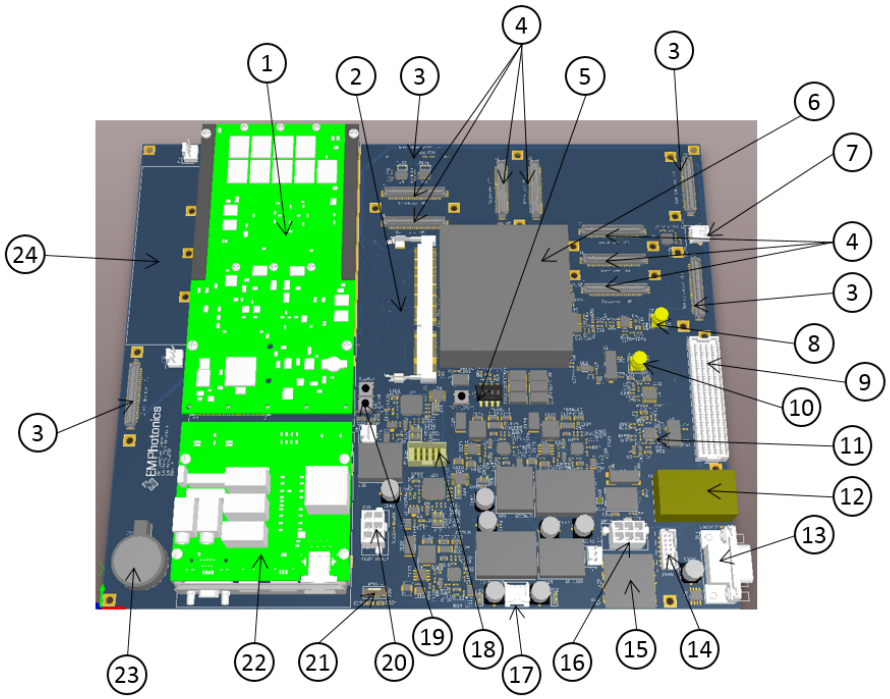

1	XMC slot for COTS processor to perform image acquisition from SWIR.
2	DDR3 SO-DIMM socket connected to FPGA. Supports up to DDR3-1066.
3	Distribution board connectors (x4).
4	Receiver board connectors (x7).
5	FPGA programming. "PROG" momentary switch and "MODE" DIP switches.
6	Xilinx Virtex-6 FF1156 FPGA. Fitted with LX365T-2 in first revision.
7	USB UART for serial debug interface.
8	14-bit reference DAC and amplifier. 12Vpp swing possible.
9	FMC expansion for future capabilities (i.e. Camera Link, cockpit MFD, …)
10	Trigger output. Connect to SWIR camera trigger input.
11	Low-jitter clock synthesizers for PCI-Express reference.
12	GigE Ethernet MAC PHY and magnetics.
13	Laser power and control. Designed to interface with EM4 EM650 laser.
14	JTAG programming connected to FPGA and FMC. Connect to Xilinx programming cable.
15	USB JTAG programming connected to FPGA and FMC.
16	SATA power (3.3V, 5.0V, 12.0V)
17	Receiver board power (3.3V, 2.5V)
18	PMBus programming. Used to program power supplies. Connect to TI USB-to-GPIO.
19	Reset switches for FPGA processor (top) and IAS processor (bottom).
20	Power input (12.0V)
21	I^2C peripheral future expansion. Planned to control power supply.
22	XMC PIM slot. Connected to XMC processor slot for peripheral breakout (video, inputs, etc.)
23	Backup battery for IAS real-time clock.
24	Fans for convection cooled IAS board. Conduction cooled configuration also supported.

Figure 4.36 Computer-Generated 3D rendering the of the Control220 PCB design with annotations and part labels

During hardware verification, we encountered an issue with the board's Ethernet interface, which we were unable to definitively identify. To provide the quickest path to operational capability, we instrumented a workaround using a COTS FMC Ethernet card made by Avnet[8], not unlike our own previous FMC expansion cards used in previous iterations of the control system. This allowed us to continue testing and complete system integration without introducing delays into the schedule.

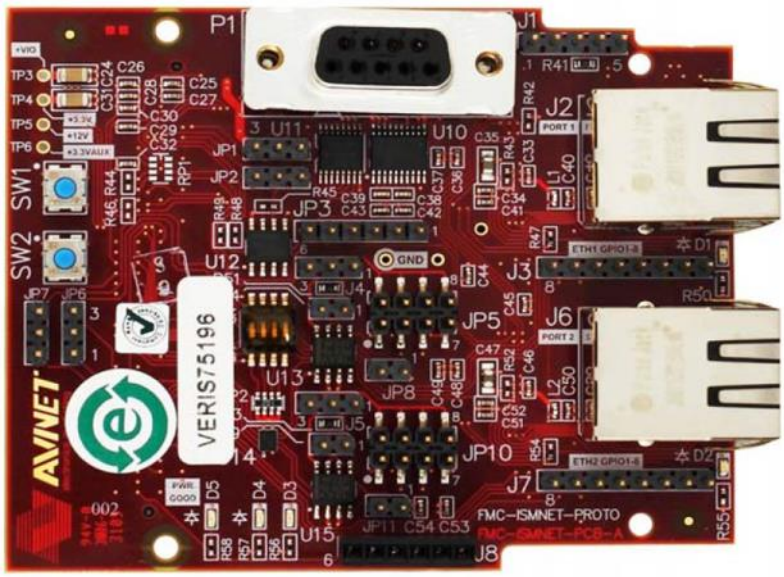

Figure 4.37 COTS Ethernet expansion card which was used as a workaround to our problematic Ethernet interface

Another minor rework was required for the Distribution board interfaces. Where previously the onboard power regulators on each Distribution board had been enabled with a 3.3V pull-up signal, we found that this signal was too weak to reach hysteresis when connected to all four Distribution boards simultaneously. To alleviate this issue we severed the connection to the 3.3V rail and installed a connection to the 5V rail.

4.4.5 Embedded Image Acquisition System

This iteration of the control system also included an embedded SBC (single-board computer) that was tasked with the operation of the IAS. This board is likewise housed within the system enclosure, and is carried and powered by the Control220 PCB along with a separate peripheral interface module. The interface to the camera requires special drivers unsupported for operating systems other than Microsoft Windows, and so our embedded board was programmed with Windows 7 Embedded, a lightweight version of Windows purposed for just such applications.

The physical board, one of the few present in the final system which was not of our design, is a COTS computing module which we acquired from Extreme Engineering Solutions. This board provides the computing power of an Intel i7 processor in a small form-factor.[9]

Figure 4.38 The COTS SBC utilized in our system to interface with the camera and collect frames

4.4.6 Power Distribution PCB

In order to provide power to the various pieces of the system housed within the enclosure, it was necessary to define and design a power-distribution network, which was implemented on another PCB. This PCB receives system power from the aircraft source via a MIL-SPEC connector and provides monitoring and over-current protection for the rest of the system. Additional circuits then translate input power to the levels required by the constituent system parts for distribution. Two 12VDC regulators are included, one to provide power to the Control220 PCB and the other to provide power to the system peripherals, namely the FLIR SC2500 SWIR camera, a Network Switch, and various System Cooling Fans. 24VDC is likewise supplied individually to each of the four Distribution55 PCBs. Power supplied to the Receiver220 PCB is conditioned and provided by the Control220 PCB, and so is not present here.

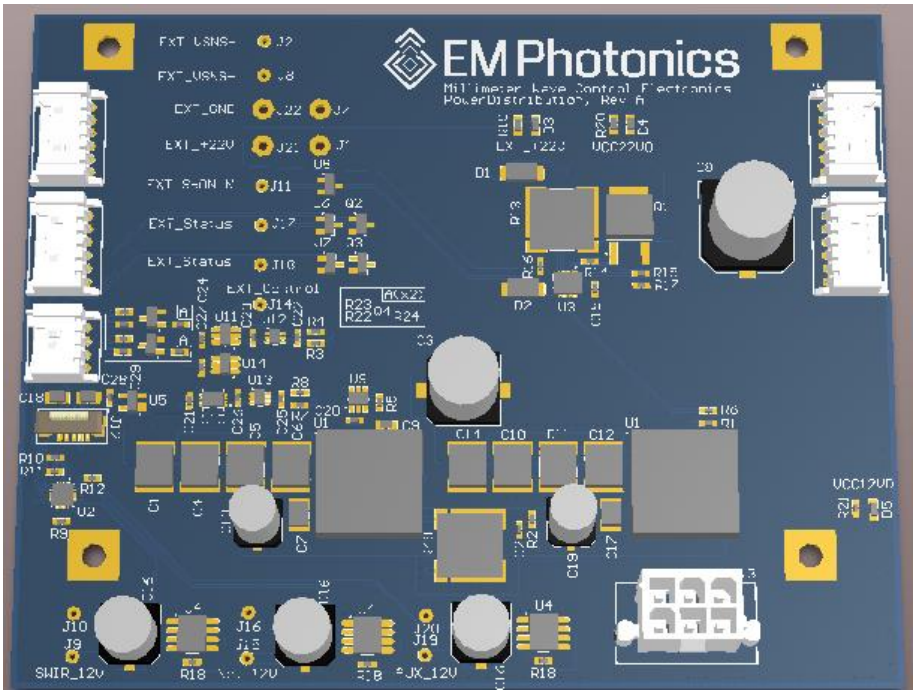

Figure 4.39 Computer-generated 3D rendering of the Power Distribution PCB

This PCB also comprises the system interface to the electronics control rack, allowing limited control functions and status display by routing these signals back to the Control board. Each of these external signals is protected with ESD diodes to prevent damage to the electronics. In addition, a flat-flex connector identical to those used on the Distribution PCB was added with connection to the Control board to allow monitoring and control of the various power supplies.

4.4.7 Firmware and Software modifications

Updates to the operational software and firmware of the control system were necessary to support the hardware expansions. Most of these were straightforward, maintaining identical interfaces to operational components as we used in the 30-Channel version of the system, thus this was largely a matter of channel replication. Other changes were necessitated by the addition of new hardware and functionality, including the Embedded IAS board and the Power Distribution board, neither of which was present in earlier iterations of the system. These changes were also reflected in the LabVIEW Control GUI.

During integration, we determined that we could provide faster framerates by windowing the frames received from the camera. This eliminated superfluous data in that the outer pixels of the image were beyond the field of view of the imager, and such we could save bandwidth by configuring the camera to ignore them. By windowing the output image from the original 320x256 to 224x224 pixels, we were able to demonstrate a capture framerate of 491 FPS, a 40% improvement over our previous performance, without sacrificing functionality.

4.4.8 System Integration

The process of fitting all the constituent components of the EC Prototype system into its enclosure required meticulous planning and execution. As such, each of the individual pieces of the puzzle and its place in the system was modeled to determine the optimal physical placement of each part. An example of one of these models, generated by PSI, showing the final configuration of the system is shown in the following figure. Because of the care involved in the planning of the system, the actual physical assembly was well organized and went smoothly.

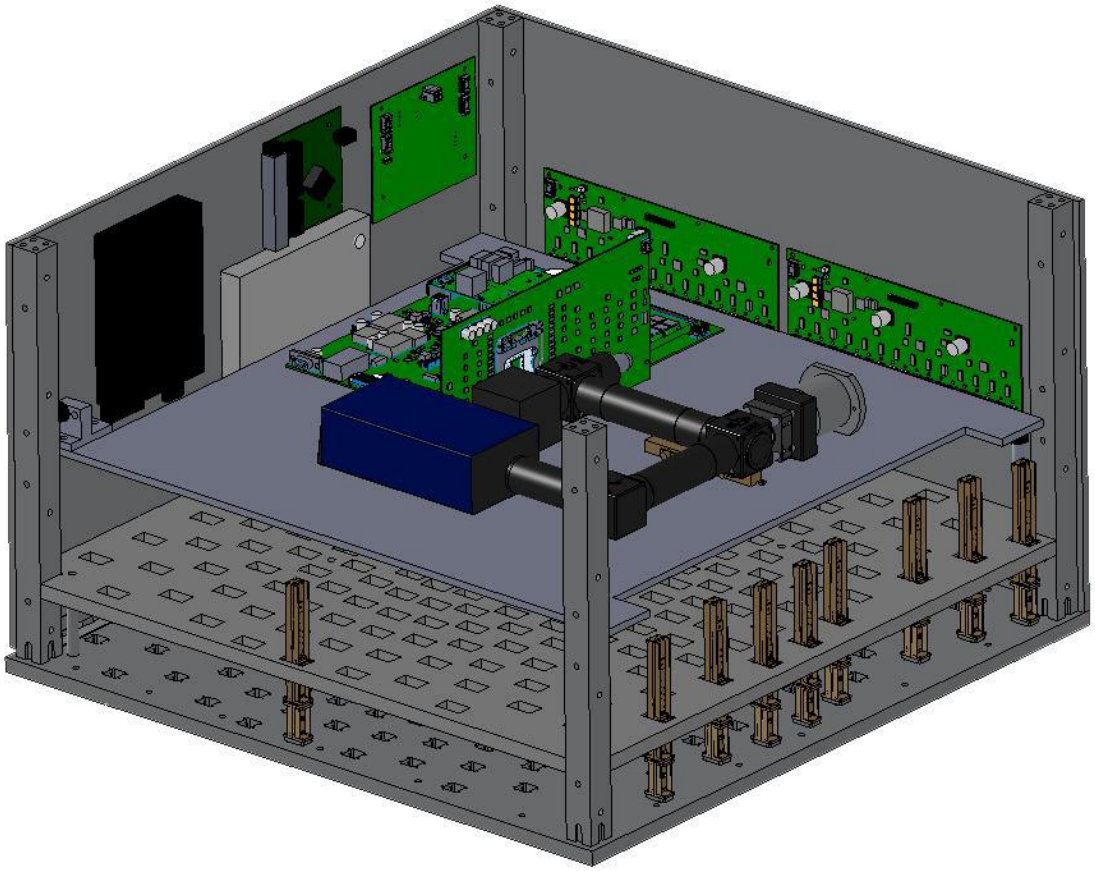

Figure 4.40 Computer-generated 3D rendering of the inside of the EC Prototype enclosure, showing the orientation of the RF modules, the Optical Processor, and the Control Electronics. Generated by PSI.

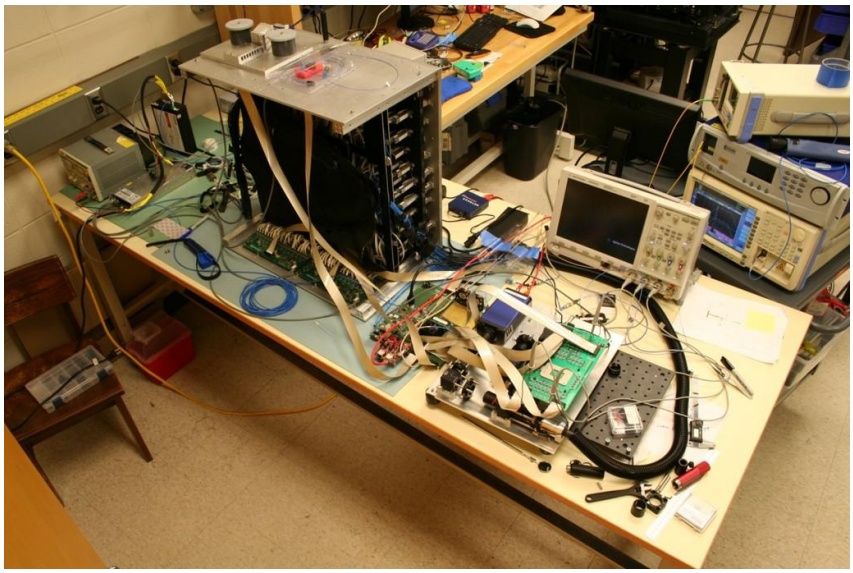

Figure 4.41 The partially-assembled EC Prototype. At this stage, all the electrical components of the system are connected, though not fully installed. This level of integration was the first to allow operational testing of the EC Prototype as a unified system.

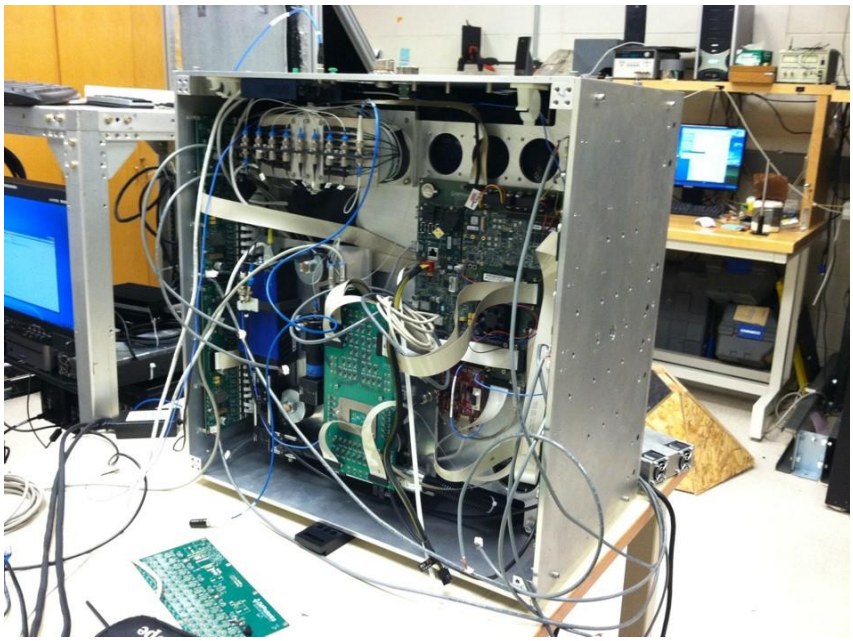

Figure 4.42 The EC Prototype, now with the major operational components of the system fully integrated.

4.4.9 Control Rack

To support flight test operations, it was necessary to devise controls which would allow power control and visual confirmation of system status with the imager casing closed and sealed. This rack would contain any necessary ancillary items such as a source laser, control PC and screen, and power conditioning. These controls would also add a quick shutdown capability in the event of an emergency. As we were unable to certify such a rack for flight testing ourselves, we were furnished with such a rack by NAVAIR which was returned after testing had been completed.

As this rack was intended for support of flight testing, power was intended to be provided by the airframe in the form of 3-phase 400Hz AC, which we would convert to conventional 60Hz AC to power our system. This power would be delivered to the system through a UPS (Uninterruptable Power Source), so as to allow continued operation or an opportunity to power-down the system in the event of power failure. Fed by the UPS were two other power supplies, our primary system power 24VDC supply and a 12VDC supply to power the source laser, which was likewise located in the control rack. A rack-mountable PC was also used, along with a mountable monitor positioned such that additional piloting crews could observe the output of our system during flight so as to provide an in-situ test, as it were, of our system's ability to provide operational awareness during flight operations.

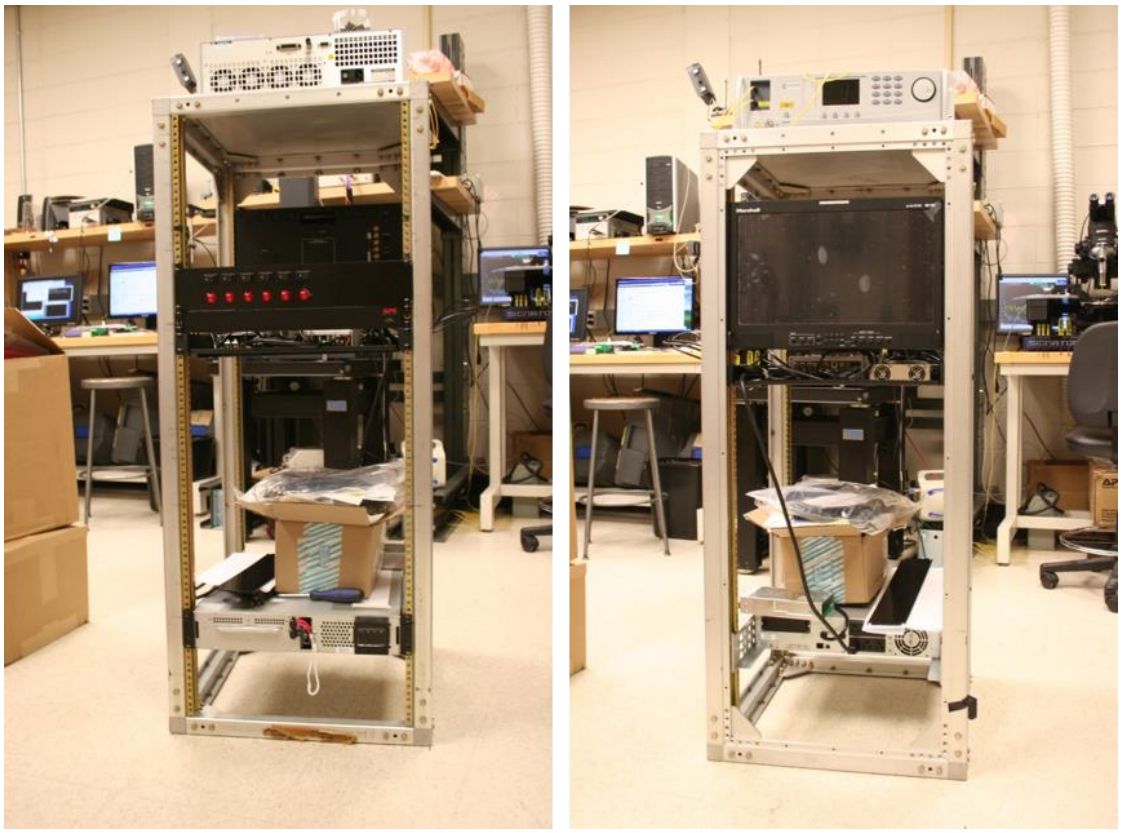

Figure 4.43 Front and back view of our custom electronics control test rack

Chapter 5

EC PROTOTYPE FIELD TESTING

The first field test of the EC Prototype was conducted at Yuma Proving Ground (YPG) in Yuma, AZ, in cooperation with the Office of Naval Research and in conjunction with a DVE testing event comprising several sensors. Originally anticipated to be a flight test, delays in the flight qualification process encountered by another sensor system precluded our system from undergoing the same qualification. As a result, our participation in the test was limited to ground-based imaging through H-53 generated dust clouds. This fundamentally changed the parameters of the test in a manner detrimental to the operation of the EC Prototype, as cold sky reflections are best seen with a downward look angle, such as would be present imaging from an aerial platform. Imaging from ground level, however, necessarily constrains the downward look-angle of the system to only a few degrees from the horizon. Nonetheless, PSI was able to prove video rate penetration of the dust clouds with their sensor and our electronics

Figure 5.1 An H-53 flyover pass during field testing in Yuma, AZ

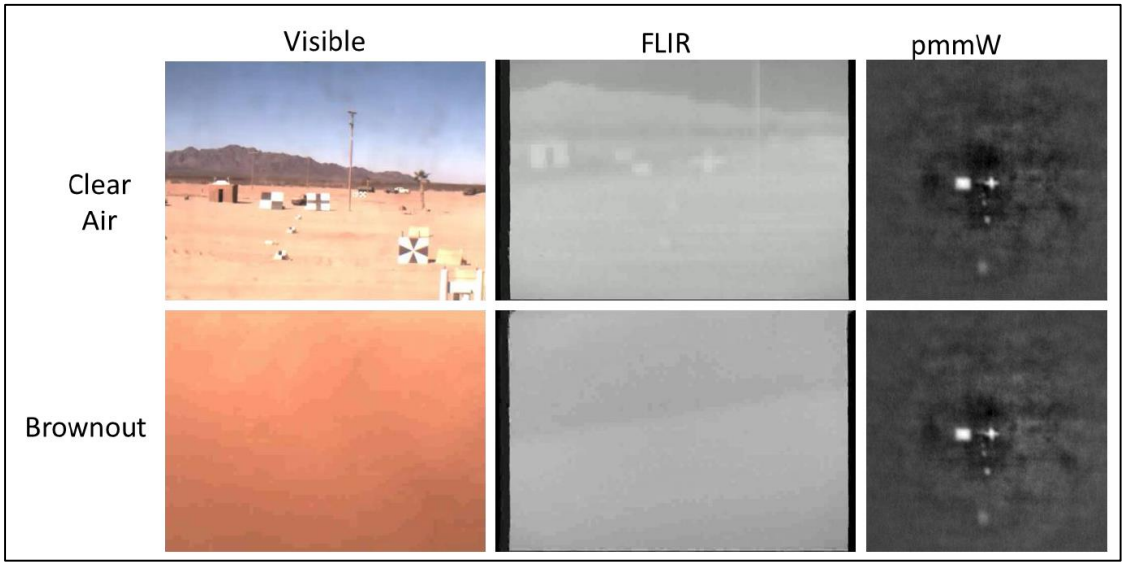

Figure 5.2 A comparison of imaging modalities used in field testing, in the presence of both clear air and brownout conditions

Testing took place over a span of two weeks with a specially-prepared H-53 helicopter operated by NAVAIR. At one point, testing was temporarily halted due to a malfunction involving the aircraft's rotor and was not resumed for several days. During this time, as we still were afforded access to the test range, we collected data using static range features and our moving automobiles as targets. Upon resuming, flight testing was concentrated into a two day period, wherein the test aircraft performed several hover and fly-by operations. Our team was able to collect a considerable amount of data, both from the flight operations and our own downtime tests, which showed very little image degradation due to brownout, even very heavy brownout.

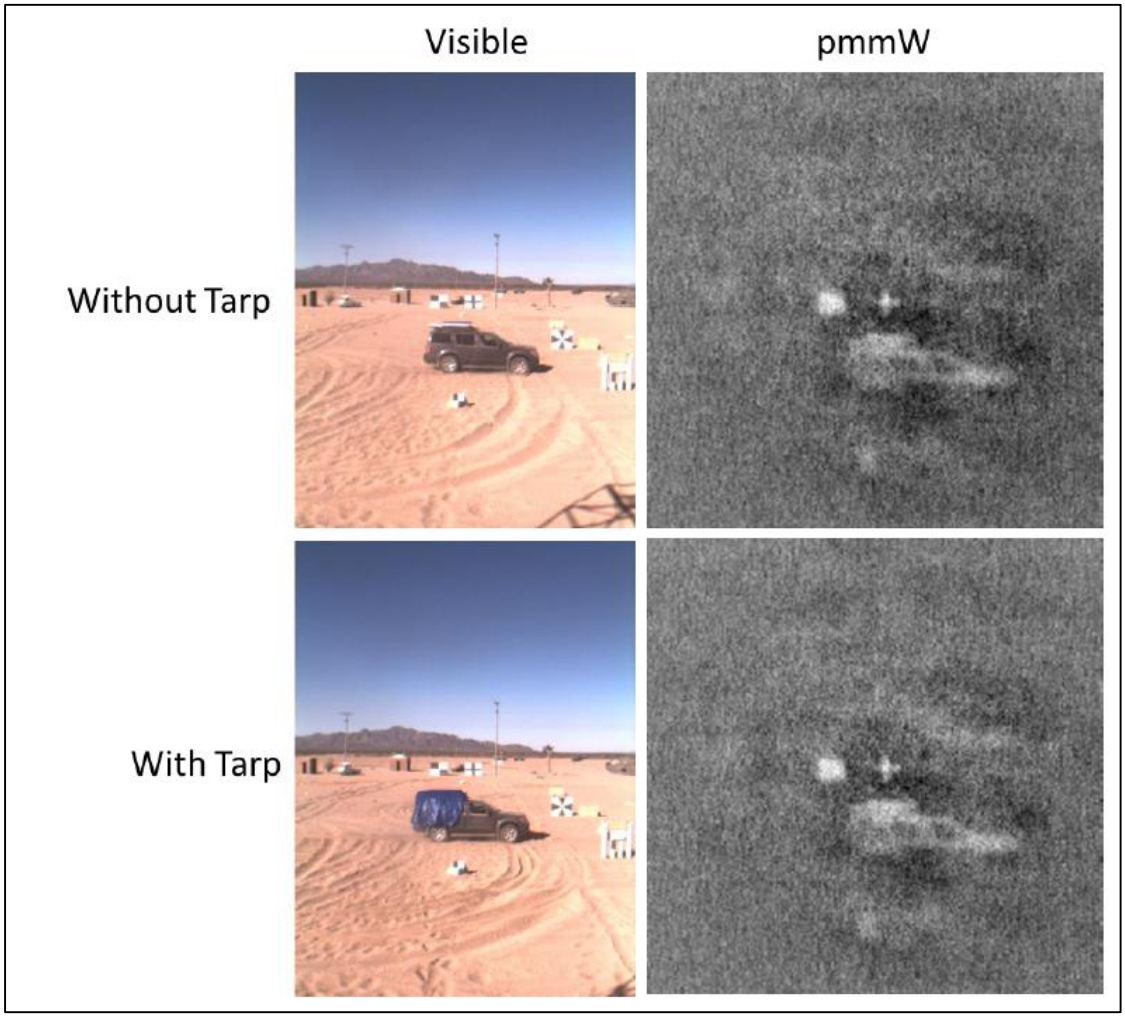

Figure 5.3 Imagery taken during field testing of an automobile with a metal tube affixed to its roof, simulating an IRAM (Improvised Rocket Assisted Mortar). The simulated IRAM is still visible in the passive MMW imagery when hidden from sight by a tarp.

Perhaps chief among the challenges encountered during testing with respect to the control system was maintaining adequate cooling. Testing in a sandy environment meant that the imager casing must be kept closed, which restricted air flow, and the ambient temperatures were likewise very high. The IAS SBC would often report CPU

temperatures in excess of 80°C and up to 100°C, which would cause throttling and occasional crashes.

Even while testing is taking place, development never stops. During the course of the testing event, we would collect data every day which we would backup to our servers in Delaware each night, providing both data security and an opportunity for other engineers at home to review data for possible enhancement capability. While reviewing captured data from the day before, such a realization was made in the observation that histogram of our video outputs were overly quantized, limiting contrast. This effect was traced to the interface between our IAS software and the LabVIEW VI. The IAS software interface was configured to send frame data over the network as floating-point data, but the LabVIEW VI performed its processing using unscaled integer data, resulting in a greatly reduced set of available levels in the image depending on the contrast of the scene. I was able to fix this issue in the field by modifying the LabVIEW processing pipeline to utilize the full-range data from the IAS software. Sample histograms that illustrate this issue, as well as the implications of this issue for final imagery, are shown in the following figures.

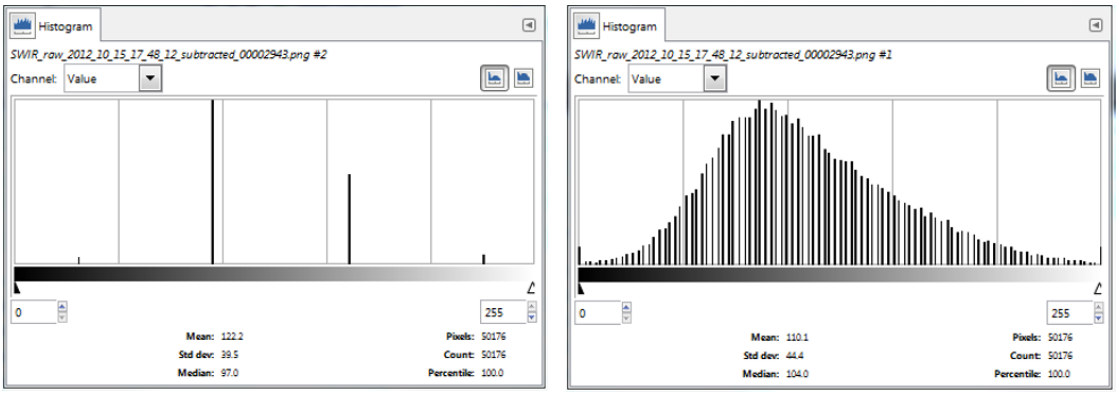

Figure 5.4 Histograms of processed images in LabVIEW, before and after field modifications to the processing pipeline. The histogram on the left is a particularly bad example, showing only 4 pixel values, but this was not typical of normal operation

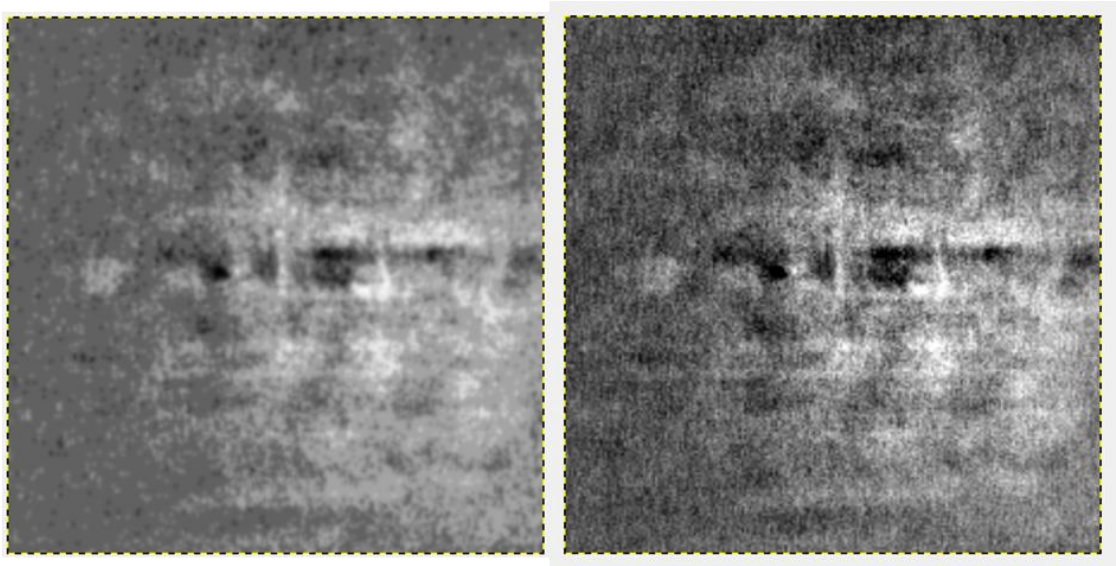

Figure 5.5 Images corresponding the histograms shown above. The figure on the left suffers a severe loss of contrast.

Chapter 6

CONCLUSION

Brownout continues to be a formidable threat to rotary-wing aircraft operating in unimproved areas. Solutions to this problem are being sought, and new imaging modalities are being developed to combat this threat. One such imaging technique has been pioneered by the engineers at Phase Sensitive Innovations who have created an imaging system which utilized passive millimeter waves to see through obscurant clouds. This system is the first of its kind, demonstrating impressive performance, and requires innovation in control methodologies to fully realize its potential. Over the course of my work on this project in conjunction with the University of Delaware, PSI, and EM Photonics, I have contributed to the development and implementation of several iterations of a phase-synchronizing control system specifically tailored to the needs of this imaging system modality. With my help, a prototype of this system has been demonstrated to penetrate brownout generated by large rotary-wing aircraft without significant degradation.

This project has taught me a great deal about engineering, the application and execution of knowledge, the formulation and implementation of ideas, and about the importance of both experience and flexibility. I hope to apply the things that I have learned to continue creating innovative solutions to important problems that make a difference to the world.

REFERENCES

[1] Daniel G. Mackrides ; Christopher A. Schuetz ; Richard D. Martin ; Thomas E. Dillon ; Peng Yao, et al. "Progress toward a video-rate, passive millimeter-wave imager for brownout mitigation", Proc. SPIE 8022, Passive Millimeter-Wave Imaging Technology XIV, 802203 (May 25, 2011); doi:10.1117/12.884164; http://dx.doi.org/10.1117/12.884164

[2] Curt, Petersen F., and Michael R. Bodnar. Synchronizing Phases of Multiple Optical Channels. EM Photonics, Inc, assignee. Patent US8897656 B2. 25 Nov. 2014. Print.

[3] Yujiri et al., IEEE Microwave Magazine, April 2003

[4] Martin et al., "Optical Up-conversion Enables Capture of Millimeter-wave Video with an IR Camera." SPIE Newsroom. 13 Aug. 2012. Web. <http://spie.org/x89063.xml>.

[5] "AN/AAQ-37 Distributed Aperture System (DAS) for the F-35." *Northrop Grumman*. Web. <http://www.northropgrumman.com/capabilities/anaaq37f35/pages/default.aspx>.

[6] "LabVIEW System Design Software." *NI LabVIEW*. Web. <http://www.ni.com/labview/>.

[7] "OpenCV | OpenCV." *OpenCV*. Web. <http://opencv.org/>.

[8] "ISM Networking FMC Module." Avnet. Web. <http://www.em.avnet.com/en-us/design/drc/Pages/ISM-Networking-FMC-Module.aspx>.

[9] "XPedite7301." Extreme Engineering Solutions (X-ES). Web. <http://www.xes-inc.com/products/view/xpedite7301/>.