

2013

Integration And Measurements of a Ka-Band Interferometric Radar in an Airborne Platform

Rockwell B. Schrock

University of Massachusetts Amherst

Follow this and additional works at: <https://scholarworks.umass.edu/theses>

 Part of the [Electromagnetics and Photonics Commons](#), [Signal Processing Commons](#), and the [Systems and Communications Commons](#)

Schrock, Rockwell B., "Integration And Measurements of a Ka-Band Interferometric Radar in an Airborne Platform" (2013). *Masters Theses 1911 - February 2014*. 1011.

Retrieved from <https://scholarworks.umass.edu/theses/1011>

This thesis is brought to you for free and open access by ScholarWorks@UMass Amherst. It has been accepted for inclusion in Masters Theses 1911 - February 2014 by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

INTEGRATION AND MEASUREMENTS OF A
KA-BAND INTERFEROMETRIC RADAR IN AN AIRBORNE PLATFORM

A Thesis Presented

by

ROCKWELL B. SCHROCK

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

MASTER OF SCIENCE IN ELECTRICAL AND COMPUTER ENGINEERING

February 2013

Department of Electrical and Computer Engineering

INTEGRATION AND MEASUREMENTS OF A
KA-BAND INTERFEROMETRIC RADAR IN AN AIRBORNE PLATFORM

A Thesis Presented

by

ROCKWELL B. SCHROCK

Approved as to style and content by:

Paul R. Siqueira, Chair

Dennis L. Goeckel, Member

Robert W. Jackson, Member

Christopher V. Hollot, Department Head
Electrical and Computer Engineering

To my parents, and their unwavering support of all my endeavors.

ACKNOWLEDGEMENTS

I am extremely grateful to Professor Paul Siqueira for bringing me on board at a time when I wasn't sure I even wanted to do a thesis. His guidance, insight, and encouragement – as well as his faith in me to complete the project – were all paramount to the success of this thesis. Thank you to professors Robert Jackson and Dennis Goeckel for providing critical feedback as members of the thesis committee, and to Professor Steve Frasier for providing an independent study to get my feet wet in MIRSL. I also owe a big thanks to Tony Swochak who encouraged me to pursue the thesis track. Perhaps it was in his own interest, as I was able to check off one of the items proposed as “Future Work” of his thesis[14]. Thanks to Caitlin Dickinson for being my lunch buddy, commiserating with me, and politely listening as I went on and on about interferometry. Finally, thank you Jess for unending encouragement, for being understanding of my long hours in the lab, and for bringing a hot meal on those late nights.

ABSTRACT

INTEGRATION AND MEASUREMENTS OF A KA-BAND INTERFEROMETRIC RADAR IN
AN AIRBORNE PLATFORM

FEBRUARY 2013

ROCKWELL B. SCHROCK, B.S., UNIVERSITY OF CONNECTICUT

M.S., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Paul R. Siqueira

The Topographic Interferometry Mapping Mission (TIMMi) instrument is a unique millimeter wave interferometric radar system operating at 35 GHz (Ka-band). It was constructed in part to advance the technology readiness level of NASA's Surface Water and Ocean Topography (SWOT) mission, a spaceborne platform that will globally map the altimetry of Earth's water to gain insight into surface water interactions and dynamics. Previous ground deployments of TIMMi were successful in demonstrating the abilities of the system from a stationary platform. The next logical step was to move TIMMi closer to space by installing it on an airborne platform prove its capability in mapping wide swaths of land at a higher incidence angle. This thesis outlines the design considerations and challenges in adapting TIMMi to a small airborne platform. Documentation is included from many points throughout the development cycle, including hardware and software development, flight planning, data acquisition, and post-flight data processing.

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	iv
ABSTRACT	v
LIST OF TABLES	viii
LIST OF FIGURES	ix
CHAPTER	
1. INTRODUCTION AND MOTIVATION	1
2.1 SWOT Mission Overview	1
2.2 The March Towards Higher Frequencies	2
2.3 Prior TIMMi Development	3
2.4 Present TIMMi Development	4
2. RADAR CONCEPTS	7
2.1 FMCW Basics	7
2.2 Interferometry Basics	11
2.3 Platform Geometry	13
3. TIMMI SYSTEM OVERVIEW AND MODIFICATIONS	15
3.1 Redesigning the Supporting Electronics	16
3.1.1 Power Amplifier Power Supply	20
3.1.2 Power Connector Pinout	20
3.2 Modifying the Aircraft Door	22
3.3 Rewriting the Acquisition Software	24
3.3.1 Digitizer Re-arm Time	24
3.3.2 RF Data File Format	26
3.4 Integration with AIMS-1	29
3.4.1 LIDAR Data File (.las)	29
3.4.2 GPS Data File (.gps)	30
3.4.3 IMU Data File (.wbd)	32
4. CHALLENGES	33
4.1 DUC/DDC Voltage Regulator Interference	33
4.2 Near-Field Antenna Coupling	35
4.3 Determining Deployment Configurations	36
4.3.1 Flight Lines and Times	36
4.3.2 Flight Altitude and Antenna Configuration	37
4.3.3 Baseband Signal Acquisition	39
4.3.4 Function Generator Configuration	41

5. DATA PROCESSING TECHNIQUES	42
5.1 Data Processing Flow	42
5.2 Workflow	42
5.2.1 Generating Previews	44
5.2.2 Setup	44
5.2.3 Projecting to Map Coordinates	45
5.3 Recovering Aircraft Roll Angle	47
6. INAUGURAL RESULTS	49
6.1 Summary of Airborne Deployments	49
6.1.1 February 3, 2012	49
6.1.2 April 13, 2012	49
6.1.3 May 18, 2012	50
6.1.4 June 11, 2012	50
6.2 Reflectivity	50
6.3 Interferometry	58
6.3.1 Variations From The DEM	59
6.3.2 Error Sources (NEW)	61
7. SUMMARY AND FUTURE WORK	62
7.1 Contributions	62
7.2 Future Work	63
7.2.1 Combining AHRS and TIMMi Datasets	63
7.2.2 Unfocused and Focused SAR	64
7.2.2.1 Work So Far	65
7.2.3 Correcting Target Range Blurring	66
7.2.4 Additional Deployments	66
7.2.5 Convert RF Datasets to a Standardized Format	67
7.2.6 Add Flexibility to Antenna Mounting Bracket	68
APPENDICES	
A. NOTATION REFERENCE	69
B. MATLAB PROCESSING CODE REFERENCE	71
C. CHECKLISTS	77
BIBLIOGRAPHY	81

LIST OF TABLES

Table	Page
3.1. Power supply 18-pin circular connector pinouts for three iterations of TIMMi power supplies. Voltages marked with an asterisk (*) are always on during standby, and blank entries are not connected (N/C). All ground (GND) terminals are connected.	21
3.2. TIMMi system power draw; during standby, all subsystems are on except the transceiver	22
3.3. TIMMi RF file format, version 1	27
3.4. TIMMi RF file format, version 2	27
3.5. An example layout of a version 2 RF data file with N_{samp} samples per pulse and N_{pulse} pulses per look	28
3.6. AIMS-1 LIDAR data file format	30
3.7. AIMS-1 IMU data file format	31
4.1. Example configurations for the look geometry of Figure 4.3 for $\theta_B = 45^\circ$; distances are approximated for convenience; the recommended configuration is in bold	38
4.2. Tektronix AFG3252 configuration settings for a 1 ms chirp (1 kHz PRF)	40

LIST OF FIGURES

Figure	Page
1.1. SWOT Ka-band interferometric diagram (Source: NASA JPL)	2
1.2. A deployment of TIMMi Ku-band on Mount Sugarloaf on June 30, 2011	4
1.3. A test deployment of the Ka-band system on the UMass Amherst campus	5
1.4. The author setting up for the first flight of the Ka-band system	5
2.1. Two targets (top) and the FMCW frequency response of those targets in comparison to the transmitted signal (bottom).	8
2.2. Block diagram of a generic FMCW transmitter and receiver	9
2.3. A time series of a single linear chirp, downconverted, compressed, and sampled at baseband	10
2.4. Interferometric antenna configuration	11
2.5. Interferometric geometry over (a) a flat earth and (b) an existing DEM	12
2.6. Airborne platform geometry nomenclature	14
3.1. System overview block diagram. Connection types between each component are given in the legend.	16
3.2. TIMMi transceiver block diagram showing DUC, DDC, oscillators, antennas, and FMCW downconversion chains.	17
3.3. The new “flat” supporting architecture	18
3.4. Layout of the new 2U power supply enclosure	19
3.5. The rear cargo doors of the Cessna 206 open like a clamshell with no center pillar	22
3.6. TIMMi Ka-band antennas mounted on a Cessna 206 cargo door. The top two slotted waveguide antennas receive, and the bottom antenna transmits.	23
3.7. The interior of the aircraft door showing mounted transceiver and flexible WR-28 waveguides	23
3.8. Residual delay in recording data due to the operating system, where zero is the nominal time	28
3.9. An example of two sets of NMEA 0183 strings with GPS data	30

4.1. Power distribution block diagram	34
4.2. Undesired voltage regulator tones from DUC and DDC in a baseband receive channel with (a) no RF chokes installed, and (b) three turns of the DC wire around the DUC choke	35
4.3. Antenna cross-track look geometry	37
5.1. Data processing flow	43
5.2. A detail view of the domain mapping scheme used to generate magnitude and interferometric images in map coordinates	43
5.3. (a) The entire map area is examined to find (b) the first beam. Then (c) a mask is created around that beam and (d) the next beam is located within this mask, a subset of the map. Steps (c) and (d) continue for every subsequent position. . . .	46
5.4. Recovering aircraft roll angle	48
6.1. The planned flight lines for June 11 overlaid with the actual platform track	51
6.2. (a), (b) Reflectivity and (c) correlation in map coordinates	52
6.3. (a) Simulated, (b) measured, and (c) differential interferometric phase in map coordinates	53
6.4. (a) Heights from interferometry, (b) the DEM, and (c) these two added together. Note that (a) has a smaller vertical scale to bring out subtle height variations. . .	54
6.5. Reflectivity for one flight line in radar coordinates, normalized to the peak return power	55
6.6. A composite reflectivity image for five flight lines taken on June 11, 2012; the solid black lines indicate the aircraft track; north is up	56
6.7. A worst-case example of bright targets spreading in range and azimuth	57
6.8. Magnitude spectrum from a single look	58
6.9. Interferogram (top) and aircraft roll angle (bottom)	59
6.10. Elevation profile from near-range (right) to far-range (left) of base DEM (solid) and the new DEM (X's)	61

CHAPTER 1

INTRODUCTION AND MOTIVATION

1.1 SWOT Mission Overview

NASA's Surface Water and Ocean Topography (SWOT) mission plans to map the altimetry of at least 90% of Earth's ocean and fresh water to within centimeter height accuracy. Ocean topography measurements are currently taken by buoys placed throughout the globe and are used to determine sea surface temperature and changes in ocean currents. Measurements made by buoys miss mesoscale trends in the vast areas between points. Gauges along rivers also give point measurements along a river's course, which can be used to determine flow rates and predict the availability of freshwater resources. Similar to ocean buoy measurements, gauges of river height are point measurements, and interpolation between measurements must be done to determine the overall gradient of river height.

Scientists are particularly interested in monitoring ocean heights for insights into climate change. Warm water expands, so it tends to bulge upward slightly. The NASA SWOT mission will be accurate enough to measure these tiny height differentials over a wide swath. A global picture of ocean topography tells us a lot about heat storage and transport in the Earth's oceans, allowing us to study how these processes tie in to climate change. Local water heights of inland rivers are also valuable for flood risk mitigation and the study of seasonal flooding patterns.

The primary instrument on this mission will be a Ka-band interferometer (Figure 1.1) with operating characteristics comparable to the instrument discussed in this thesis. The Microwave Remote Sensing Laboratory (MIRSL) here at the University of Massachusetts in Amherst has been funded in part by NASA to develop prototype RF subsystems for SWOT. In doing so, we hope to increase the technology readiness level (TRL) for this relatively new technology. Developing, building, and testing our own system enables us to document real-world challenges




Figure 1.1: SWOT Ka-band interferometric diagram (Source: NASA JPL)

and measurements throughout the process. This is an ongoing project at MIRSL, with many different areas being explored at this time of writing.

1.2 The March Towards Higher Frequencies

Spaceborne interferometry has historically been conducted at frequencies lower than Ka-band. In 2000, the Shuttle Radar Topography Mission (SRTM) provided the first global digital elevation model (DEM) at a spatial resolution of 90 m using C-band technology (5.3 GHz) [4]. At the time of writing, another ambitious global mapping mission is underway by the German Aerospace Center (DLR), TanDEM-X, whose twin X-band (9 GHz) satellites flying in formation will produce a global DEM at 15 m pixel resolution with unprecedented vertical resolution[10].

As solid-state RF technology matures, it becomes more reliable and power-efficient to move microwave systems to higher frequencies. Consequently, the transition up the spectrum is a moving target for radar systems, and it offers a number of advantages. Most importantly, the range resolution of a radar improves linearly with bandwidth, which is an important metric in determining the accuracy and error of a radar to take range measurements. A 1% fractional bandwidth at Ka-band is 350 MHz, but the same fractional bandwidth at C-band is just 50 MHz

– a range resolution that is seven times worse. In addition, radio spectrum licensing follows the same trend as higher-frequency devices become commercially available, so radars must move even higher to escape spectrum congestion.

Size and weight constraints are extremely important considerations for spaceborne applications, especially in the absence of the Space Shuttle to deliver large payloads to low Earth orbit. Regardless, heavier satellites will always be more expensive to put into orbit, and cost is always a consideration. RF circuitry, transmission lines, and antennas can be scaled down to provide similar operating characteristics in a smaller package at reduced wavelengths.

For interferometry, the baseline separation between the two antennas can also be reduced. SRTM was flown by NASA in February of 2000 at an orbital height of 230 km, requiring a baseline of 60 m at a wavelength of 5 cm. In contrast, SWOT is meant to have an altitude of 800 km, which would require a baseline of roughly 240 m to maintain the same performance. Targets that change quickly over time (e.g. bodies of water) must be imaged in a single pass, because multi-pass interferometry is subject to temporal decorrelation. However, 240 m is not realistic for a single stable platform, so it is necessary to use a higher frequency so a proportionally shorter baseline can be used. This is another driving factor in the move towards higher-frequency instruments. In fact, the compactness of TIMMi’s antennas was the enabling factor in deploying this instrument on a small airborne platform, as discussed later.

1.3 Prior TIMMi Development

Development of the TIMMi system thus far has been documented by UMass students Karthik Venkatasubramanian[16], Harish Vedantham[15], Anthony Swochak[14] for the Ka-band and Ku-band subsystems. Development of the Ku-band system was done first, with the design of the dual upconverter (DUC) and dual downconverter (DDC) circuitry, the supporting systems, and a set of planar patch array antennas. Then the Ka-band system was designed and built. The Ku-band patch antennas were then replaced with the slotted waveguide antennas. Both systems share data acquisition hardware, software, and data processing code, reducing the overhead of having two completely separate systems.

Both systems have been deployed at a number of ground locations, including the roof of the Lederle Graduate Research Center on the UMass Amherst campus, Mount Holyoke in



Figure 1.2: A deployment of TIMMi Ku-band on Mount Sugarloaf on June 30, 2011

Hadley, MA, and Mount Sugarloaf in Sunderland, MA (Figure 1.2). These locations were chosen on the merit of having low incidence angles in the interest of maximizing signal return. In addition, a high lookout enables the mapping of a wide swath of land.

One interesting undergraduate experiment demonstrated the ability of TIMMi to resolve small displacement over longer periods. By pointing the radar at the 300-foot tall W. E. B. Du Bois library during a windy day on campus, a periodic building sway on the order of millimeters could be detected.

1.4 Present TIMMi Development

Given a working Ka-band interferometric system, the challenge has been to adapt the ground-based system (Figure 1.3) onto an airborne platform (Figure 1.4). This seemingly simplistic goal was met with a number of challenges in all areas of the radar. Before doing anything new, intimate knowledge of the RF subsystems had to be obtained in order to resolve any issues with these components. Platform constraints required rethinking of the mechanical support structures and power supply architecture. The acquisition of long, continuous data sets was contingent upon rewriting the data acquisition software that could also merge in real-time GPS data. Post-flight data processing programs had to be written to take into account platform motion from discrete onboard instruments. These programs also had to handle the sheer amount of



Figure 1.3: A test deployment of the Ka-band system on the UMass Amherst campus



Figure 1.4: The author setting up for the first flight of the Ka-band system

data from these long acquisitions.

This thesis touches on all of these areas, as well as showing some promising results from some actual airborne deployments. The topographic measurements presented are primarily images of the Earth and not of water. In order to achieve coherence over a body of water, the incidence angle must be very steep. In other words, the radar must be looking almost straight-down, which is not physically possible with our current antenna hardware. Despite this, we are confident that the results are indicative of the overall performance of this SWOT prototype.

CHAPTER 2

RADAR CONCEPTS

2.1 FMCW Basics

The word radar was derived as an acronym of RAdio Detection And Ranging. The name alone reveals radars are very good at: determining range to a target. A traditional pulsed radar sends a short, high-powered pulse of RF energy and records the time response of the echo. The time-of-flight of the pulse,

$$t = \frac{2r}{c}, \quad (2.1)$$

can be measured very accurately by sequential sampling of the received signal. Given prior knowledge of c , the speed of light, the range r can be determined[12].

By contrast, a frequency-modulated, continuous wave (FMCW) has a 100% duty cycle, which means it transmits continuously. The transmitted signal is a frequency-modulated linear “chirp” that spans a wide bandwidth. The received signal can be integrated over the entire chirp duration τ to provide significant signal gain. In an abstract sense, the chirp essentially encodes additional information into the transmitted signal which can be extracted in the receiver to increase range resolution.

Figure 2.1 shows an FMCW radar with a transmit antenna, receive antenna, and two targets at ranges r_1 and r_2 . The targets are assumed to be stationary relative to the radar. At time t_0 , the transmitter begins emitting a linear FM chirp of bandwidth B in the frequency range from f_1 up to f_2 . The chirp has a time duration τ , so the slope of the sweep is B/τ [Hz/s]. The echo from each target is a replica of the transmitted chirp, but it is delayed due to the time of flight of the signal. In particular, for a target at range r ,

$$\Delta t = t - t_0 = \frac{2r}{c}. \quad (2.2)$$




Figure 2.1: Two targets (top) and the FMCW frequency response of those targets in comparison to the transmitted signal (bottom).

The transmit frequency is constantly changing, so at time t the frequency of the echo differs from the frequency of the currently-transmitted signal. This frequency difference, also known as the beat frequency, is related to the slope of the chirp by

$$\Delta f = \frac{B}{\tau} \Delta t. \quad (2.3)$$

For a target at a given range, a faster chirp (smaller τ) will result in a bigger Δf , pushing the frequency response of that target higher. Targets at zero-range have $\Delta f = 0$, or DC.

Because we are sampling continuously, it is not possible to measure Δt directly. However, it is very easy to measure Δf by simply multiplying the received echo by the transmitted signal. Recall that a mixing operation produces the sum and difference frequencies between two signals. By adding a low-pass filter to the output of the FMCW mixer, the sum frequencies are rejected, and the desired difference frequencies remain. A block diagram of this operation is shown in Figure 2.2. The mapping between Δf and Δt , above, means the range of that particular frequency can be recovered.

Given that sampling occurs at f_s , the maximum recoverable Δf is $\Delta f_{max} = f_s/2$ per the

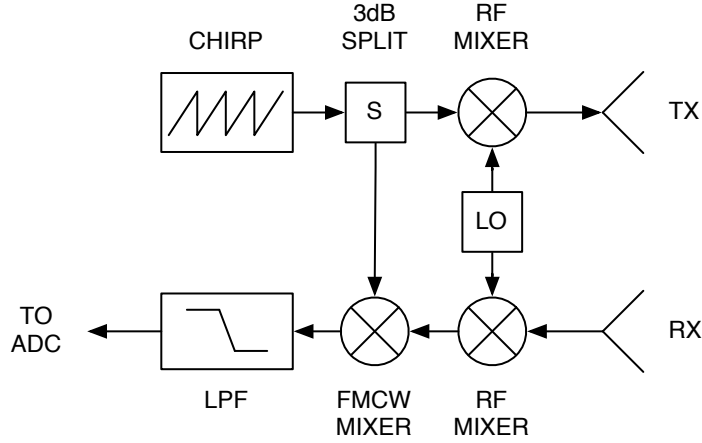


Figure 2.2: Block diagram of a generic FMCW transmitter and receiver

Nyquist-Shannon sampling theorem. Thus,

$$\Delta f_{max} = \frac{f_s}{2} = \frac{2r_{max}}{c} \cdot \frac{B}{\tau}, \quad (2.4)$$

which means

$$r_{max} = \frac{c\tau}{2B} \cdot \frac{f_s}{2}. \quad (2.5)$$

One chirp of duration τ sampled continuously at a rate of f_s results in $N_s = \tau f_s$ time samples. Taking the fast Fourier transform (FFT) and discarding the redundant frequency components, we are left with $N_s/2$ frequency bins ranging from DC to $f_s/2$. Each frequency bin corresponds to a range of

$$\Delta r = \frac{r_{max}}{N_s/2} = \frac{c}{2B}. \quad (2.6)$$

This is known as the range resolution of the radar. In fact, this equation is fundamental to all radar modes including pulsed, pulse-compressed, and FMCW. As long as the bandwidth is known, the range resolution can be determined.

In the case of TIMMi, the bandwidth is fixed at $B = 100$ MHz for a range resolution of $\Delta r = 1.5$ m. Adjustments to r_{max} are made by varying f_s or τ . A time series taken from the radar is seen in Figure 2.3. Low frequency components dominate the signal, and some clipping is present since the sampling threshold was set to $1 V_{p-p}$.

The primary advantages of FMCW are twofold. First, the sampling rate is much lower




Figure 2.3: A time series of a single linear chirp, downconverted, compressed, and sampled at baseband

than pulsed radars that achieve similar ranges because Δf_{max} is relatively low – on the order of 1 MHz. The compression of the signal is done in hardware, which means that slower, cheaper analog-to-digital converters (ADC's) may be used, and processing of the time-domain samples is as easy as taking an FFT in software. Second, FMCW radars offer compression gain over pulsed radars having the same bandwidth. The compression gain G_c is the time-bandwidth product of the chirp

$$G_c = B\tau \quad (2.7)$$

and is often on the order of tens of decibels. This enables the use of solid-state power amplifiers that are lightweight, reliable, and consume considerably less power than a high-peak-power amplifier like a Klystron.

It was mentioned that targets are assumed to be stationary. Targets moving in the range dimension inherently have a Doppler shift associated with them. We are entirely dependent on the frequency of the signal to determine its range, so a Doppler shift will cause the target to be misrepresented in range. A target moving towards the radar ($\frac{dr}{dt} < 0$) will have a higher frequency and will appear closer (smaller Δf) than its true range. Conversely, targets moving away from the radar ($\frac{dr}{dt} > 0$) have a lower-frequency echo (larger Δf) and will appear to have a farther range. Targets moving perpendicular to the beam have no Doppler velocity component, but they may suffer from blurring due to movement across the beam throughout the relatively long chirp duration.




Figure 2.4: Interferometric antenna configuration

2.2 Interferometry Basics


A traditional monostatic radar has a single antenna to both transmit and receive signals. An interferometric radar adds additional information by having two receive antennas physically separated by a baseline B_a (Figure 2.4). We are using the far-field approximation where R is very large, so we can assume that parallel plane waves are incident at both antennas. Echoes reach antennas A_1 and A_2 at different times, but this time delay is too small to measure directly. It is convenient instead to measure the phase difference between the received signals of the two antennas.

The free-space wave number


$$k = \frac{2\pi}{\lambda} \quad (2.8)$$

permits conversion between signal phase (radians) and physical wavelength (m)[11]. The phase difference between the two channels is $\phi = k\Delta R$. It is calculated by taking the complex correlation of the received waveforms,

$$\gamma e^{j\phi} = v_1 v_2^* = \gamma_1 e^{-j\phi_1} \gamma_2 e^{j\phi_2}. \quad (2.9)$$



(a)



(b)

Figure 2.5: Interferometric geometry over (a) a flat earth and (b) an existing DEM

In general, the extra path length ΔR is larger than one wavelength, and there are phase ambiguities at multiple wavelengths where the phase wraps from 0 to 2π .

The relationship between phase and target height is known as the phase sensitivity,

$$k_z = \frac{2\pi B_a \cos(\theta - \alpha)}{\lambda r \sin \theta}, \quad (2.10)$$

for a target at range r and look angle θ . The baseline B_a is tilted at angle α relative to the ground. k_z is an important factor that permits the generation of height maps based on interferometric information.

Height must be measured in relation to some reference, so we create a fictitious flat earth – a plane of constant height at some altitude beneath the platform (Figure 2.5a). We can then create a simulated interferogram for this reference plane that has an interferometric phase of ϕ_f . By calculating

$$e^{j\phi_i} = e^{j\phi} e^{-j\phi_f}, \quad (2.11)$$

we can then unwrap the differential phase ϕ_i in two dimensions to get a height measurement directly using

$$z = \frac{\phi_i}{k_z}. \quad (2.12)$$

Another option is to use an existing digital elevation model (DEM) in place of the flat earth (Figure 2.5b). In this case, z is calculated in reference to the DEM and is thus relatively small. If k_z is small enough such that $|k_z z| < \pi$, then there are no phase ambiguities present, and 2-D phase unwrapping becomes unnecessary. A detailed DEM may be generated by simply adding the heights of the base DEM and the differentially-measured heights.

2.3 Platform Geometry

Figure 2.6 defines the common geometric language used throughout this thesis. Beginning with the along-track view, clockwise aircraft roll is indicated by positive Ω and is measured parallel to the ground. We have already defined baseline tilt as α , and it is related to Ω by

$$\alpha = \Omega + 90^\circ + \alpha_B \quad (2.13)$$

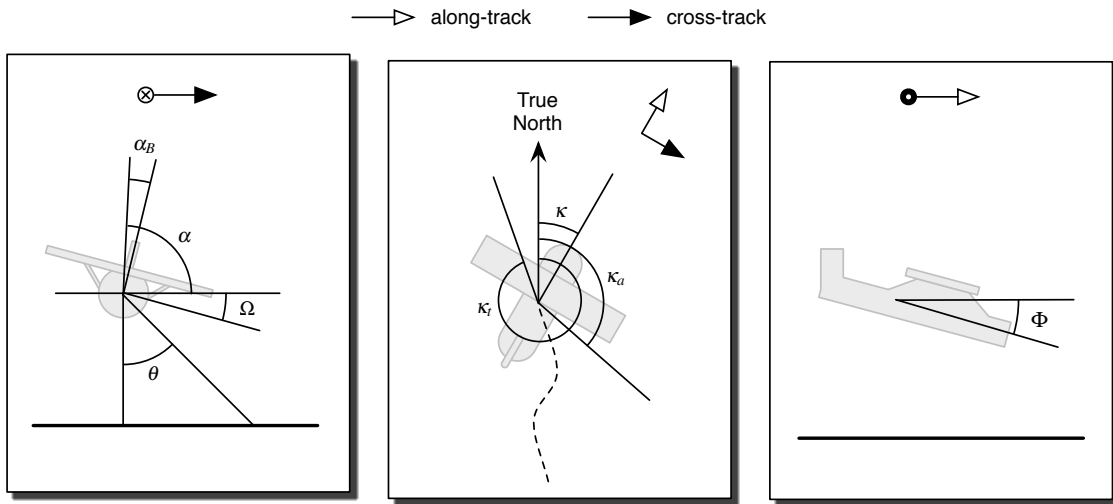


Figure 2.6: Airborne platform geometry nomenclature

where α_B is the baseline tilt off of vertical. If the baseline were perfectly vertical in horizontal flight, then $\alpha_B = 0^\circ$. Otherwise, α_B is a correction factor that can be determined by measuring the baseline when the aircraft is in the hangar. It may also be inferred from interferometric images. The look angle to a target on the ground is given by θ , which varies continuously across the swath.

From above, κ is the compass direction where the aircraft's nose is pointing. The antenna pointing angle κ_a is also an absolute measurement, but it is related to κ by an additive factor. Ideally, the antennas point directly broadside to the aircraft, and $\kappa_a = \kappa + 90^\circ$, but in reality this is not the case. Just like for α_B , the appropriate correction factor can be measured on the tarmac or inferred from the data. In this example, the aircraft's track along the ground is κ_t . This number is not very valuable, since we only care where the antennas are pointing at any moment in time, and not where the aircraft is heading. The important thing to note is that, in general, $\kappa_t \neq \kappa$. This effect is called yaw, and must be measured and accounted for.

The final panel shows a side view of the aircraft. The nose pitching down is represented by positive Φ , and is known as the pitch angle.

CHAPTER 3

TIMMI SYSTEM OVERVIEW AND MODIFICATIONS

A system-level block diagram is given in Figure 3.1. The signal source is a Tektronix AFG3252 arbitrary function generator that supports a 240 MHz bandwidth. The baseband (BB) signal is a linear FMCW chirp from 5 to 105 MHz. At the beginning of each chirp, a TTL rising-edge trigger is sent to the acquisition computer to specify the start of a new record. The baseband signal travels to the transceiver (Figure 3.2), where the dual upconverter (DUC) generates the Ka-band signal via a set of internal oscillators at IF (1.3 GHz) and RF (33.7 GHz). These oscillators are phase-locked to a standard 10 MHz reference signal which is provided externally. The RF signal is amplified and sent to the transmit antenna. The transmit and receive antennas are a matching set of three slotted waveguide antennas.

Upon signal reception, the dual downconverter (DDC) in the transceiver brings the RF signal back to baseband. For FMCW signal compression to occur, the baseband signal is then mixed with a copy of the transmit signal. The difference frequency, or beat frequency, is filtered, amplified, and passed on to the acquisition computer for each of the receive channels. The FMCW receive chain can be seen at the top of Figure 3.2.

The digitizer is a National Instruments PXI-5152 that records 8-bit samples on two channels at rates up to 2 GS/s. The acquisition computer is a National Instruments PXIe-8130 controller in a NI PXIe-1062Q chassis. The chassis contains the controller, digitizer, a NI 8262 Redundant Array of Independent Disks (RAID) controller, and a NI PXI-6259 data acquisition card for low-frequency analog and digital signals. The former two cards are not used by the airborne TIMMi. To maintain portability, the computer is controlled remotely over an ethernet connection using Windows Remote Desktop.

The GPS antenna is a standard, circularly-polarized “puck”. The GPS receiver sends NMEA-formatted strings over a serial connection to the acquisition computer. In addition to

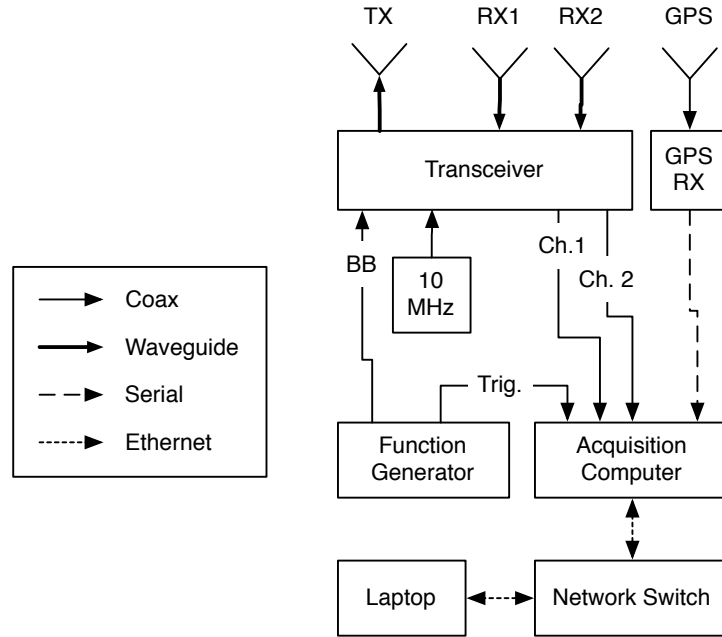


Figure 3.1: System overview block diagram. Connection types between each component are given in the legend.

the standard position updates, the receiver is configured by the acquisition software to send additional strings such as platform track direction and dilution-of-precision (DOP) figures.

3.1 Redesigning the Supporting Electronics

The first working prototype for TIMMi was initially designed for ground deployments, where weight and size were not constrained. A 15U rack-mount aluminum case housed a display, keyboard, power supply, computer, function generator, and four-disk RAID.

In the rear of the Cessna 206, floor space is at a premium, so the supporting electronics were rearranged into a vertical configuration. A 26 × 32 in. plywood board was outfitted with tie-down straps to attach the computer, function generator, and a network switch. The board is placed on its end and tied down in the rear of the aircraft cargo area.

The power supply and GPS receiver were housed in a 2U enclosure for protection (Figure 3.4). The large, heavy, linear power supply was replaced with a Vicor “Flat Pack” switching power supply, which has two independent 15 V outputs, each rated at ___ A. One downside of

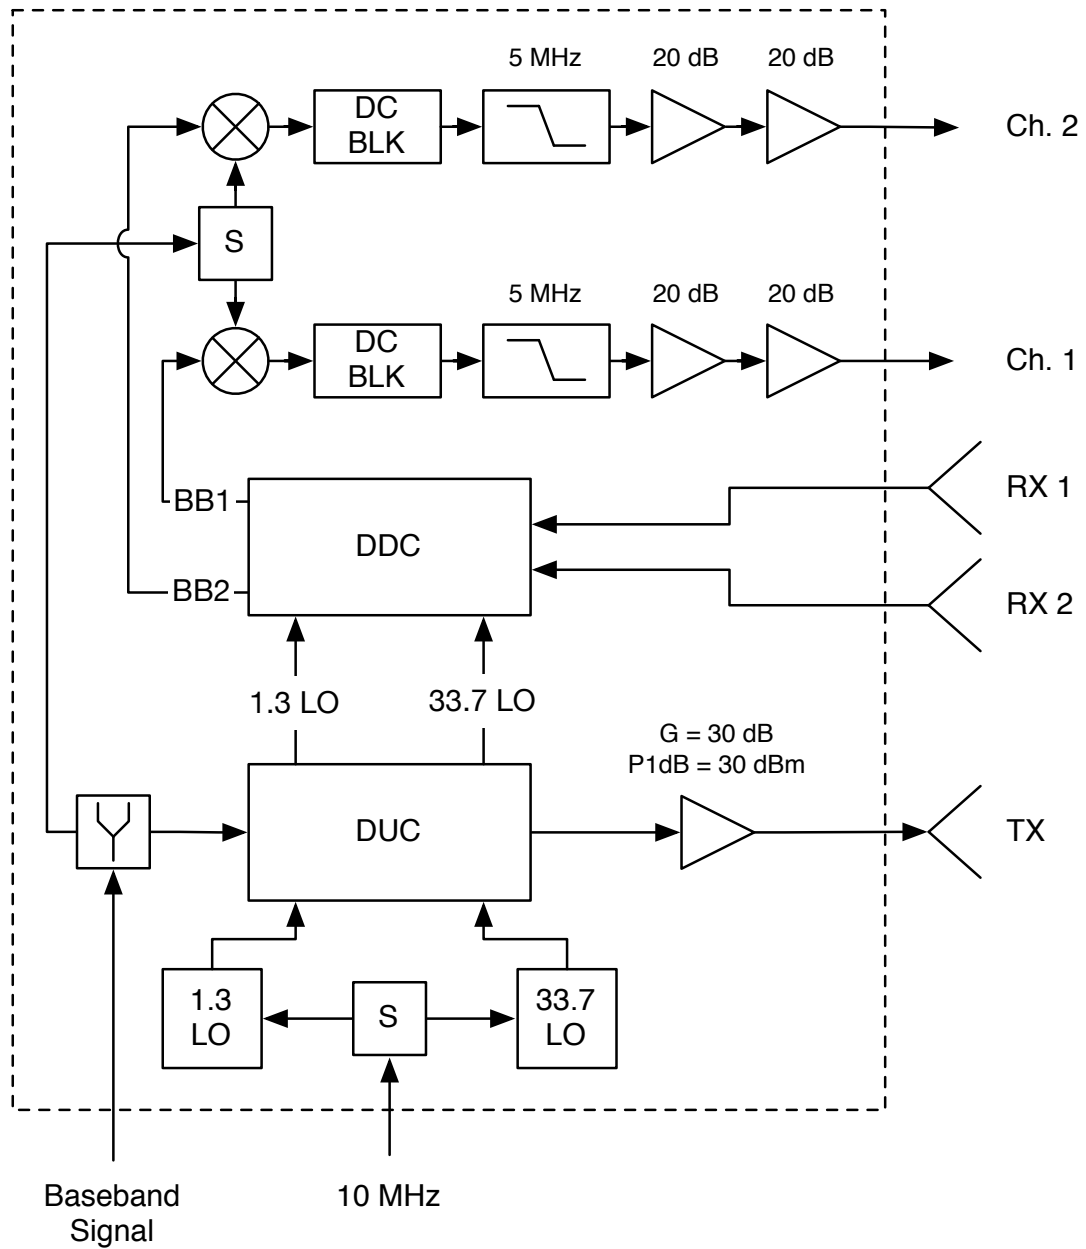


Figure 3.2: TIMMi transceiver block diagram showing DUC, DDC, oscillators, antennas, and FMCW downconversion chains.



Figure 3.3: The new “flat” supporting architecture

switching power supplies is that they can introduce RF interference on the DC outputs, providing noisy power that may introduce undesired signals into the system. To reduce these effects, each output of the Vicor is wired to a ripple attenuator module (RAM) to clean up the power supply voltage and to also provide isolation between the two outputs. The enclosure also includes the 9 V regulator for the power amplifier, as well as the GPS receiver and a USB-to-serial adapter. There is also a 10 MHz oscillator in the enclosure which provides the local oscillators in the transceiver with a standardized reference signal.

There is no display connected directly to the computer. Instead, a laptop and the computer are connected to a network switch, and then the laptop is used to control the acquisition computer using Windows Remote Desktop control software.

The acquisition data rate is about 4 MB/s (§4.3.3), which is sufficiently slow to record to a single disk in real time, so the RAID was removed completely. The 80 GB hard disk drive in the computer was upgraded to a 256 GB solid-state drive, providing increased storage, read/write speed, and reliability.


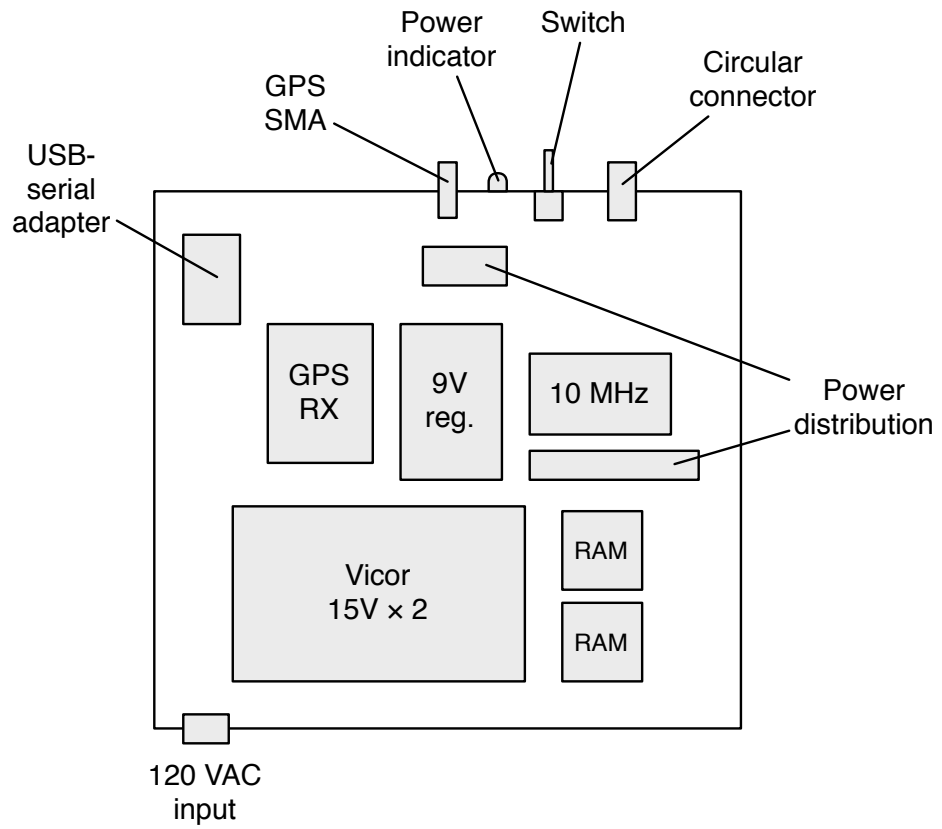


Figure 3.4: Layout of the new 2U power supply enclosure

3.1.1 Power Amplifier Power Supply

All of the components in the transceiver require 15 V DC except for the final stage power amplifier. The Millitech AMP-28-01090 operates from 8–13 V (8 V nominal) and has a maximum current draw of 1.5 A. Until now, an external lab bench power supply was used to power the amplifier.

The L78S09 is a 9 V regulator with up to 2 A of output current. The device experiences a voltage drop of nearly 1 V during high current draw, so a 9 V regulator was chosen to keep the supply voltage within the operating range of the amplifier. The regulator, while small, requires a large heat sink to dissipate the heat and prevent damage to the device. A thermal shut-down built into the device itself prevents damage if the temperature rises too high.


The regulator itself is a standard TO-220 package that is quite small. Most of its footprint in Figure 3.4 is occupied by a large heat sink. Without the heat sink, the circuit thermal protection kicks in after only a few seconds of operation, causing a failure. However, with sufficient heat dissipation, the L78S09 performed well under continuous heavy load, and did not show any sign of degradation or failure in any testing or deployment.

3.1.2 Power Connector Pinout

To accommodate the 9 V regulator and the separate 15 V supplies, the pinout of the circular military connector had to be modified. The connector and pinout are listed in Table 3.1. The pinouts for two other TIMMi connectors are also included for completeness.

There are now three separate rack-mount power supply units floating around the lab. One was the original supply for the Ka-band ground deployments. It also includes a separate supply and serial connections for the positioner, but does not include 9 V for the power amplifier. The second supply was developed for Ku-band ground deployments in an attempt to filter out spurious tones, but for the original supply is also needed alongside it to power the positioner. The new supply is designed exclusively for the Ka-band transceiver and is not compatible with the Ku-band system.

Overall system power measurements are listed in Table 3.2.



Pin	Ka Airborne		Ka Ground		Ku Ground	
	V	Use	V	Use	V	Use
A	+15	LOs, BB amps	+15	LOs, BB amps	+15	BB amps, power amp
B	GND	DUC & DDC				
C	+15*	DUC & DDC			+5	L-band amp, L-band LO
D						
E					+15	DDC
F			+15*	Unused	+12	Ku-band LO, fans
G						
H					GND	L-band amp, L-band LO
J	+9	Power amp	+9	Power amp		
K	GND	Power amp	GND	Power amp	GND	BB amps, power amp
L	GND	LOs, BB amps	GND	LOs, BB amps	GND	Ku-band LO, fans
M						
N						
P						
R						
S						
T						
U					GND	DDC

Table 3.1: Power supply 18-pin circular connector pinouts for three iterations of TIMMi power supplies. Voltages marked with an asterisk (*) are always on during standby, and blank entries are not connected (N/C). All ground (GND) terminals are connected.

Mode	Current		Power
	115 VAC	12 VDC	
Standby	2 A	20 A	230 W
Radiate	2.5 A	24 A	290 W

Table 3.2: TIMMi system power draw; during standby, all subsystems are on except the transceiver



Figure 3.5: The rear cargo doors of the Cessna 206 open like a clamshell with no center pillar

3.2 Modifying the Aircraft Door

An early challenge of the airborne TIMMi project was finding a way to physically mount the antennas to the aircraft. Permanent modifications to the Cessna 206 were out of the question. They would be very expensive, would take a long time to complete, and would be difficult or impossible to get approved by the FAA. Instead, we opted to make modifications to a more temporary part of the aircraft: the rear cargo door.

The Cessna 206 Stationair is designed for carrying large, heavy cargo. Its cargo doors (Figure 3.5) open outward in a clamshell configuration that does not have a center pillar, and each door can be taken off by simply removing two pins from its hinges.

A spare door for the 206 was purchased and adapted for the radar. Thomas Liimatainen, a machinist at Mount Holyoke College, made modifications to the existing antenna mount (used for ground deployments) so that it is easily attached to the door. The mount itself allows the antennas to be positioned vertically in increments of 2 cm, and can accommodate both the Ka-band and Ku-band slotted waveguide antennas. The boresight elevation can be adjusted to any



Figure 3.6: TIMMi Ka-band antennas mounted on a Cessna 206 cargo door. The top two slotted waveguide antennas receive, and the bottom antenna transmits.



Figure 3.7: The interior of the aircraft door showing mounted transceiver and flexible WR-28 waveguides

angle. The slotted waveguide antennas have a low profile that is naturally aerodynamic, so they do not negatively impact the aerodynamics of the aircraft.

Three holes were drilled through the door and fitted with bulkhead waveguide adapters. Flexible WR-28 waveguides route the transmit and receive channels through the door to the transceiver, which is mounted directly on the inside of the door (Figure 3.7). The waveguides were kept as short as possible to minimize conductive losses in the feedline that would hurt the overall receiver noise figure.

The fully-assembled door weighs about 50 pounds and is easily carried.

3.3 Rewriting the Acquisition Software

The software for acquiring data in the field was written in LabVIEW. The previous application controls a positioner atop a tripod that sweeps out a predetermined number of looks for each acquisition. The code was modified to support continuous acquisition and to display additional diagnostic information such as waveforms and backscatter power for each channel.

The software also reads NMEA sentences from the GPS receiver and stores the data inline with the binary RF data. The primary reason for recording the GPS data is to provide a global timestamp so that the RF data can be matched up with the motion compensation data provided by AHRS. There is an inherent delay between the one second GPS “ticks” and the beginning of the next look. The LabVIEW program keeps track of this time offset and adds it to the look timestamp when the next look begins.

3.3.1 Digitizer Re-arm Time

In order for the RF data at every chirp to be acquired at the correct time, a systemwide rising-edge trigger is employed at the beginning of each FMCW chirp. The trigger is supplied by the Tektronix AFG3252 arbitrary waveform generator, which also generates the baseband FM sweep. The trigger is read by the digitizer card and is used to begin a new acquisition.

The NI PXI-5152 digitizer card cannot continuously acquire triggered data. There is a hardware delay between the final acquired sample and the beginning of the next acquisition, called the re-arm time. The re-arm time for this digitizer is specified at $t_r = 8 \mu s$, which was confirmed via experimentation. This means that the trailing samples of each acquisition must be

truncated to prepare for the next trigger, which results in a slight undersampling of the received signal.

Recall that the range resolution of a radar is

$$\Delta r = \frac{c}{2B}. \quad (3.1)$$

By truncating the trailing samples, only a portion of the full chirp bandwidth is sampled, which degrades the range resolution. The simple linear relationship is given by

$$\Delta r_{trunc} = \frac{r_{max}}{N_s/2} = \Delta r \frac{N_s}{N_{s,trunc}}. \quad (3.2)$$

Here, r_{max} is the maximum range in the ideal case, which does not change, and N_s is the number of samples..

For example, with $f_s = 2$ MHz and $\tau = 1$ ms, each pulse would nominally consist of $N_s = \tau f_s = 2000$ samples and have $\Delta r = 1.5$ m. The trigger re-arm time takes the place of $t_r f_s = 16$ samples, which is rounded up to 20 samples as a buffer. The final sample count is $N_{s,trunc} = 1980$. After taking the FFT and removing the duplicate frequency components, we are left with 990 range bins covering the full range up to $r_{max} = 1.5$ km. Therefore, each range bin has a resolution of $\Delta r_{trunc} = 1.515$ m.

The compression gain is also degraded a small amount because the both B and τ are reduced. The truncated bandwidth is given by

$$B_{trunc} = B \left(1 - \frac{t_r}{\tau}\right) \quad (3.3)$$

and the truncated chirp duration is

$$\tau_{trunc} = \tau \left(1 - \frac{t_r}{\tau}\right), \quad (3.4)$$

so the effective compression gain is

$$G_{c,trunc} = B\tau \left(1 - \frac{t_r}{\tau}\right)^2. \quad (3.5)$$

As long as $t_r \ll \tau$, the impact on G_c is minimal. In this example, $t_r = 10 \mu s$ including a buffer in

the delay, which is just 1% of the full acquisition time. This results in a loss of only 0.09 dB over the ideal case.

3.3.2 RF Data File Format

The LabVIEW acquisition program writes raw RF samples to disk in a binary format. This file is not self-descriptive and has no metadata attached to it, so the file format specification must be followed very closely. There are currently two versions of this file format. Version 1 was only used on the February 3 flight, and version 2 was used for all subsequent flights. The first 4 bytes of the file indicate its version so that it can be properly loaded into memory. If any modifications are made to this file format, the LabVIEW program must be modified to write a new file version number in the file header, and the MATLAB processing code must be able to correctly read the new version. The byte order of the file is little-endian.

The format specifications are given in Tables 3.3 and 3.4. RF data is merged with GPS data and transferred from the digitizer to disk once every “look”. A look is simply a collection of consecutive pulses. Looks are usually configured to occur once every second so that no GPS data is missed. Every look has a header that contains some basic GPS information so that it can be merged back in with the AIMS data in post-processing. Then the look contains a sequence of samples, first from channel 0, then from channel 1. Knowledge of N_{samp} and V_{p-p} – obtained from the configuration text file – are necessary to read the file. A concrete example of the file format is given in Table 3.5.

The NI PXI-5152 ADC takes 8-bit (1 byte) time samples, so each channel’s pulse is N_{samp} bytes long. Each sample is a scaled measurement based on the peak-to-peak voltage V_{p-p} specified at the beginning of the acquisition, where $-V_{p-p}/2$ corresponds to signed byte value -128 , and $+V_{p-p}/2$ is equal to signed byte value $+127$. Thus, the scaling for a sample byte b to a true voltage V is

$$V = \frac{b + 0.5}{2^7 - 0.5} \cdot \frac{V_{p-p}}{2}. \quad (3.6)$$

Previous versions of the acquisition software written for the Ku-band interferometer wrote 16-bit values to disk despite only sampling 8 bits, resulting in unused bits and an unnecessary doubling of file size.

The Format column is given in terms of the MATLAB representation of that data type.

Bytes	Format	Description	Notes
File header			
4	uint32	File version	Equal to 1
Look			
23	char	Timestamp	ASCII, UTC time, no null terminator
8	double	Latitude	Zero if no GPS fix
8	double	Longitude	Zero if no GPS fix
Pulse			
n_{samp}	int8	Ch. 0 samples	
n_{samp}	int8	Ch. 1 samples	

Table 3.3: TIMMi RF file format, version 1

Bytes	Format	Description	Notes
File header			
4	uint32	File version	Equal to 2
Look			
23	char	Timestamp	ASCII, UTC time, no null terminator
4	uint4	Length of UTC string	Equal to 9
9	char	UTC time string	HHMMSS.ss
8	double	Latitude	Zero if no GPS fix
8	double	Longitude	Zero if no GPS fix
1	uint8	Position type	0 = no fix, 1 = fix, 2 = diff.-corr. fix
1	uint8	Num. of satellites	
8	double	HDOP	Horizontal dilution of precision
8	double	Meters above MSL	WGS-84 datum
8	double	Meters above ellipsoid	WGS-84 datum
8	double	Age of differential correction	Zero, since we are not using diff. corr.
Pulse			
n_{samp}	int8	Ch. 0 samples	
n_{samp}	int8	Ch. 1 samples	

Table 3.4: TIMMi RF file format, version 2

Bytes		Description
Start	Length	
1	4	File format version
5	23	Look 1 UTC timestamp string
28	4	Length of following UTC string (9)
32	9	UTC time string (9 characters)
41	8	Latitude
49	8	Longitude
57	1	Position fix type
58	1	Number of satellites
59	8	HDOP
67	8	Meters above MSL
75	8	Meters above ellipsoid
83	8	Age of differential correction
91	N_{samp}	Look 1, pulse 1, channel 0 samples
$91 + N_{samp}$	N_{samp}	Look 1, pulse 1, channel 1 samples
$91 + 2N_{samp}$	N_{samp}	Look 1, pulse 2, channel 0 samples
$91 + 3N_{samp}$	N_{samp}	Look 1, pulse 2, channel 1 samples
... repeat for N_{pulse} pulses...		
$91 + 2N_{pulse}N_{samp}$	23	Look 2 UTC timestamp string
$91 + 2N_{pulse}N_{samp} + 23$	4	Length of following UTC string (9)
... repeat for an indeterminate number of looks...		

Table 3.5: An example layout of a version 2 RF data file with N_{samp} samples per pulse and N_{pulse} pulses per look




Figure 3.8: Residual delay in recording data due to the operating system, where zero is the nominal time

The format for the Timestamp field is a 23-character string of the format `'YYYY-mm-dd HH:MM:SS.uuu'`. The fractional second part, `uuu`, is added on to the GPS timestamp that comes from TIMMi (and not the AIMS system). It is the differential time between when the NMEA sentence arrives over the serial port and when the RF data arrives from the digitizer. We are not using a real-time operating system, so sometimes there is an indeterminate delay in transferring the RF samples from the digitizer into system memory (Figure 3.8). By keeping track of the “extra” time that elapses between GPS acquisition and RF acquisition, we can more accurately characterize the platform’s position.

3.4 Integration with AIMS-1

The Airborne Imaging Multispectral Sensor (AIMS) system developed at Mount Holyoke is used for forest research[9], and is currently deployed on the Cessna 206 alongside TIMMi. Its Altitude and Heading Reference System (AHRS) is an inertial measurement unit (IMU) that records aircraft pitch, yaw, roll at a rate of 71 Hz. The yaw angle is measured by using a differential GPS receiver that has two antennas separated by a baseline.

AIMS-1 also features a LIDAR rangefinder to measure aircraft height above ground level at approximately 240 Hz. However, the LIDAR pulse power is weak at the altitudes we want to fly for TIMMi (larger than 1000 feet), so GPS altitude data is used instead.

The TIMMi system itself also features a GPS receiver which is used to timestamp every RF pulse. Post-flight, AHRS data are interpolated into the RF time basis for processing.

The AIMS instruments produce three separate plain-text data files, documented in the following sections.

3.4.1 LIDAR Data File (.las)

The LIDAR rangefinder data is parsed using the `read_lidar` function. The LIDAR generates a large text file with a `.las` extension. Each line in the ASCII text file is one LIDAR measurement, 39 bytes long plus carriage return and line feed characters (ASCII 13 and 10) for 41 bytes total per line. The format for each line is shown in Table 3.6.

The Windows timestamp, used throughout the AIMS data files, is specified as the number of 100 ns intervals since January 1, 1601. The function `win_timestamp_to_datetime` will

Description	Bytes	Example
Height (ft)	7	02597.3
First or last return (f or F)	1	F
Separator	1	,
Strength (units?)	4	0276
Separator	6	__UTC:
Windows timestamp	20	00129787911606979800
New line	2	\n

Table 3.6: AIMS-1 LIDAR data file format

```

$GPGGA,130115.00,4224.353127,N,07228.103662,W,1,09,0.9,919.24,M,-33.74,M,,*5E
$GPVTG,200.8,T,,,141.92,N,262.84,K,A*4F
$GPGSA,A,3,09,05,27,08,15,21,26,29,18,,,,,1.8,0.9,1.5*36
$GPRMC,130115,V,4224.353127,N,07228.103662,W,141.92,200.8,130412,14.7,W,N*08
_TIME:000000129787956750000000
$GPGGA,130116.00,4224.316464,N,07228.123224,W,1,09,0.9,914.50,M,-33.74,M,,*54
$GPVTG,202.4,T,,,142.49,N,263.90,K,A*40
$GPGSA,A,3,09,05,27,08,15,21,26,29,18,,,,,1.8,0.9,1.5*36
$GPRMC,130116,V,4224.316464,N,07228.123224,W,142.49,202.4,130412,14.7,W,N*07
_TIME:000000129787956760000000

```

Figure 3.9: An example of two sets of NMEA 0183 strings with GPS data

convert this number into a MATLAB “datenum” representation. This is used in the `read_lidar` function.

If the size of the data file is not an exact multiple of 41 bytes, `read_lidar` will throw an exception. Sometimes the AIMS logging software will truncate a line if it is closed in the middle of writing that line to disk. In this case, the user must manually edit the `.las` file and remove the truncated line at the end of the file, ensuring that the format of newline characters (`\n`) is maintained, and that the last line of the file is blank.

3.4.2 GPS Data File (.gps)

The GPS data file has a `.gps` extension, and is parsed using the `read_nmea` function. This text file is simply a concatenated list of all the NMEA-formatted strings collected from the serial port of the GPS receiver. The National Marine Electronics Association protocol NMEA 0183 is the standard method to capture information from a GPS receiver. Each sentence consists of a sentence identifier, a series of comma-delimited parameters, and a checksum. Figure 3.9 has an example listing from an actual data file.

Description	Bytes	“Old” col.	“New” col.	Example
Data quality	1	1	1	I
Time since turn-on (sec)	5	2	N/A	00288
Clockwise roll (Ω , deg)	5	3	2	+02.7
Nose pitch down (Φ , deg)	6	4	3	-003.0
Azimuth (κ , deg)	5	5	4	053.8
Roll rate (units?)	5	6	5	+0.05
Pitch rate (units?)	5	7	6	+0.05
Azimuth rate (units?)	5	8	7	+1.00
Velocity (units?)	6	9	8	+178.7

Table 3.7: AIMS-1 IMU data file format

There are four types of sentences in the data file: `$GPGGA`, `$GPVTG`, `$GPGSA`, and `$GPRMC`. NMEA 0183 is widely adopted and well documented, so refer to one of the many references¹ for descriptions of these sentences. The log file also contains lines beginning with `_TIME:`, which are not NMEA sentences, but are Windows timestamps interleaved with the NMEA data.

Only some of the NMEA sentences contain UTC timestamps, so we must choose a standard timestamp to apply to any ancillary sentences. Every time `read_nmea` encounters a `$GPGGA` string, it uses that timestamp and applies any sentences following it to the same instant in time, until a new `$GPGGA` string is encountered. In reality, all the data contained in the following sentences was captured and measured at the same time, but it is just transmitted serially in separate sentences, so this is a safe assumption to make.

The NMEA checksums are not computed or checked, and the Windows timestamps are ignored. The UTC timestamp strings are converted into the MATLAB datenum representation. As before, sentences at the end of the file may be truncated, so manual editing may be required to remove incomplete sentences at the end of the file.

There is no date information contained in the NMEA sentences, so when calling `read_nmea`, you must also specify a string to indicate on what date the data file was taken, e.g. `'July 11, 2012'`. This does not handle the edge case where the UTC date changes in the middle of an acquisition.

3.4.3 IMU Data File (.wbd)

The third instrument in the AIMS package is the inertial measurement unit (IMU) manufactured by Watson Industries. It measures acceleration in the pitch and roll axes and measures heading using magnetic sensors.

The Watson IMU generates a `.wbd` file that can be read using the `read_imu` function. There are two versions of the IMU device, an “old” one and a “new” one. The ASCII text data file has a number of space-delimited fields (Table 3.7), followed by the Windows UTC timestamp, followed by a `\n` newline character.

The file format for each IMU version is similar, but there are a few slight differences. The old IMU has an extra column to specify the number of seconds since system turn-on. The first column in both cases is a boolean value indicating the quality of data received from the LIDAR for that particular pulse. For the old IMU, the lower-case “i” indicates low-quality data, and an upper-case “I” indicates high-quality data. The new IMU uses the letters “g” and “G”, respectively.

The roll measurement, Ω , is taken with a positive value indicating a clockwise roll, and a positive pitch Φ corresponds to the nose pitching down. Azimuth is an absolute compass direction. It was found that the June 11 azimuth measurement was the complement angle measurement to what was expected, perhaps because the differential GPS antennas were accidentally switched. For this flight, the azimuth angle is calculated as $360^\circ - \kappa$.

¹<http://aprs.gids.nl/nmea/>

CHAPTER 4

CHALLENGES

4.1 DUC/DDC Voltage Regulator Interference

An FMCW received baseband signal is at relatively low frequencies, beginning at DC for targets at the first range bin. The Nyquist sampling theorem states that the signal must be sampled at least twice its maximum frequency in order to properly reconstruct the original signal without any aliasing. This means that the sampling frequency can also be kept low. For example, a common configuration for TIMMi is to use $f_s = 2$ MHz, so the maximum beat frequency is $f_{max} = 1$ MHz and the chirp length is $\tau = 1$ ms. This yields a maximum range of $r_{max} = \frac{c\tau}{2B} f_{max} = 1.5$ km.

One problem with this approach is that many sources of signal interference occur within the passband of the receiver. When these signals leak into the receive chain, they often cannot be simply filtered out, because they occur where a signal is expected.

The DUC and DDC printed circuit boards both contain voltage regulators that bring the 15 V supply down to lower voltage levels for integrated circuits on the board. These are switching regulators that operate by turning the supply voltage on and off very quickly. An oscillator internal to the voltage regulator has a specific duty cycle that enables the voltage regulation. The datasheet for the LM2676 device specifies an internal oscillator of 260 kHz. Experimentally, it was determined that the DUC regulator has a principal tone of 269 kHz, with weaker images at 539 kHz and 809 kHz. The DDC regulator showed up at 266 kHz, with an image at 530 kHz. All of these signals appeared in both the baseband RF channels, as well as on the 15 V DC power supply lines (Figure 4.2a).

It was suspected that the undesired regulator tones were leaking through the DC power supply, which also supply power to the final stage baseband amplifier blocks right before the

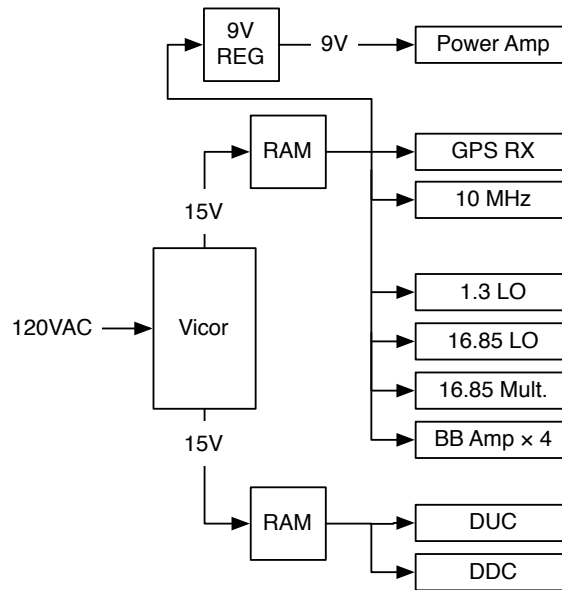



Figure 4.1: Power distribution block diagram

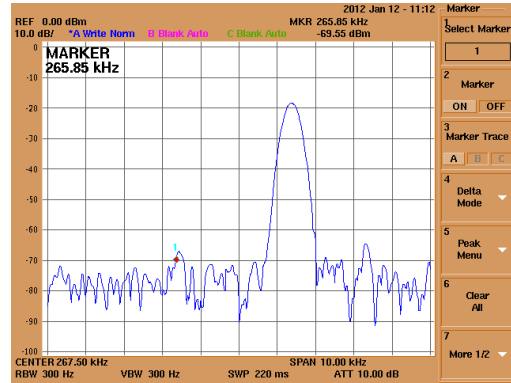
National Instruments digitizer. This was confirmed by powering the DUC and DDC from a bench-top power supply, effectively isolating the two power channels and removing all traces of the tones from the baseband signal.

A common way to prevent undesired radio frequency interferences (RFI) is to install an RF choke in which the wire is wrapped around a ferrite toroid. The ferrite has an inductive effect that blocks RF signals, effectively behaving like a low-pass filter. The part LFB143064-000 from Laird Technologies was selected for its high impedance at low frequencies, although impedances were only specified at the lowest frequency of 500 kHz, suggesting that this material might not provide enough attenuation at 270 kHz. Using a spectrum analyzer, it was shown that a single turn of the wire through the toroid yielded an attenuation of 33 dB, and three turns achieved 43 dB of attenuation, with no appreciable attenuation for additional turns after that (Figure 4.2).

When first implemented, this seemed like the ideal solution until some odd measurements started showing up. The choke on the transmitter line was emitting a high-pitched audible tone when powered on. This was an indication that the choke itself was actually oscillating at an audible frequency. The 5.8 kHz tone showed up in many places. It even mixed with the Ka-band LO at 33.7 GHz and showed up as sidebands in the LO output signal. The impact of this ringing turned out to be far more troublesome than the original problem, so the RF chokes



(a)



(b)

Figure 4.2: Undesired voltage regulator tones from DUC and DDC in a baseband receive channel with (a) no RF chokes installed, and (b) three turns of the DC wire around the DUC choke

were removed.

The second attempt to isolate the tones from the other circuitry involved integrating a Vicor power supply with two independent 15 V outputs. Vicor also manufactures Ripple Attenuator Modules (RAM), which combine passive and active filtering elements to smooth out the power supply, since the Vicor itself is a switched supply. VE-RAM-E1 ripple attenuator modules were installed on both outputs in an attempt to reduce the coupling between the two 15 V sources. This same configuration was used by Anthony Swochak for the Ku-band interferometer [14]. This configuration attenuated – but did not completely remove – the interference.

4.2 Near-Field Antenna Coupling

Another major source of interference in the baseband receive signal is the presence of strong, correlated, low-frequency signals that dominate from DC up to 150 kHz. These signals only show up in field deployments of the radar, which suggests that they are the result of interaction between the transmit and receive antennas. The exact nature of this coupling has not yet been explored. At the time of this writing, it is suspected that it is due to the physical spacing of the antennas acting like a filter that is oscillating or ringing.

The pattern of directional antennas is usually made up of one main lobe and many sidelobes. They are typically designed such that the sidelobe levels are low enough (in relation to the main beam) to be disregarded. However, a weak signal leaking from the transmit antenna

into the receive antenna's sidelobes will be many times stronger than a desired signal received from the main beam simply due to the proximity of the antennas. This is a problem that can be solved by physical isolation of the transmit antenna from the receive antennas.

It should be noted that this is a problem that is exclusive to FMCW radar, which is constantly transmitting and receiving. For a pulsed or pulse-compressed radar, the transmitter is off during the receive cycle, so this coupling does not have an effect.

For most of the flights, the transmit antenna was placed as far as possible from the closest receive antenna for the given antenna mount, 20 cm. For the April 13 flight, the antennas were set up in an alternate configuration, where the transmit antenna was placed in between the two receive antennas. This was done to increase the baseline and thus the rate of interferometric fringes. However, this also reduced the distance between the transmit and receive antennas to 10 cm, and therefore increased coupling. The consequence was a decreased dynamic range in both receive channels, and poor SNR that did not have enough phase information in the correlation between the two channels.

This ringing dominates low frequencies in the range bins between the airplane and the ground, so it is easy to filter them out in post-processing. However, the large amplitude of these signals is what causes the biggest problem. The digitizer has a limited dynamic range that is specified at the time of acquisition in terms of peak-to-peak voltage, V_{pp} . It is also limited to 8 bits per sample. As V_{pp} is increased, quantization error for small signals increases, reducing the ability of the radar to measure weak frequency components. Eliminating the ringing would decrease the dynamic range of the received signals, essentially compressing the magnitude response and bringing out weaker echoes.

4.3 Determining Deployment Configurations

4.3.1 Flight Lines and Times

Finding suitable flight lines began as a trial-and-error process and evolved into a more methodical approach as confidence in TIMMi's performance improved. To maximize the performance of the overall instrument, a tradeoff between SNR and swath width is made.




Figure 4.3: Antenna cross-track look geometry

Among the adjustable factors that influence these quantities, altitude is the most influential. A higher flight line yields a wider cross-track swath along the ground, but the SNR decreases because the received power degrades at a rate of r^4 . We have seen so far that a 3000 feet AGL (above ground level) is too high, but 2000 feet AGL (approximately 600 m) is a reasonable compromise, as the signal approaches extinction at the far end of the beam. 1000 feet AGL also proved successful, although the cross-track range is significantly shorter.

The time of day also plays a role in the overall quality of the data. The first flight was conducted in the afternoon, when daytime heating created pockets of turbulence and made it difficult to maintain a steady course and altitude. Subsequent flights were flown early in the morning, around 7:30 AM, when winds were much calmer. Interpolation of platform orientation data is more accurate when the deviation of the aircraft motion from a straight flight path is small. In particular, a number of SAR processing equations assume the platform follows a straight course, which is easier to achieve in calm air.

4.3.2 Flight Altitude and Antenna Configuration

The antenna boresight angle, θ_L , is one of the few system parameters that cannot be adjusted during flight. Adjusting the boresight angle simply controls where power from the antennas is delivered. If the boresight angle is too high, the power is delivered at grazing incidence angles, and most of it is reflected away from the receiver. If the boresight angle is too low, there is no return power at farther ranges where the gain is needed most. Additionally, because of the vertical alignment of the baseline, there is no interferometric information directly below the

θ (deg)	h (m)	x_{min} (m)	x_{max} (m)	r_{min} (m)	r_{max} (m)	Notes
45	300	125	725	325	785	Good; LIDAR coverage
45	600	250	1450	650	1570	Recommended
45	900	370	2170	975	2350	Okay
60	300	230	2300	380	2300	Good; LIDAR coverage
60	600	460	4560	760	4600	Alt. and angle too high
60	900	690	6800	1135	6900	Alt. and angle too high

Table 4.1: Example configurations for the look geometry of Figure 4.3 for $\theta_B = 45^\circ$; distances are approximated for convenience; the recommended configuration is in bold

aircraft, and k_z is very large for small θ , so it is very difficult to unwrap phase ambiguities close to nadir.

The cross-track look geometry is shown in Figure 4.3, and a number of typical configurations are listed in Table 4.1. The numbers in this table are approximations designed to inform a range of acceptable values for the maximum unambiguous range of the radar. For flight planning, it was necessary to keep in mind that h is the altitude above ground level (AGL) which varies as the ground rises and falls beneath the aircraft. The aircraft altitude also varies by 50 m or more despite the best efforts of the pilot to maintain a constant altitude. The values from the table are derived from simple trigonometric analysis of Figure 4.3, and are given by

$$x_{min,max} = h \tan \left(\theta \mp \frac{\theta_B}{2} \right) \quad (4.1)$$

and

$$r_{min,max} = \frac{h}{\cos(\theta \mp \theta_B/2)}. \quad (4.2)$$

In order to provide a buffer so that measurements are not missed, it is important select a PRF and sampling rate such that r_{max} is greater than the maximum desired value. The variables x_{min} and x_{max} are useful for pre-flight mapping in determining the swath along the ground that will be imaged throughout a flight line.

The slotted waveguide antennas have a 3 dB elevation beamwidth of 45° [13], so reasonable boresight angles are in the range $22.5^\circ < \theta_L < 67.5^\circ$. Two flights each have been completed for $\theta_L = 45^\circ$ and 60° , revealing that the 60° boresight angle is too high to achieve sufficient SNR for good interferometric data. $\theta_L = 45^\circ$ appears to be a good compromise.

The baseline separation, B_a , determines the rate of phase wrapping in interferometric fringes. A larger baseline is more sensitive to small height variations, which makes it difficult to unwrap the phase ambiguities in areas with steep or abrupt changes in height. The maximum unambiguous height change is given by $\Delta h_u = 2\pi/k_z$. The baseline must be chosen such that Δh_u is much larger than the maximum expected difference in height between two adjacent pixels. For example, flying at a height of 600 m and looking at a pixel in the center of the beam at $\theta = 45^\circ$, setting the baseline to $B_a = 8$ cm results in $\Delta h_u = 91$ m, which is more than sufficient for the rolling hills, treelines, and buildings in the acquisition area.

4.3.3 Baseband Signal Acquisition

The duration of the baseband chirp, τ , impacts the performance of the radar in several ways. In general, τ should be as long as possible to maximize compression gain, $G_c = B\tau$. However, just like every other radar parameter, a trade-off must be taken into account. The maximum range is given by

$$r_{max} = \frac{c\tau}{2B} f_{max}. \quad (4.3)$$

Due to the Nyquist criteria, $f_{max} = f_s/2$, so in order to increase r_{max} , either f_s or τ can be increased. The bandwidth, B , is fixed at 100 MHz to take advantage of the full bandwidth of the system. If τ is too long, platform motion throughout the duration of the chirp begins to become a concern. In an FMCW radar, the beat frequency response of a target is accumulated during the sweep. However, the target appears to move across the beam throughout τ , so its beat frequency coherently adds although the response is changing, resulting in distortion of the target response. If SAR processing is to be considered, this also imposes an upper bound on τ (§7.2.2).

As τ is decreased, f_s must be increased in order to maintain a constant maximum range. The drawback to increasing f_s is an increased data rate. However, the data rate for FMCW is so low that there is a lot of flexibility in increasing f_s . A typical deployment has $f_s = 2$ MHz for each 8-bit receive channel – a total data rate of 4 MB/s. Write speeds of up to 40 MB/s have been demonstrated on the data acquisition computer.

Function	
Sine	
Run Mode	
Sweep	
Amplitude Menu	
Amplitude	4 Vpp
Offset	0 V
Sweep Menu	
Page 1	
Start Frequency	5 MHz
Stop Frequency	105 MHz
Sweep Time	1 ms
Return Time	0 ms
Page 2	
Center Frequency	55 MHz
Span	100 MHz
Hold Time	0 ms
Type	Linear
Page 3	
Mode	Repeat
Source	Internal
Slope	Positive
Trigger Interval	1 ms

Table 4.2: Tektronix AFG3252 configuration settings for a 1 ms chirp (1 kHz PRF)

4.3.4 Function Generator Configuration

The Tektronix AFG3252 is a full-featured arbitrary function generator that has a lot of flexibility and, thus, a lot of settings. In order to aid future users of the system, Table 4.2 specifies the desired settings for a typical deployment. The configuration may be stored in one of four memory banks by using the *Save/Menu* button.

CHAPTER 5

DATA PROCESSING TECHNIQUES

5.1 Data Processing Flow

Post-flight, the raw data is processed using a series of MATLAB functions and toolboxes (Figure 5.1). The inputs consists of the AHRS time-domain platform motion information, a DEM provided by MassGIS, and the TIMMi time-domain RF samples. Every radar pulse is assigned an interpolated timestamp from the TIMMi GPS, and this time domain is used as the destination domain for interpolating all the other data sets.

The interpolation process is shown in more detail in Figure 5.2. First the raw time sample data from the digitizer (Figure 5.2a) is converted to the frequency domain via the Fast Fourier Transform (FFT) (Figure 5.2c). These frequency samples correspond directly to range. Data in this coordinate system is termed as being in “radar coordinates.” In the map domain, we begin with the DEM (Figure 5.2b) and interpolate the heights to an arbitrary grid in map coordinates which will be used for the final image (Figure 5.2d). We then iterate over every radar pulse, taking platform position and antenna beam pattern into account to project the range bins onto ground coordinates.

During this final step, we use real aperture radar (RAR) processing and simply average the responses from overlapping pulses. Future work will incorporate synthetic aperture radar (SAR) processing techniques (§7.2.2).

5.2 Workflow

This section will walk the reader through the entire data processing workflow, beginning with raw data and ending up with geocoded reflectivity and interferometric images. Each

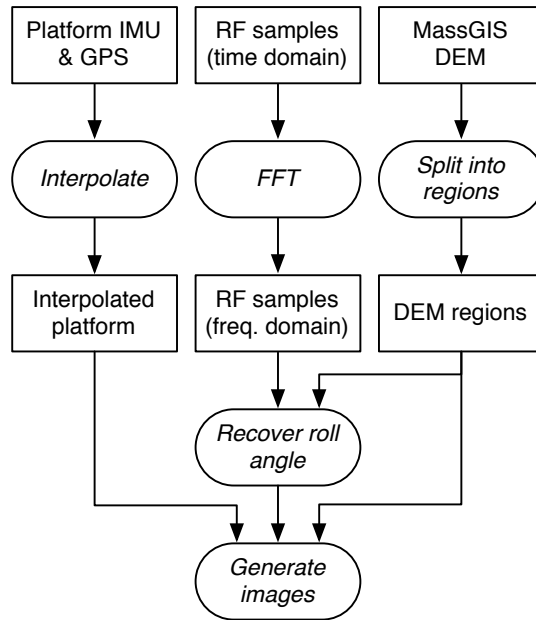


Figure 5.1: Data processing flow

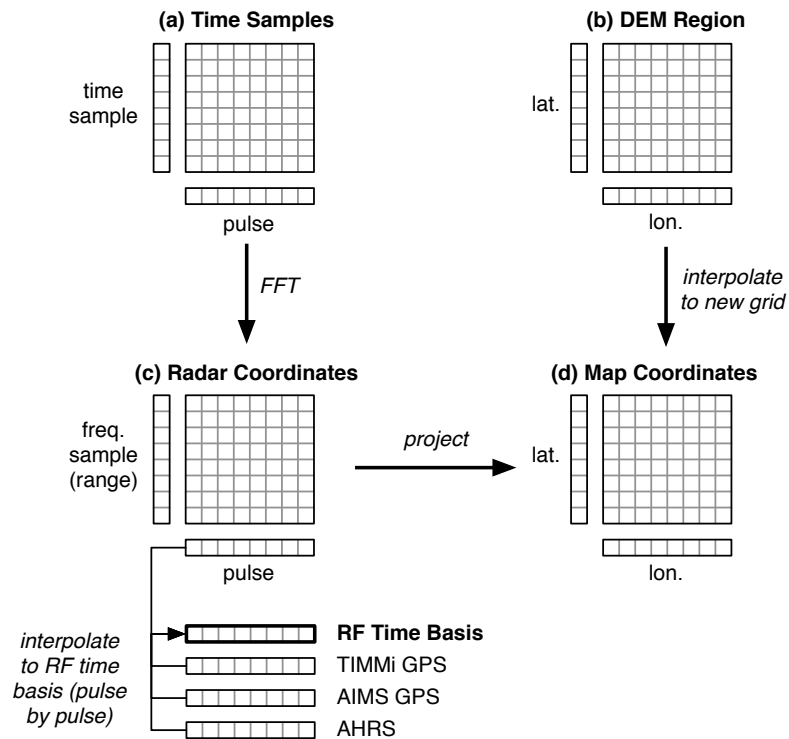


Figure 5.2: A detail view of the domain mapping scheme used to generate magnitude and interferometric images in map coordinates

flight yielded an incrementally better code base, so we will examine the processing routines for the most recent flight on June 11, whose scripts and functions are found in the `flight4_processing` directory of the Microwave and Remote Sensing Laboratory's (MIRSL's) code repository. They are prefixed with the string "flight4_".

MATLAB structure arrays are used heavily to separate each data source into its own namespace. For a more detailed understanding of these MATLAB structures, the reader is referred to Appendix ??.

5.2.1 Generating Previews

Immediately after a flight, it is critical to have an overall understanding of the quality and location of the acquisitions taken while in the air. The function `timmi_airborne_generate_previews` takes a `source` and generates a number of PNG images in radar coordinates for quick visual analysis. There are magnitude and phase images for channel 1, channel 2, and the cross-correlation between the channels. The script `flight4_generate_previews` is a wrapper for this function that generates previews for all flight 4 data sets in one batch.

The function also plots the platform's position in a Google Earth KML file, labeling points on an interval with the look number corresponding to the look index in the data set. This makes it convenient to determine the start and stop looks for a `region` of interest by simply examining the KML file projected over the visible imagery available from Google Earth or Google Maps.

5.2.2 Setup

The starting point is `flight4_setup`, a script that populates the `process` structure with all necessary data for the time-intensive geocoding and image creation routines that will follow. In other words, the setup script does everything short of creating magnitude and interferometric products in map coordinates. It is designed to batch process many `region` structures in sequence so that the routines may be run unattended.

The MATLAB script `flight4_setup` begins by calling `flight4_regions` to define the mapping regions of interest. These regions are rough outlines defined by the user that specify an outer boundary over which processing will occur. The setup script also loads the raw MassGIS DEM that is a superset of these regions. Next, the script parses the AIMS IMU and GPS

text files. Since these files do not change, this step only needs to be done once, and the results are saved to disk as a `.mat` workspace variable file for faster loading in the future.

Once the preprocessing step has been performed, each `region` is mapped to a `source` to tell the script where the data files are located on disk and how to process them. An important parameter is the `source.averaging_rate` which specifies how many output looks are generated per input look. As the number decreases, more pulses are averaged together and fewer output looks result. The minimum value for this variable is 1, which means that all pulses per look are averaged together, and the maximum value is `PULSES PER POSITION` as defined in the configuration text file. The `averaging_rate` must also be a factor of `PULSES PER POSITION`, so that an integer number of pulses can be averaged together.

Now that the map region, DEM, AIMS data, and data sources are configured, the `process` structure can be populated. The structure members `process.region` and `process.dem` are assigned to the region and DEM under consideration. The map coordinate space is configured by the `generate_map` function and is assigned to the member `process.map`.

The reading of RF samples from the binary data file has the longest duration among all setup steps. The source structure `process.source` is set and then `timmi_airborne_raw_read` reads the binary file specified by the source and applies the requested amount of pulse averaging. The raw time-series data is converted into the frequency domain via the Fast Fourier Transform (FFT), and the results are stored in the `process.rf` structure (§B.1.5). If the boolean variable `process.source.cached` is `true`, then the `rf` structure is saved to disk as a `.mat` workspace file to reduce overall processing time in future runs of the setup script. If a previously-cached `.mat` file is found on disk, it is loaded directly into memory and the function returns immediately. Manual deletion of the `.mat` file may be done to recover disk space or to force the reprocessing of raw FMCW data.

The final step of the setup script sets the time domain `process.t` and merges the platform data from the discrete instruments to create the `process.platform` structure. The `process` structure is now ready for further analysis.

5.2.3 Projecting to Map Coordinates

Currently, the only function that generates meaningful map products is `generate_mag_kml`. Initially designed to compile only reflectivity images, it has been expanded to also




Figure 5.3: (a) The entire map area is examined to find (b) the first beam. Then (c) a mask is created around that beam and (d) the next beam is located within this mask, a subset of the map. Steps (c) and (d) continue for every subsequent position.

produce interferometric images as well.

The projection of radar coordinates into map coordinates relies on the function `elevation` from the MATLAB Mapping Toolbox. `elevation` takes in the position of the aircraft and returns the elevation angle, slant range, and azimuth angle to every point in map coordinates. Because a DEM is supplied as the base map, variations in elevation are taken into account when calculating elevation angle and slant range. With a priori knowledge of the antenna beamwidth (1°), azimuthal angle of the antenna boresight, and the maximum unambiguous range of the radar, we can filter the results of `elevation` to the pixels seen on the ground by an individual look or pulse.

The `slant_range` can be easily converted into a range bin by dividing by the range resolution, and then complex frequency samples can be simply averaged together for real aperture radar processing. The `elevation_angle` is related to θ in Figure 2.4, and the platform roll angle corresponds to α . These values can be used to calculate the expected interferometric phase for every point along the ground, given by

$$\phi_{sim} = \arg [\exp (j2\pi B_a \sin (\pi/2 + \theta - \alpha))]. \quad (5.1)$$

MATLAB's `elevation` function is very processor-intensive takes a long time to complete, especially when done over every pixel in a map region. For the first iteration of the first radar look, `elevation` is computed over the entire map coordinate space (Figures 5.3a and 5.3b). For future iterations, we rely on knowledge of the previous look's antenna pattern and

the fact that the beam does not move very far along the ground from one look to the next. The ground-projected beam from the previous look is dilated using the `imdilate` function from the Image Processing Toolbox. This creates an “expanded” version of the beam pattern, which is called the `map_mask` (Figure 5.3c). It is assumed that the beam for the current look is a subset of the pixels in this mask. Then `elevation` and the subsequent beam pattern are calculated, but this time from within the context of `map_mask` (Figure 5.3d). This is much faster because it only considers the limited set of pixels within the mask. If the beam is not found within the mask, an `'Empty mask'` exception is thrown.

If the platform’s position ever falls outside the bounds of the `region`, an exception is thrown and the process aborts, so `region` must be defined with padding around the platform track to account for human measurement error.

In addition to generating 2-D map matrices like `process.map.mag1` and `process.map.int`, the function saves out a series of Google Earth KML images for these products. Each pair of PNG images and KML files are compressed into KMZ files for portability. Examples of these images are presented in Chapter ??.

5.3 Recovering Aircraft Roll Angle

One interesting consequence of the interferometric information is the ability to extract the roll angle of the aircraft platform. This is important because the interferometer is very sensitive to roll angle when measuring height.

Recall that for an interferometer’s phase sensitivity to height is

$$k_z = \frac{2\pi B_a \cos(\theta - \alpha)}{\lambda r \sin \theta} \quad (5.2)$$

for a target at range r and look angle θ , baseline separation B_a , and a baseline tilt of α , as shown in Figure 2.4.

The TIMMi antennas are located on the right-hand side of the airplane and point in that direction (also known as the starboard side). As the airplane rolls right (clockwise when looking from the tail towards the nose), the baseline tilt angle α decreases, decreasing k_z and resulting in faster-varying interferometric fringes in the range direction. Conversely, as the airplane rolls left




Figure 5.4: Recovering aircraft roll angle

(counterclockwise), the fringes become less frequent as the phase sensitivity to height increases.

For every platform position, we use the existing MassGIS DEM to model the expected differential phase along the ground. We then vary α until the difference between the modeled and measured phase is minimized. Sometimes the global minimum is discontinuous in the along-track direction of the aircraft and is therefore not a correct solution, in which case a local minimum closer to the previous roll angle is found.

Figure 5.4 shows how well the estimated roll angle matches the direct measurements from the inertial measurement unit (IMU), with errors on the order of 0.5° . It is noted that the IMU shows both a time delay and amplitude dampening effect. Because the IMU relies on inertial changes to make its measurements, it takes time for the IMU to “catch up” to the true roll angle of the platform. In addition, it has limited ability to respond to quickly-changing conditions, which makes it act like a low-pass filter. This accounts for the dampening effect on the angle measurements.

However, it was determined in processing that this result should not be used directly to generate interferometric products. By modifying the roll angle to fit what is observed, we are essentially removing the effects of what we are trying to measure. Instead, this analysis may be used to inform the user of the quality of the IMU data and the time offset between the two data streams.

CHAPTER 6

INAUGURAL RESULTS

6.1 Summary of Airborne Deployments

6.1.1 February 3, 2012

The first airborne deployment was designed to be an engineering flight to determine the feasibility of the radar as an airborne instrument and to work out any major problems in the system. The rack-mount configuration from previous ground deployments was reduced in size and weight, but it was still too large to accommodate AHRS in the cabin, so there was no platform orientation data for this flight. As we were sitting on the runway prepared for takeoff, the coaxial connector for the navigation computer's GPS receiver broke, so we were not able to fly the planned flight lines.

The baseline separation was 8 cm, with a look angle of 45° from nadir. Beginning at an altitude of 1000 feet AGL with a 4 ms chirp, it was determined there was sufficient SNR to achieve higher altitudes and shorter chirps. We ended up at 2000 feet AGL with a 1 ms chirp.

Despite the setbacks, this flight achieved its goal. In all, 45 minutes of RF data were collected, including GPS latitude and longitude (but not altitude).

6.1.2 April 13, 2012

This flight was meant to test the new "flat" supporting architecture, as well as collect AHRS data alongside TIMMi. The aircraft flew at 3000 feet AGL with a 2 ms chirp. The antennas were pointing 60° from nadir, and were reconfigured with the transmit antenna in between the two receive antennas, achieving a baseline separation of 28 cm.

Unfortunately, coupling between the antennas dominated the receive channels. To correct for these large-amplitude signals, the digitizer was configured in-flight to truncate to a certain peak voltage. This caused high-frequency returns to be effectively “thrown out”, reducing the SNR to a point where the interferometric data is was unrecoverable.

6.1.3 May 18, 2012

This flight was intended to test the behavior of baseband high-pass filters designed to reduce the near-field antenna ringing. We flew at 2000 feet AGL with a 1 ms chirp with a 60° look angle. The baseline separation was returned to 8 cm with the transmit antenna on the bottom.

This deployment suffered from aircraft power issues. A new AHRS unit was put into place, but combined with TIMMi it drew too much current from the airplane. The 35 A circuit breaker was tripped multiple times, so only a few minutes of data were collected. There was enough data to reveal that that the 60° boresight angle did not provide sufficient signal return for good interferometric data.

6.1.4 June 11, 2012

This was intended to be another attempt of the May 18 flight. The power issues on the Cessna were resolved by installing a 50 A circuit breaker in the aircraft. As on the February 3rd flight, the antenna boresight angle was adjusted to 45°.

This was the most complete dataset to date. Near the end of the flight, it was determined that the baseband high-pass filters installed for the May 18 flight were causing a major phase bias in the interferometric information, so they were removed from both receive channels in-flight. The final three flight lines of the deployment contain the best combined data of all flights so far. An overview of this flight’s coverage is shown in Figure 6.1.

Figures 6.2, 6.3 and 6.4 are subsets of this flight that illustrate results in the following section.

6.2 Reflectivity

Reflectivity is the most basic image that can be generated from a radar. While not the most valuable product of an interferometer, reflectivity images are easy-to-read and are useful




Figure 6.1: The planned flight lines for June 11 overlaid with the actual platform track

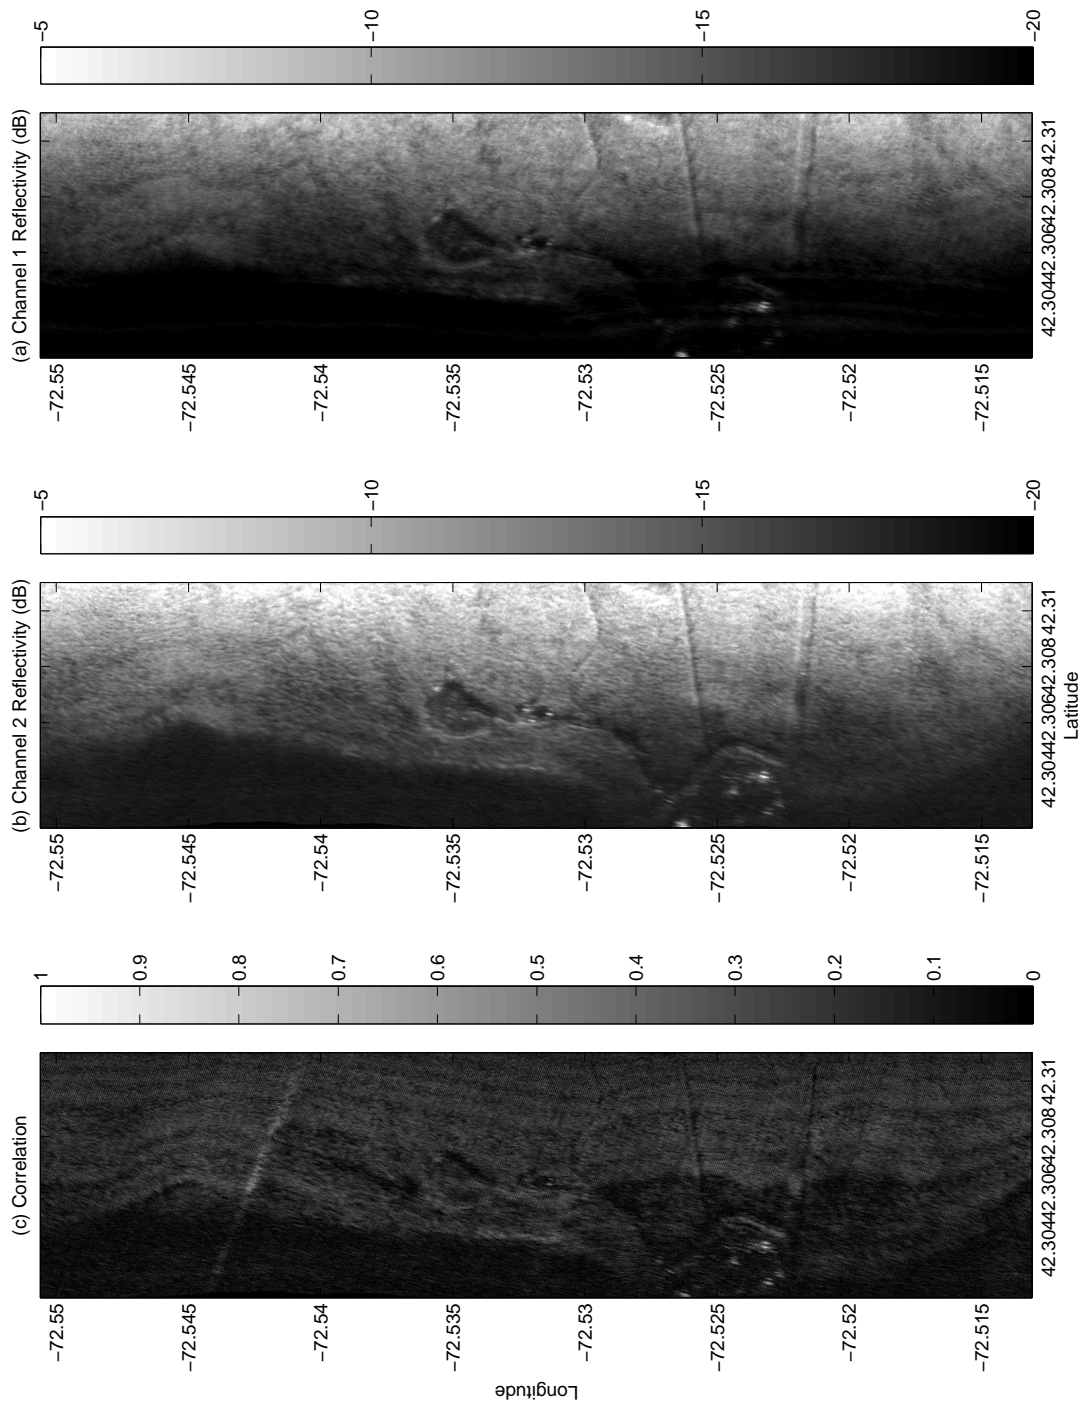


Figure 6.2: (a), (b) Reflectivity and (c) correlation in map coordinates

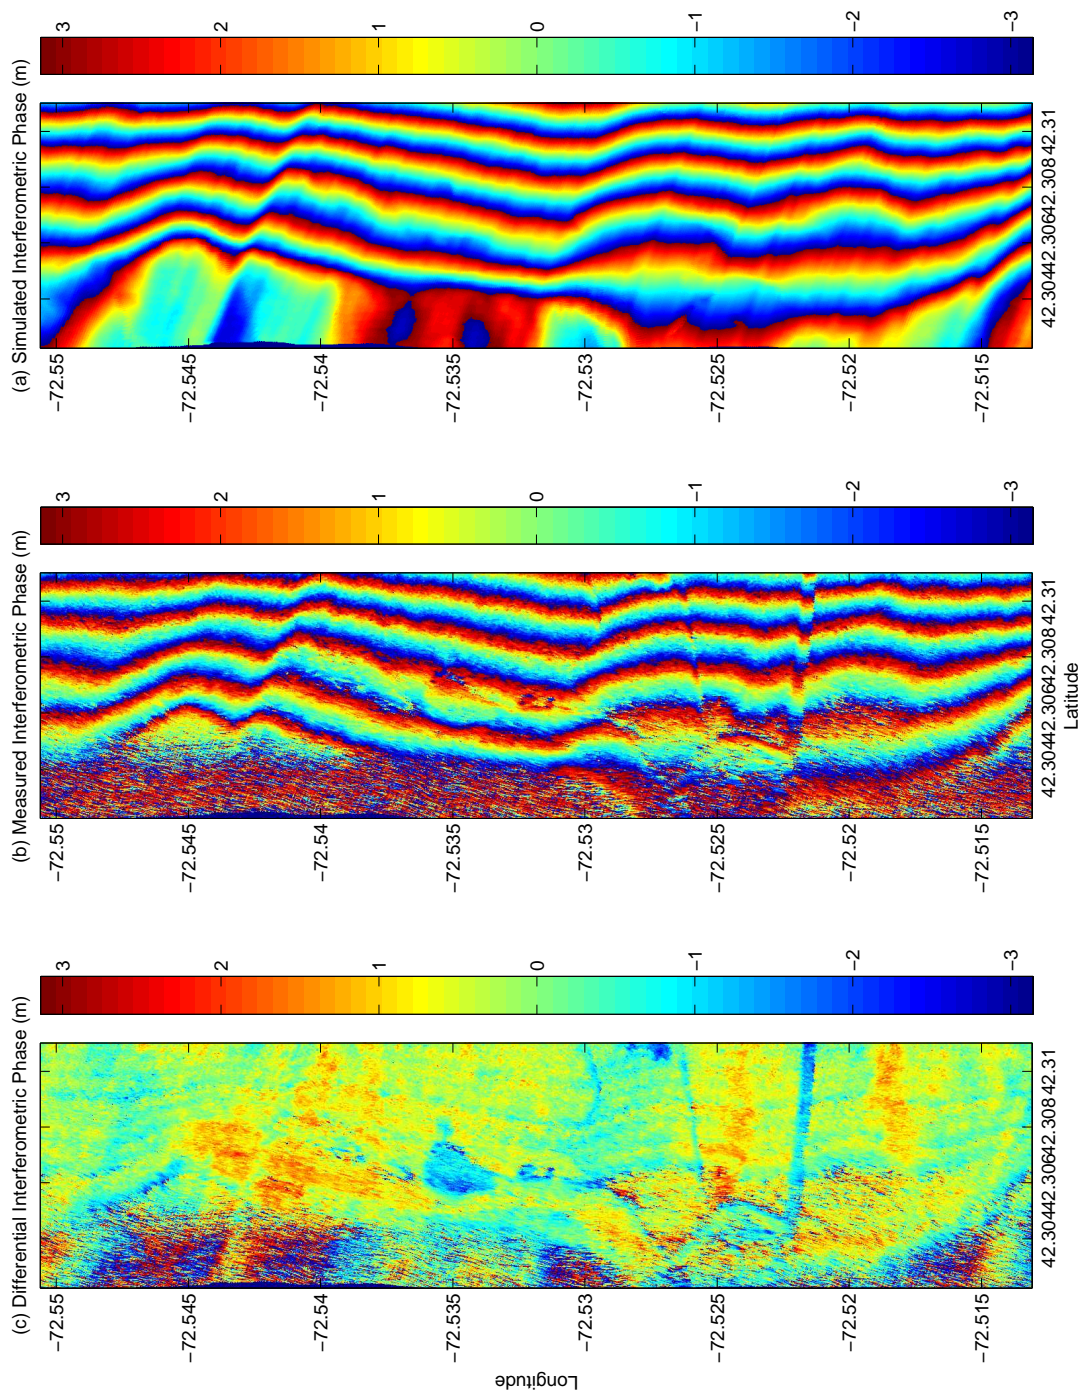


Figure 6.3: (a) Simulated, (b) measured, and (c) differential interferometric phase in map coordinates

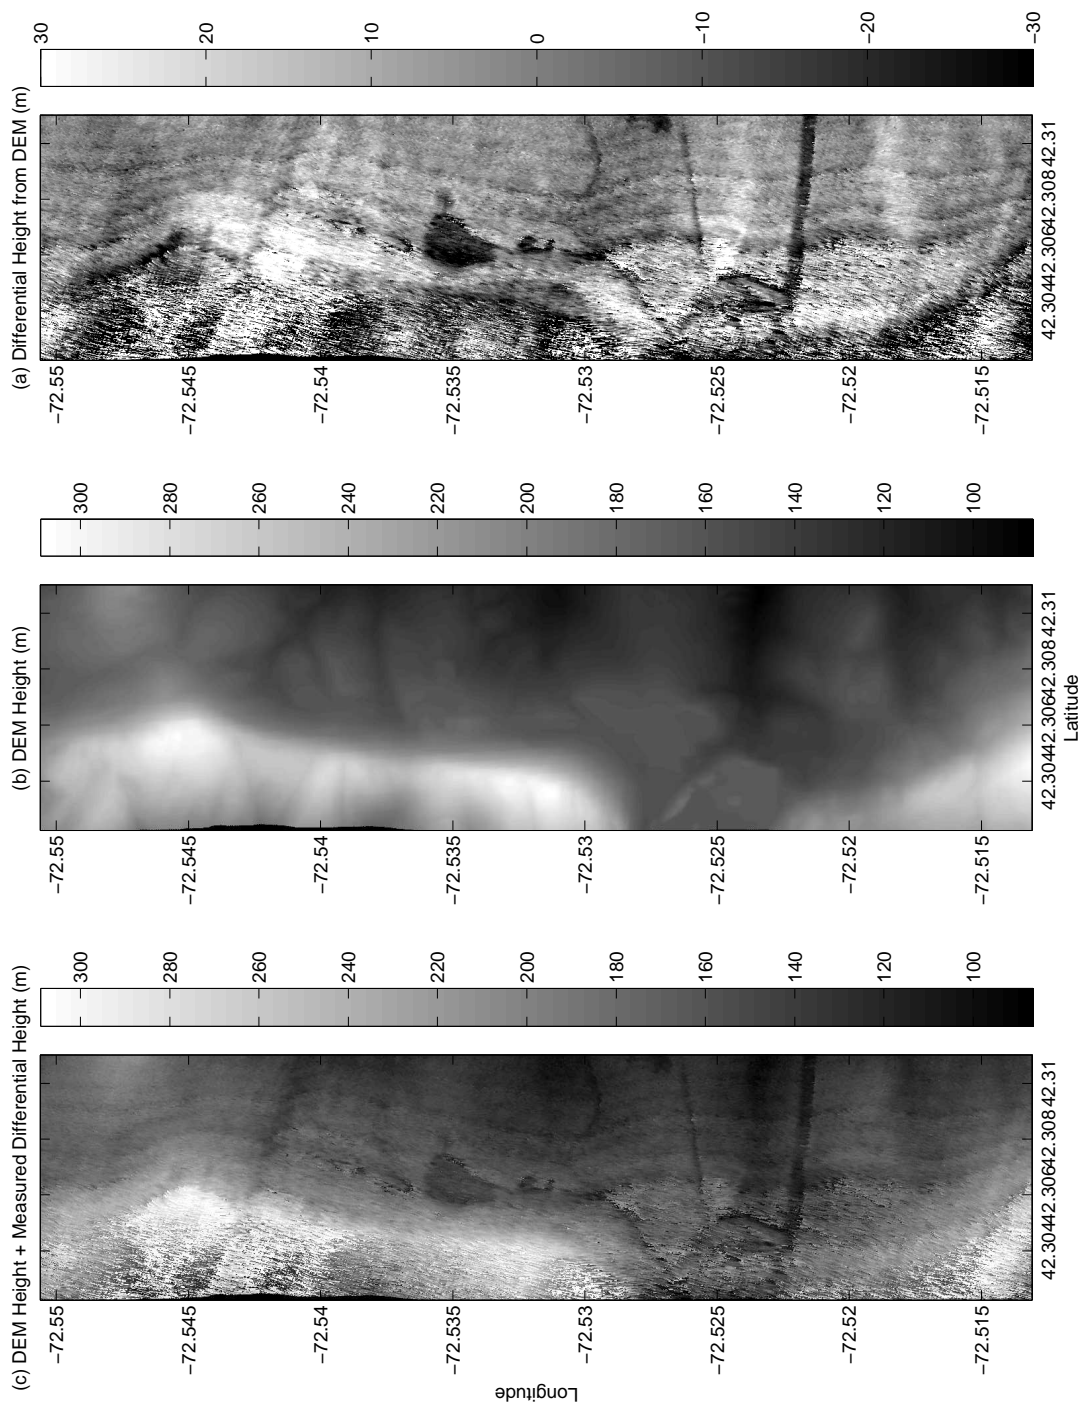


Figure 6.4: (a) Heights from interferometry, (b) the DEM, and (c) these two added together. Note that (a) has a smaller vertical scale to bring out subtle height variations.

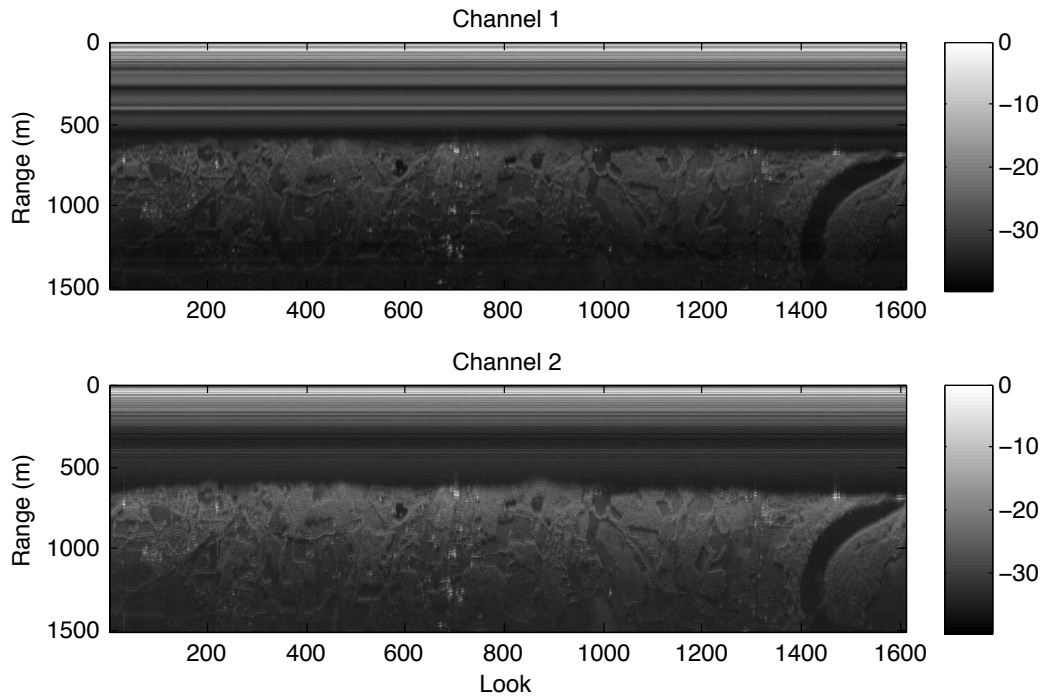


Figure 6.5: Reflectivity for one flight line in radar coordinates, normalized to the peak return power

for understanding what the radar sees. Figure 6.5 illustrates reflectivity images for both channels before map projection. Compared to the real world, there is a geometric distortion in the cross-track direction because cross-track resolution along the ground decreases as the look angle increases, resulting in a “stretching” of the image with range. FMCW coupling effects are clearly seen prior to the nadir return. The distance to nadir may be interpreted as a direct measurement of the aircraft height above the ground. This particular set of images shows a range of distributed targets such as fields and treetops, as well as some bright point targets. The dark area to the right is a portion of the Connecticut River. The incidence angle is not steep enough to obtain backscatter from water[1].

Another set of reflectivity images is given in Figures 6.2a and 6.2b. The RF data has been projected into map coordinates. This region is the foothills of the Holyoke Range in southern Amherst. This area is largely comprised of trees. Some relief is seen where shadowing occurs on the far side of the mountain ridge and at the edges of the tree line. The point targets are buildings and machinery in a quarry. Depressions are also visible for a road, two power lines, and a large clearing. There exists a disparity in dynamic range between the two channels. Channel 2, the



Figure 6.6: A composite reflectivity image for five flight lines taken on June 11, 2012; the solid black lines indicate the aircraft track; north is up

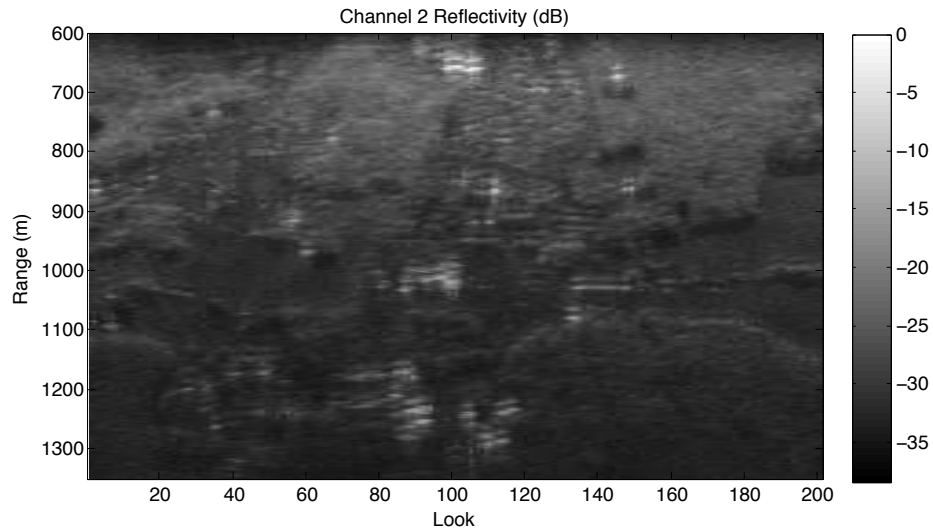


Figure 6.7: A worst-case example of bright targets spreading in range and azimuth

bottom antenna, has consistently exhibited more gain than channel 1, the top antenna. This may be due to the physical placement of the antennas, or it may be simply due to gain differences in the channels in the receiver. The receiver is presently in the process of being fully characterized from end-to-end to determine the cause.

Figure 6.6 is a large composite image from channel 2 of five separate flight lines all in the same region, and demonstrates the ability of the system as a whole to gather and process images over a large area. This image, along with additional interferometric products not shown here, took about 4 hours to generate on a personal computer. Notable geographic features include the Connecticut River flowing from north to south and Mount Toby in the southeast part of the image.

These images give a glimpse into the high resolution available to a 35 GHz system. However, image quality suffers primarily due to two factors: real aperture radar processing, and blurring of targets in range (Figure 6.7). Techniques for resolving these issues are discussed in §7.2.2 and §7.2.3.

Also evident is the dependence of backscatter power on look angle. Close to nadir where the look angle is steep, much of the power transmitted is reflected back towards the radar, resulting in a strong signal and a high SNR. As the look angle increases and approaches grazing angles, more power is reflected away from the radar to dissipate into space. This is especially apparent on the far side of Mount Toby, where the ground slopes away from the radar. When

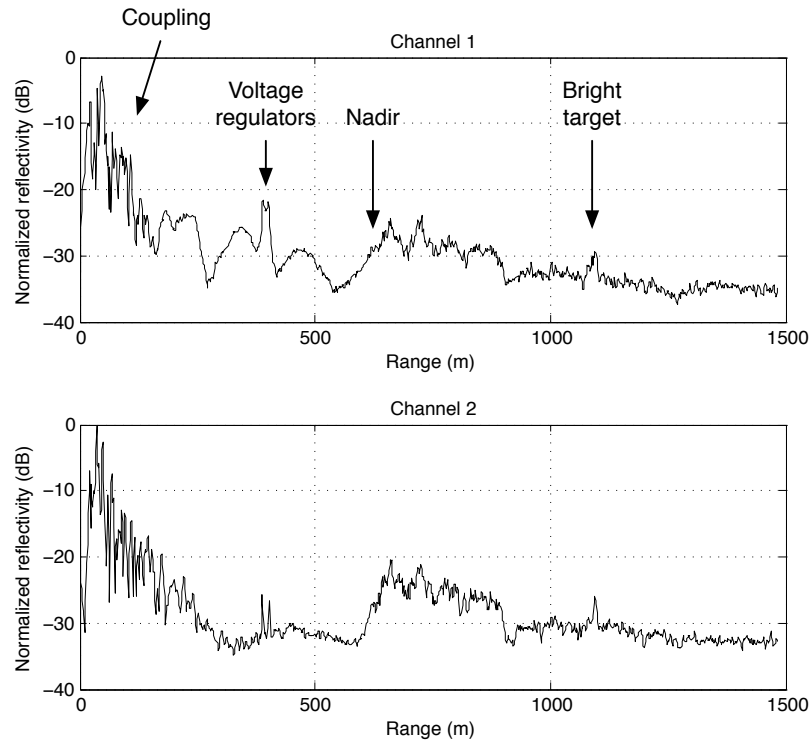


Figure 6.8: Magnitude spectrum from a single look

the terrain is so steep that it completely obscures the transmit signal, shadowing occurs, and no information can be retrieved. This effect can also be seen in the Mount Holyoke datasets.

The magnitude spectrum in Figure 6.8 highlights how different features appear after consecutive pulses have been averaged together to a single look. In this case, 100 pulses at a PRF of 1 kHz have been averaged. The decibel scale is measured relative to the strongest frequency component. The near-range FMCW coupling and the undesired voltage regulator tones occur before the nadir return, so they are not seen in the previous figures.

6.3 Interferometry

An example of a raw interferometric phase product is given in Figure 6.9. The quality of the fringes is considered to be exceptionally good, and has excellent coherence as compared to prior ground deployments of the Ku- and Ka- band radars. As expected, the phase becomes more noisy at shallow incidence angles in the far range, coincident with the aforementioned low

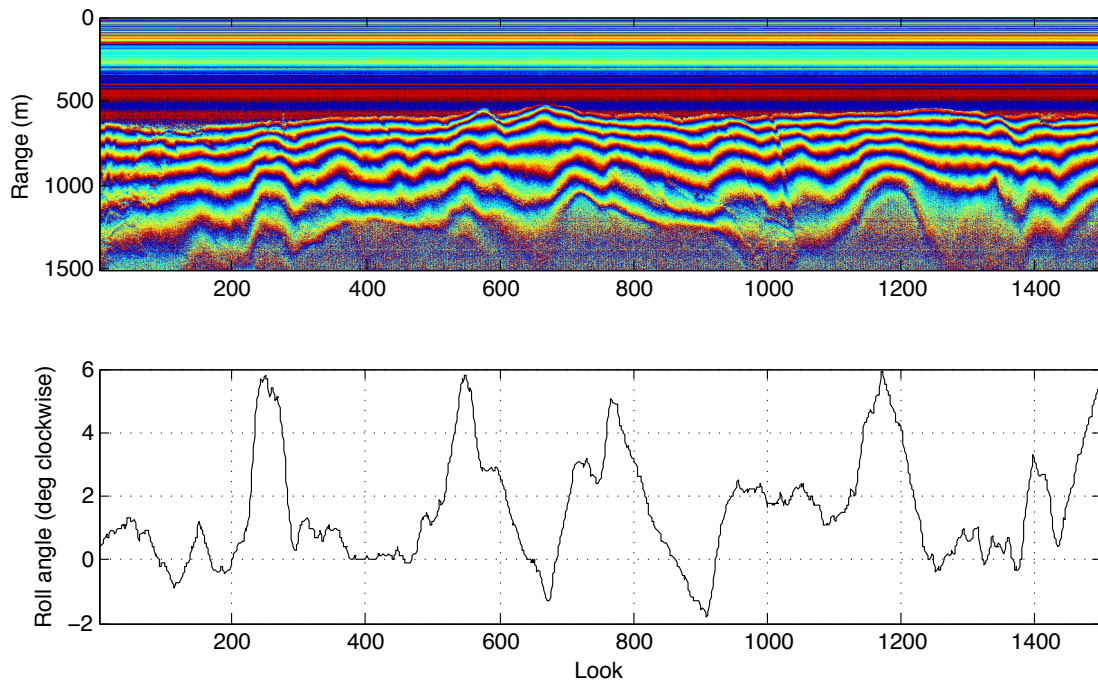


Figure 6.9: Interferogram (top) and aircraft roll angle (bottom)

backscatter power. This is an image of a mountain range, so shadowing is also present on the far side of the ridge. In these regions, the phase is nearly random, though there is a bias present in the noise.

The dominant factor in the fringe rate is not actually due to topography, but instead to the quickly-changing roll angle α . This is evident by comparing the aircraft roll angle for the same sequence of looks, which is plotted below the interferogram. In this figure, it can be seen that as the aircraft momentarily rolls clockwise, the fringe rate increases accordingly. These variations must be accounted for, as discussed next.

6.3.1 Variations From The DEM

A series of interferometric images taken during the June 11 flight are presented in Figure 6.3. The data shown in this figure are an excerpt from a single 10 km flight line north of the Holyoke Range in Hadley, MA. Seven other flight lines were mapped during this deployment for a total along-track length of 140 km. The area shown is about 3 km in the along-track direction and 1 km in the cross-track direction. The aircraft was travelling west to east with the antennas pointing out the right-hand side, so the aircraft tracked to the north of the area.

Image (a) shows a simulated interferogram as projected on the MassGIS DEM, and (b) shows the measured interferogram. From a wide perspective, the images are very similar. However, smaller features not captured in the DEM stand out as phase shifts in (b). In particular, the large clearing, power lines, and road stand out as variations in the fringes.

The differential phase between (a) and (b) is shown in (c). The depressions of forest clearings are very visible as the cool, blue colors. There are also some areas of increased height in the forested regions. Note that this image is a simple phase measurement. The factor of k_z to convert phase to relative heights has not yet been taken into account.

Figure 6.4 takes the interferometric products a step further by converting them to actual heights. Figure 6.3c is divided by k_z on a pixel-by-pixel basis to create the differential height map in Figure 6.4a. Image (b) is the base DEM altitude on a larger vertical scale, and the mountain range is very obvious. Finally, by adding (a) to (b), we can generate a new height model (c) with a high spatial resolution based off the measurements.

Some features in the base DEM, such as the drainage depressions around 72.54° W longitude, are actually smoothed over by the measured data. This makes sense because the DEM measures altitude at the ground level, whereas the radar's high frequency does not allow it to penetrate the canopy. The trees in this area have grown to similar heights and obscured the ground variation beneath the tree tops.

Conversely, variations in the forest are not represented in the DEM, so features such as the aforementioned forest clearing (72.535° W), road (72.530° W), and power lines (72.525° W and 72.522° W), stand out exceptionally well.

The interferometric phase is random in shadowed regions, resulting in the random height jumps that seen in the far range. These numbers are not meaningful, and it would be beneficial for higher-level products to filter out these areas based off a correlation threshold.

Figure 6.10 shows an elevation profile comparing the base DEM with radar measurements. This is a cross-track profile with the radar located at the right of the plot. There is good agreement between the plots from 1600 m to 800 m. At this point, the peak of the mountain blocks the signal from reaching the descending slope, and the measurements begin to deviate from the ground truth, quickly becoming a fully noisy signal.




Figure 6.10: Elevation profile from near-range (right) to far-range (left) of base DEM (solid) and the new DEM (X's)

6.3.2 Error Sources (NEW)

The distinct ripple in the differential phase in the cross-track direction (Figure 6.3c) is an outstanding source of error. This translates to an inaccurate height measurement (Figure 6.4a). The pattern of the ripples closely follows the fringes in Figure 6.3b but at twice the fringe rate, indicating that the measured phase falls above or below the expected value twice with every 2π phase wrap.

(TODO: explain here)

Errors in platform orientation and position have a significant impact on image (c). If there is a time offset between TIMMi and AHRS, then the simulated interferogram can “lag” or “lead” the measured interferogram, generating incorrect differential phase. Errors in yaw angle of the aircraft also play a similar role. This is the cause of the along-track height ripples in (c).

CHAPTER 7

SUMMARY AND FUTURE WORK

7.1 Contributions

The following is a bulleted list of contributions arising from this thesis work which began in September 2011.

- Redesigned the radar's supporting electronics, achieving portability and effortless deployment. This included work on the power supply, reduction of system size, and the upgrading of system data storage to a solid-state hard drive.
- Rebuilt the entire transceiver, diagnosing and fixing problems.
 - Characterized transient tones from voltage regulators onboard the DUC and DDC and attempted to filter them out of the baseband signal in several ways.
 - Reconfigured waveguide components in the transceiver to accommodate the Ka-band power amplifier from within the enclosure.
 - Rewired power supply blocks and connectors to accommodate additional 9 V and 15 V supplies.
- Designed and implemented baseband high-pass filters to reduce effects of antenna coupling and increase the system's dynamic range. However, these filters were found to have a severe impact on interferometric data quality, and are no longer used.
- Transitioned the mechanical antenna support structure to the aircraft door.
- Planned and executed deployments on the aircraft platform (Cessna 206). Developed a set of tables, formulas, and mapping tools to aid in designing flight lines.

- Developed a comprehensive codebase for the analysis of radar data and generation of higher-level products.
 - Used these tools to develop an intimate understanding of interferometric radar concepts, which in turn informed iterative improvements in the code itself.
 - Focused on achieving adequate performance on modest hardware by reducing processor load and memory usage.
 - Implemented a workflow that has a quick turnaround time, permitting same-day image generation following an acquisition.
 - Documented code thoroughly via inline comments and Appendix ??.
- Modified acquisition software to support continuous acquisition, integrated a GPS data stream, and added user interface elements for real-time feedback on the system's behavior.
- Completed four field deployments of the system.
- Completed a number of static, ground-based test deployments to characterize system performance and feasibility.
- This work was presented at the IEEE International Geoscience and Remote Sensing Symposium (IGARSS) 2012 conference in Munich, Germany. A four page paper was submitted and published in the conference proceedings.

7.2 Future Work

7.2.1 Combining AHRS and TIMMi Datasets

There have been a number of problems encountered in the merging of the AHRS platform orientation data with the TIMMi RF data. Despite the GPS time synchronization, there appears to be a delay on the order of 1 second between the two systems. This may also be due to the aforementioned inertial delay inherent in the IMU. Finding a way to accurately correct for this time offset will be necessary in the merging of these datasets.

Another difficulty comes from the measurement of yaw angle. AHRS provides aircraft yaw angle via differential GPS measurements, but the antennas do not point directly broadside

due to the tapered shape of the fuselage. Determining the angle between the aircraft nose and the antenna pointing angle using ground targets as a reference has proved to give inconsistent results.

This information is critical to the augmentation of interferometric heights over the existing DEM. If a measured height is mapped to the wrong location on the DEM, the differential height between the DEM “ground truth” and the measurement will be skewed. Any persistent roll, yaw, or time bias will impact every measurement throughout a flight line.

7.2.2 Unfocused and Focused SAR

For real aperture radar (RAR) processing, the cross-track resolution of the radar is limited by the azimuthal beamwidth of the antenna pattern, in this case $\phi_B = 1^\circ$. The resolution is range-dependent because the beam widens as range increases. For example, when $r = 1.5$ km, the beam is 26 m wide. This is much worse than the range resolution of 1.5 m.

Synthetic aperture radar (SAR) processing techniques may be employed to significantly improve the cross-track resolution. The theoretical resolution achievable through focused SAR processing is half the along-track antenna length[3][5], which is $\Delta x = L_a/2 = 0.24$ m for the Ka-band slotted waveguide antennas. As an additional benefit, the resolution is not range-dependent.

A limiting factor in SAR data acquisition is that the radar must transmit at least one pulse each time the platform travels forward at one-half the antenna length. At a typical cruising speed of $V_{st} = 120$ kts and a PRF of 1 kHz, we are able to achieve 3.89 pulses per half-antenna-length, which is more than sufficient for focused SAR. The Cessna 206 used for this work is also outfitted with a short take-off and landing (STOL) kit that provides extra lift and enables it to fly at speeds as low as 55 kts if necessary.

In order for fully focused SAR to be successful, compensation for platform motion must be well known throughout a target’s entire illumination period. Unfocused SAR may be seen as a compromise that offers improved resolution over real aperture radar at the cost of a lower SNR and range dependence. For a target at closest cross-track approach range R_0 , the achievable resolution is $\Delta x = \sqrt{\lambda R_0}/2$ assuming accumulation of phase error no greater than $\pi/4$. For a worst-case comparison at $R_0 = 1.5$ km, this resolution is about 2.5 m – about ten times worse than focused SAR.

7.2.2.1 Work So Far

SAR may be conceptualized from several different perspectives. We have attempted to construct SAR images using three distinct approaches with limited success. The first pass was done in map coordinates. With a priori knowledge of the DEM and, thus, range to every ground pixel at every platform position, the expected signal phase can be calculated. The focusing occurs by taking out “extra” phase due to pixels being away from broadside, forming the synthetic aperture.

The second iteration approached the problem from the Doppler perspective. Ground targets experience a Doppler frequency shift due to the relative motion of the aircraft, and can be calculated easily. An FFT is taken of a sequence of samples from a particular range bin, yielding the Doppler response of those targets. Each frequency bin from the FFT is mapped onto a Doppler frequency bin on the ground. For large yaw angles, the Doppler frequency of the approaching or receding targets may exceed the rate of PRF/2, resulting in aliasing in the frequency domain. However, knowledge of the platform velocity and approximate yaw angle can be integrated to resolve this ambiguity. For example, the centroid of one FFT spectrum with $f_s = 1$ kHz was found to be at 400 Hz, but the yaw angle indicated that Doppler shift should be negative. A back-of-the-envelope calculation revealed the true shift to be -600 Hz.

The third attempt was undertaken in RF coordinates in a similar fashion to the first. For each range bin, a matched filter is constructed based on the antenna beamwidth, range, and distance between consecutive pulses. Each matched filter is then convolved with a series of all the samples in its range bin. This method is very fast to compute, but it will only focus a narrow set of ranges. I suspect this is due to incorrect assumptions in the geometry of the scene when generating of the matched filter.

An odd issue appeared with the phase measurements for individual channels that was not a problem until SAR was attempted. It seems that the phase is incremented by a factor of exactly π on every consecutive pulse. Multiplying the samples of *every other* pulse by $e^{j\pi} = -1$ yields phase measurements that vary smoothly with time. This is consistent for both channels because interferometric phase has been reliable up to this point. This problem likely originates from the conversion of time samples into the frequency domain in `timmi_airborne_raw_read`.

7.2.3 Correcting Target Range Blurring

A persistent problem seen across all ground and airborne deployments is the replication of targets across many range bins. This is especially noticeable for bright targets, whose sharp response can be seen to have many ripples in range. The cause of this blurring in the cross-track direction is not understood right now.

Meta et al. [7][8] identify the major cause of range blurring in FMCW radars as nonlinearities in the transmitted waveform. We have characterized the waveform of the Tektronix AFG3252 in the lab, and have found that it indeed has a very nonlinear response. However, the undesired images due to these nonlinearities fall out of the pass-band of the receiver, so they should not impact the image. More work must be done in order to understand the full impact of these nonlinearities, and if any distortion is introduced during upconversion. We also plan to do a ground deployment with an Agilent N8242A arbitrary function generator, which we have already seen has a very different behavior across the 100 MHz chirp, but we are unsure how this will impact detection of target range.

Another potential cause of this blurring is due to multipath reflections. The incident RF energy impinging on the ground is re-radiated in all directions. These re-radiated signals can reflect off of different parts of the aircraft and coherently (or incoherently) add with the desired signals, hence creating an along-track banding in the reflectivity and interferometric images.

The blurring degrades the quality of the backscatter images. More importantly, this alters the interferometric phase from its true value because the phase responses of independent targets are accumulated coherently when they are not correlated in reality.

7.2.4 Additional Deployments

There are a number of hardware changes that might improve the quality of data in future deployments. Most notably, the Tektronix AFG would be replaced by an Agilent N8242A function generator in an attempt to reduce blurring in the range direction due to nonlinearities in the baseband chirp. The Agilent device also supports a 200 MHz bandwidth to which TIMMi is transitioning.

We would also like to increase the physical isolation between the transmit and receive antennas by placing a barrier between them. A piece of RF-absorbing foam would be a likely candidate for the separating material. This analysis can be performed prior to flight, and is most

easily tested by pointing the entire antenna assembly at the sky in an area free of obstructions.

The Harvard Forest is an area of ongoing research for many members of MIRSL and the scientific community. It is an important testbed for understanding the use of remote sensing techniques in estimating how much biomass is contained in wooded areas. At MIRSL, we are currently analyzing lidar footprints and radar images over the Harvard Forest taken from NASA's LVIS[2] and UAVSAR[6] instruments. The UAVSAR backscatter images were taken at L-band. A TIMMi flight over the same area would create a unique interferometric dataset that could be corroborated with the existing datasets, ground truth, and new instruments being flown this year.

7.2.5 Convert RF Datasets to a Standardized Format

As it stands, the handling of the raw RF data files is somewhat cumbersome. The files are not self-descriptive and rely on prior knowledge of the specific format in order to be read. This comes with a high startup cost if one wishes to load the data into another application, and makes it difficult for future users to read archived data if documentation cannot be found.

I propose an additional step immediately following acquisition to convert the raw RF file into a standardized file format designed for storing multidimensional data for scientific applications. The file could include additional metadata related to radar acquisition software configuration (e.g. PRF, sampling rate), hardware configuration (e.g. baseline separation), and additional notes from the operator. Its primary payload would be an array of waveforms for each pulse and an array of GPS measurements. The concept of a "look" would no longer be tied to the data format, and would instead just be an input parameter to the acquisition software to tell it how often to offload samples from the ADC.

NetCDF¹ is one such format that meets all these specifications. It is commonly used in the geosciences, and is an open standard that has cross-platform support and is widely supported. HDF5² has similar features and may also be satisfactory.

¹<http://www.unidata.ucar.edu/software/netcdf/>

²<http://www.hdfgroup.org/HDF5/>

7.2.6 Add Flexibility to Antenna Mounting Bracket

In order to achieve a strong return from the surface of water at Ka-band, TIMMi must be able to take measurements within a few degrees of nadir. The physical configuration of the antenna mounting bracket does not currently allow for this. The antenna mounting bracket must be modified in such a way that the baseline is not perpendicular to the ground. It may be possible to manufacture an extension that allows the mount to tilt away from the aircraft. Another option is to utilize the instrument port in the bottom of the aircraft, placing the antennas parallel to the ground and looking in the cross-track direction.

The easiest way to test what incidence angles are required is to complete a circle flight over a body of water at varying banking angles. The roll angle allows you to tilt the baseline without physically tilting it in relation to the aircraft. The Quabbin Reservoir would be an ideal candidate for such an experiment.

APPENDIX A

NOTATION REFERENCE

A.1 Symbolic Notation

α	Baseline tilt angle (rad)	Ω	Platform roll angle; clockwise is positive (rad)
B	Bandwidth (Hz)	ϕ	Platform pitch angle; up is positive (rad)
B_a	Antenna baseline separation (m)	ϕ_B	Azimuthal antenna beamwidth (rad)
c	Speed of light in free space ($3 \cdot 10^8$ m/s)	r	Range (m)
f	FMCW target frequency (Hz)	Δr	Range resolution (m)
f_{IF}	Intermediate frequency (Hz)	R_0	Closest cross-track target range (m)
f_{LO}	Local oscillator frequency (Hz)	P_{out}	Output power (W)
f_p	Pulse repetition frequency (Hz)	t_r	Trigger re-arm time (s)
f_{RF}	Radio frequency (Hz)	τ	Chirp/pulse length (s)
f_s	Sampling rate (Hz)	θ	Cross-track boresight angle from nadir (rad)
Δh_u	Unambiguous height change (m)	θ_B	Elevation antenna beamwidth (rad)
k_z	Phase sensitivity to height (rad/m)	θ_L	Antenna boresight angle (rad)
κ	Platform yaw angle; clockwise is positive (rad)	V_{pp}	Peak-to-peak voltage (V)
L_a	Along-track antenna length (m)	V_{st}	Along-track platform velocity (m/s)
λ	Wavelength (m)	Δx	Cross-track resolution (m)
N_p	Number of pulses		
N_s	Number of samples		

A.2 Abbreviations

AFG	Arbitrary Function Generator	NetCDF	Network Common Data Form
AGL	Above Ground Level	NI	National Instruments
AHRS	Altitude and Heading Reference System	NMEA	National Marine Electronics Association
AIMS	Airborne Imaging Multispectral Sensor	PCI	Peripheral Component Interconnect
BB	Baseband	PXI	PCI eXtensions for Instrumentation
DDC	Dual Down-Coverter	RAID	Redundant Array of Independent Disks
DEM	Digital Elevation Model	RAM	Ripple Attenuator Module
DOP	Dilution Of Precision	RAR	Real Aperture Radar
DUC	Dual Up-Converter	SAR	Synthetic Aperture Radar
FFT	Fast Fourier Transform	SNR	Signal to Noise Ratio
FMCW	Frequency-Modulated, Continuous-Wave	TIMMi	Topographic Ice Mapping Mission
HDF5	Hierarchical Data Format, version 5	STOL	Short Take-Off and Landing
IMU	Inertial Measurement Unit	SWOT	Surface Water Ocean Topography (Mission)
JPL	Jet Propulsion Laboratory	UAVSAR	Uninhabited Aerial Vehicle Synthetic Aperture Radar
LVIS	Laser Vegetation Imaging Sensor		

APPENDIX B

MATLAB PROCESSING CODE REFERENCE

B.1 The `process` structure

`process` is an all-encompassing structure that contains all input and output information for every processing step. It is convenient to be able to persist a single instance of the `process` struct out to disk using `save` and `load`, and to separate different processing workflows into different variables. It is made up of the following members.

B.1.1 `region`

The `region` structure is the smallest unit of data processing, and typically corresponds to an individual flight line. The DEM base maps are typically much larger than the flight line of interest, so the `region` defines a subset of that large DEM, which will be more efficient to process in later steps. Its fields are as follows:

`name` A canonical name for this region, used to generate unique output file names

`dataset` The RF dataset name, e.g. `'2012-06-11_120540Z_1'`

`looks` An array of look indices to process in the dataset, e.g. `570:985`

`UL.lat` The latitude of the upper-left corner of this region

`UL.lon` The longitude of the upper-left corner of this region

`LR.lat` The latitude of the lower-right corner of this region

`LR.lon` The longitude of the lower-right corner of this region

`dlat` The pixel size in latitude to interpolate the DEM to, in degrees, e.g. `2.5E-5`

`dlon` The pixel size in longitude to interpolate the DEM to, in degrees, e.g. `2.5E-5`

B.1.2 source

Once a `region` structure is populated with information about the DEM and map coordinates, the `source` structure is generated to give the processing scripts specific instructions on how to handle the reading of the raw data files. The `source` is the sole parameter passed to `timmi_airborne_raw_read` for the initial reading and FMCW processing of raw data files.

Typically, each `region` has a corresponding `source`, although there may be unforeseen cases where multiple `source` structures are created for an individual `region`.

directory The full path to the dataset

averaging_rate The number of averaged looks to generate from each raw “look”; when equal to 1, produces one output look per input look; when equal to N_{pulse} , produces one output look per pulse; must be a factor of number of N_{pulse}

start_look The first look index to process; use 1 for the first look

end_look The last look index to process; use -1 for the last look

gps_only Optional; set to `true` to only read GPS data and skip all RF data

waveforms_only Optional; set to `true` to only read raw time-series waveforms without performing any additional processing such as windowing and FFT

cached Optional; set to `true` to save processed RF data to a MAT file on disk for quick loading in the future

name Optional; if `cached` is `true`, specify a unique name for the cached file

B.1.3 dem

The `dem` structure handles the raw MassGIS DEMs generated using the ENVI mapping software. The ENVI header is a plain text file that can be read to manually populate these fields.

file The full path to the DEM data file

nsamp The number of elevation samples per line

nline The number of elevation lines

dlat The size of the elevation pixels in latitude degrees

dlon The size of the elevation pixels in longitude degrees

UL.lat, UL.lon The latitude and longitude of the upper-left corner of the DEM

LR.lat, LR.lon The latitude and longitude of the lower-right corner of the DEM

lat_vec, lon_vec Arrays of latitude and longitude coordinates for the grid

alt The 2-D matrix of elevation values

coords A 4-by-2 matrix, each row contains the lon/lat of one corner of the DEM bounds

B.1.4 map

The `map` structure is intended to store all data products that have been projected into map coordinates. It is first created from a region and a DEM using `generate_map`. This function sets up the geographic coordinate space specified by the `region` and interpolates the DEM onto this grid.

dlat, dlon The size of the grid pixels in latitude/longitude degrees

lat_vec, lon_vec Arrays of latitude and longitude coordinates for the map grid

lat, lon 2-D matrices of the coordinates for each pixel, created by `meshgrid`

alt The DEM interpolated into map coordinates

flat A synthetic flat Earth initialized at the minimum DEM elevation

B.1.5 rf

The `rf` structure is populated with information read that `timmi_airborne_raw_read` loads from disk. It reads the raw time samples and performs a fast Fourier transform (FFT) to transform the samples into the frequency domain – the primary way of determining range for FMCW radars. TIMMi GPS data present at the header of each look is also read, and parameters from the configuration text file are parsed and stored in the `config` struct.

mag1, mag2 Linear power magnitude for each channel in RF coordinates, averaged

slc1, slc2 Single Look Complex data for each channel in RF coordinates, averaged

corr The complex correlation between the two channels; `corr = slc1 .* conj(slac2)`

dn The datenum representation for each look, read from the 23-character Timestamp field (Table 3.4)

config Metadata and other radar configuration parameters; see §B.1.5.1

coordinates A two-column matrix of latitude/longitude coordinates for each look

gps A structure encapsulating all GPS information recorded in the look header, e.g. latitude, longitude, altitude, timestamp, etc.

looks An array of all look indices

B.1.5.1 rf.config

The `config` structure contains metadata about the acquisition that is necessary for processing. Most of the values are parsed from the configuration text file created during acquisition, but some values are hard-coded or must be filled in manually.

nlooks Total number of looks

nsamp Number of FFT samples per look (half the number of time samples)

npulse Number of pulses per look

dr Range resolution (1.5 m)

BW Radar system bandwidth (100 MHz)

fsamp ADC sampling rate

B Antenna baseline separation

source_dir, fname RF data file directory and name

time0, time1 Start and end acquisition times (first and last `dn` values)

time_str A human-readable date string representing this data set (the first `dn` value)

version The RF data file version (1 or 2)

averaging_rate The number of output looks per input look (see `source.averaging_rate`)

B.1.6 `t`

The function `rf_time_basis` takes an `rf` structure and generates the array `t` to act as a standardized time domain for all processing steps. It begins by taking every `datenum` value from the `rf.gps` struct and smoothing it to have a linear, monotonically-increasing trend. This is done under the assumption that the ADC trigger is extremely well-timed, whereas there is a varying delay in transferring this data from the ADC's memory into program memory and then to disk. `normalize_gps_time` finds a straight line that closely matches all the timestamps. Then `rf_time_basis` takes this new GPS time series and interpolates it to every output look, so that we have good idea of the time that every output look occurred.

B.1.7 `ames`

Unfortunately, this structure is a misspelling of AIMS, but it has not yet been refactored. The `ames` structure is populated by `interpolate_ames`, which takes as an input the time series `t` as well as the IMU and GPS structures read directly in from the data files using `read_imu` and `read_nmea`. Sometimes, there is an offset between the AIMS timestamps and what is expected on the order of one second. If there is a discrepancy in daylight savings time, this offset will be plus or minus one hour ($\pm \frac{1}{24}$ `datenum`). Errors in geocoding and interferometric products are very sensitive to an incorrect offset, because this factor determines how well the IMU and GPS data from AIMS line up with the GPS and RF data recorded separately by TIMMi.

Again, all these measurements are taken from the raw data files and interpolated to the universal RF time basis `t`, so they can be matched directly to RF coordinates on a pulse-by-pulse basis.

`imu.roll` Roll angle in degrees, positive is clockwise

`imu.pitch` Pitch angle in degrees, positive is down

`imu.azimuth` Absolute compass direction in degrees, only reliable for the "new" IMU

`gps.lat, gps.lon` Latitude and longitude

`gps.alt` Altitude above geoid plus geoid height above ellipsoid

`gps.track` "Track made good", the true direction of the platform's motion along the ground

`gps.speed_kts, gps.speed_kph` Speed along the ground in knots and kilometers per hour

B.1.8 platform

The `platform` structure is simply a layer of abstraction between the AIMS instrument data, the RF GPS data, and the processing algorithms. This structure is populated on a per-flight basis, and allows the mixing and matching of various data sources to arrive at the final picture of the platform's position and orientation at every instant in t .

For example, perhaps it was discovered that the AIMS GPS data stream was interrupted, then it could be supplemented or replaced by the TIMMi GPS data by simply substituting the appropriate members of the `platform` structure. In a real-world case, a correction factor had to be applied to the azimuth measurements of the June 11 flight because the AIMS differential GPS antennas were connected to the wrong ports on the GPS receiver.

Each flight has its own function, e.g. `flight4_platform`, that returns a `platform` struct created from any number of ancillary parameters.

alt Platform altitude in meters

lat, lon Platform latitude and longitude in degrees

pitch Platform pitch in degrees, positive is nose-down

roll Platform roll in degrees, positive is clockwise

look_direction A string indicating the broadside look direction, either `'right'` or (unsupported) `'left'`

speed Platform along-ground speed in meters per second

baseline_tilt Baseline tilt angle to vertical, positive means the top antenna tilts away from the aircraft, in degrees

nplat The number of platform positions (essentially the length of t)

heading Platform azimuthal look direction in true degrees

antenna_pointing Antenna azimuthal look direction in true degrees, typically `heading + 90 + fudge_factor`, where `fudge_factor` is the antenna mounting bracket's angle away from broadside

APPENDIX C

CHECKLISTS

C.1 Hardware Checklist

- Devices
 - NI PXI acquisition computer
 - Tektronix AFG3252 function generator
 - Power supply unit with GPS receiver and USB serial adapter
 - Linksys network switch
 - Notebook with ethernet port and Remote Desktop client

- RF
 - TIMMi Ka-band transceiver, tested
 - 3 × slotted waveguide antennas, tuned
 - 6 × flexible WR-28 waveguides
 - 6 × locking pins for positioning antennas
 - 1 × GPS puck antenna
 - Assorted WR-28 waveguide components

- Cables
 - 2 × CAT-5 ethernet cables
 - 4 × long BNC cables
 - 1 × short BNC cable

- 1 × 19-pin circular military power cable
- Tools
 - Hex drivers for machine screws
 - Adjustable wrench
 - Flat-head screwdriver
 - Multimeter
 - Oscilloscope and/or spectrum analyzer (optional)
- Miscellaneous
 - 4-32 nuts and bolts for connecting waveguides
 - SMA and BNC adapters
 - Extra BNC cables
 - Extra WR-28 waveguide components
 - Power strip, power cords, and AC adapters for every device
 - USB storage (FAT32 or NTFS) for offloading AIMS data
 - Aircraft door

C.2 Pre-Flight Checklists

C.2.1 In the Lab

- It is much easier to assemble the door hardware in the lab than to do it in the hangar attached to the aircraft. This includes the transceiver, waveguides, bulkhead adapters, antenna mount, and antennas. The entire unit may be carried like a briefcase using the handle on top of the transceiver.
- Note which COM port the USB serial adapter has acquired.
- Test the transceiver end-to-end. Connect everything to the aircraft door as it would be in situ – transceiver, antennas, everything. Insert a baseband signal, radiate a room full of your peers, and make sure baseband signals return out of both receive channels.

- Configure the acquisition computer to use a static IP address and write it down. Remote Desktop will be configured to connect to this IP address. If you do not do this, the computer will assign itself a random 169.x.x.x address that you cannot determine without connecting a monitor directly to the computer, and you won't necessarily have one in the hangar. Which brings us to...

C.2.2 In the Hangar

- Remove the rear cargo door, then place the supporting electronics in the rear of the aircraft and make all the power and coax connections. Only then should you install the fully-assembled radar door and make the final connections. The door can be installed single-handedly, but having an extra person helps.
- Do not open the rear cargo door once it is installed. This puts undue physical stress on the flexible waveguides when they push up against the fuselage.
- Test the full end-to-end RF path by simulating an acquisition. Alternatively, enable the function generator and run the `quick_fmcw_test` VI to check that everything is working.
- Push the aircraft out of the hangar and run `quick_gps_test` to make sure the COM port is correct and the GPS antenna is connected and working.

C.2.3 On the Tarmac

- Once the aircraft is started and idling, immediately boot the acquisition computer as soon as power is available. Only instruct the pilot to take off once the computer is fully booted and a LabVIEW acquisition has been started.
 - Sometimes the LabVIEW initialization gets stuck at “___”. In this case, starting a second instance of LabVIEW will be successful as long as the frozen instance remains in the background.
- Load the saved configuration into the Tektronix AFG3252 and immediately turn channel 1 output on.

C.3 Post-Flight Checklists

C.3.1 On the Tarmac

- When the aircraft engine stops, power will cut out abruptly. Make sure that the LabVIEW acquisition has been stopped before the aircraft shuts down to prevent corrupt or incomplete data sets. Ideally, stop the final acquisition immediately following touchdown.

C.3.2 In the Hangar

- Download data from AIMS computer to an external USB storage device.
- Disconnect connections to the transceiver and remove the system from the aircraft in the reverse order it was installed: remove the door, remove the supporting systems, then replace the stock door. Don't lose those cotter pins!

BIBLIOGRAPHY

- [1] W.R. Alpers, D.B. Ross, and C.L. Rufenach. On the detectability of ocean surface waves by real and synthetic aperture radar. *J. Geophys. Res.*, 86(C7):6481–6498, 1981.
- [2] J.B. Blair, D.L. Rabine, and M.A. Hofton. The laser vegetation imaging sensor: a medium-altitude, digitisation-only, airborne laser altimeter for mapping vegetation and topography. *ISPRS Journal of Photogrammetry and Remote Sensing*, 54(2):115–122, 1999.
- [3] J.C. Curlander and R.N. McDonough. *Synthetic Aperture Radar: Systems and Signal Processing*. New York: John Wiley & Sons, Inc, 1991.
- [4] T.G. Farr, P.A. Rosen, E. Caro, R. Crippen, R. Duren, S. Hensley, M. Kobrick, M. Paller, E. Rodriguez, L. Roth, et al. The shuttle radar topography mission. *Reviews of Geophysics*, 45(2):RG2004, 2007.
- [5] L.C. Graham. Synthetic interferometer radar for topographic mapping. *Proceedings of the IEEE*, 62(6):763–768, 1974.
- [6] S. Hensley, K. Wheeler, G. Sadowy, C. Jones, S. Shaffer, H. Zebker, T. Miller, B. Heavey, E. Chuang, R. Chao, et al. The uavsar instrument: Description and first results. In *Radar Conference, 2008. RADAR'08. IEEE*, pages 1–6. IEEE, 2008.
- [7] A. Meta, P. Hoogeboom, and L.P. Ligthart. Range non-linearities correction in fmcw sar. In *Geoscience and Remote Sensing Symposium, 2006. IGARSS 2006. IEEE International Conference on*, pages 403–406. IEEE, 2006.
- [8] A. Meta, P. Hoogeboom, and L.P. Ligthart. Signal processing for FMCW SAR. *IEEE Transactions on Geoscience and Remote Sensing*, 45(11):3519–3532, 2007.
- [9] T.L. Millette and C.D. Hayward. Detailed forest stand metrics taken from aims-1 sensor data. URL: <http://www.mtholyoke.edu/dept/earth/facilities/Millette-.pdf>, 2004.
- [10] A. Moreira, G. Krieger, I. Hajnsek, D. Hounam, M. Werner, S. Riegger, and E. Settelmeyer. Tandem-x: a terrasars-x add-on satellite for single-pass sar interferometry. In *Geoscience and Remote Sensing Symposium, 2004. IGARSS'04. Proceedings. 2004 IEEE International*, volume 2, pages 1000–1003. IEEE, 2004.
- [11] David M. Pozar. *Microwave Engineering*. Wiley, 4th edition, 2011.
- [12] Merrill Skolnik. *Introduction to Radar Systems*. McGraw-Hill, 3rd edition, 2002.
- [13] Warren L. Stutzman and Gary A. Thiele. *Antenna Theory and Design*. Wiley, 2013.
- [14] Anthony Swochak. Development, Deployment, and Characterization of a Ku-band Interferometer. Master's thesis, University of Massachusetts, September 2011.
- [15] H.K. Vedantham. Design and Development of a Ka-band Interferometer for Cryospheric Applications. Master's thesis, University of Massachusetts, 2009.
- [16] Karthik Srinivasan Venkatasubramanian. Design and development of timmi - an interferometric radar. Master's thesis, University of Massachusetts, September 2007.