Multi-Aperture Coherent Change Detection and Interferometry for Synthetic Aperture Radar

David D. Madsen

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

David G. Long, Chair
Brian D. Jeffs
Karl F. Warnick

Department of Electrical and Computer Engineering

Brigham Young University

April 2010

ABSTRACT

Multi-Aperture Coherent Change Detection and

Interferometry for Synthetic Aperture Radar

David D. Madsen

Department of Electrical and Computer Engineering

Master of Science

Interferometry and coherent change detection (CCD) utilize phase differences between complex SAR images to find terrain height and to detect small changes between images, respectively. A new method for improving interferometry and CCD using multiple sub-apertures is proposed. Using backprojection processing, multiple sub-aperture images are created for a pair of flights. An interferogram and coherence map is made from each sub-aperture. For CCD, each sub-aperture coherence map offers an independent estimate of the coherence over the same area. By combining coherence maps, low coherence areas associated with residual motion errors are reduced, shadowed areas are minimized, and the overall coherence of stationary objects between images is increased. For interferometry, combining independent estimates of a scene's height offers a more accurate height estimate. For repeat-pass interferometry, multiple apertures are shown to increase the coverage of valid height estimates. The benefits of multi-aperture interferometry and CCD are shown using examples with real data.

Keywords: SAR, synthetic aperture radar, image registration, coherent change detection, CCD, interferometry, multi-aperture, repeat-pass interferometry, L-Band, backprojection

ACKNOWLEDGMENTS

I would like to first thank Dr. David Long for his support in completing this thesis and his mentoring, which has helped me become a better engineer. I also thank Evan Zaugg and Matt Edwards, fellow SAR team members at BYU, for helping me learn about SAR. Thanks to the Space Dynamics Laboratory (SDL) and Naval Research Laboratory for use of NuSAR data. A special thanks to Chad Knight at SDL for his willingness and ability to modify SAR processing to help my work. I greatly thank my wife Laurie for her support.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1. INTRODUCTION

Synthetic aperture radar (SAR) is a versatile remote sensing instrument. The most basic use of SAR is imaging. SAR creates images by transmitting and receiving electromagnetic waves. Based on radar reflections alone, an image can be created day or night and through clouds or dust. Additional uses of SAR include finding moving targets, imaging through foliage, detecting man-made objects, identifying targets based on scattering properties, detecting small disturbances in a scene, mapping terrain, measuring surface deformation, and measuring surface movement. This thesis focuses on the use of radar for coherent change detection (finding small disturbances between two images) and interferometry (terrain mapping).

SAR produces complex images. Both the magnitude and the phase of the returned signal relative to the transmitted signal are recorded to produce these images. coherent change detection (CCD) and interferometry utilize phase differences between two complex images. For CCD, images of the same area are taken at different times and compared; changes to the area on the scale of the radar wavelength are detected. Interferometry uses the phase difference between two images of the same area taken from different angles. This phase difference is related to the height of the imaged area.

This thesis develops registration, coherent processing, and new methods of CCD and interferometry for a new radar system. The registration algorithm implemented aligns and warps two images so that the cross-correlation between the images is maximized. A list of coherent processing techniques using backprojection processing is given. These techniques ensure images formed from different flight paths are as similar as possible. An innovative method of utilizing multiple antenna sub-apertures to provide independent estimates of the coherence and height is introduced and applied to real data.

1.1 Past Work

This thesis builds on the scientific community's research into registration, CCD, and interferometry. Also, this work could not have been accomplished without the previous research and SAR development at Brigham Young University.

The correlation-based registration algorithm implemented in Chapter 3 is widely used in SAR processing. Specifically, [2] and [3] outline the algorithm and use it to register images before performing interferometry and CCD. Often, the cause of misregistration between images is errors in the recorded flight track. [4] and [5] have developed new methods of registration based on removing motion errors. These methods require that each flight be in interferometric mode and, therefore, do not work when only a single receive channel is used.

Past work with CCD is well documented with examples and theoretical development in [2] and [3]. There are several resources on collecting and processing highly coherent images. In [6], a theoretical development of the effect of flight path differences on the correlation between images in undertaken. Several practical suggestions and "rules-of-thumb" for image collection and processing are given in [7].

This thesis performs interferometry on images formed with backprojection processing. Interferometry using other image formation algorithms has been studied. For example in [2], interferometry is derived based on tomographic principles with approximations unique to spotlight-mode imaging. A specific source that traces the interferometric phase through backprojection was not found in the literature; however, the following resources are helpful: (1) A detailed description of the received SAR signal through range compression is given in [8] and [9]; (2) In [1] and [8], the backprojection integral is shown; and (3) A derivation of the interferometric phase given different assumptions about the ground elevation and antenna positions is found in [10]. This work contributes to a better understanding of interferometry with backprojection.

Raw data for this thesis was obtained by the NuSAR sensor. NuSAR is a pulsed radar system capable of transmitting at L-Band and X-Band with a bandwidth up to 500 MHz. In the primary imaging mode, NuSAR is designed to image wide swaths (3 km) from low altitudes (1000 to 1500 m). Development of the radar was a collaboration between Brigham

Young University (BYU), Artemis, Inc., Space Dynamics Laboratory (SDL) and the Naval Research Laboratory.

Processing of images is done with time-domain backprojection as implemented by BYU and SDL. The computational burden of the algorithm is efficiently managed by multi-threading the processing on a graphics processing unit (GPU). The accuracy and flexibility of backprojection allow for the use of arbitrary flight paths and wide-beam antennas. Sub-aperture formation and multi-aperture processing are easily implemented into backprojection.

Experience from previous SAR work at BYU also contributed to this thesis. Specifically, data from YINSAR is used. YINSAR is a low-cost, pulsed, interferometric radar with a 200 MHz bandwidth at X-band [11]. Between 1995 and 2003, YINSAR gathered many good interferometric data sets. Other radar experience at BYU included design and operation of the microSAR [12] and the microASAR [13].

1.2 Thesis Outline

This thesis is divided into the following chapters:

- Chapter 2 provides a background to SAR image characteristics. A theoretical study of sources of image decorrelation is also given.

- Chapter 3 describes the implementation of a correlation-based registration algorithm. Examples using real data show that registration removes noise from the interferogram and increases the coherence between images.

- Chapter 4 discusses coherence estimation. Steps are given to increase the coherence between images through careful processing. A new method of using multiple sub-aperture images to maximize the coherence between two passes is explained and demonstrated with data at both L-Band and X-Band.

- Chapter 5 develops the steps and assumptions made when creating a repeat-pass interferogram. By combining the height estimates from multiple sub-aperture interferograms a combined height estimate is formed that overcomes some of the drawbacks of

repeat-pass interferometry. An example of multi-aperture interferometry demonstrates this new method.

- Chapter 6 summarizes the results of the thesis and suggests avenues for future research into multi-aperture CCD and interferometry with backprojection.

CHAPTER 2. BACKGROUND

The algorithms for image registration, interferometry, and coherent change detection (CCD) all start with fully-processed, complex SAR images. Because of this, it is not necessary to understand all the details behind SAR image formation. Rather, this chapter pulls together only the relevant parts of SAR theory and provides them as background. The following sections include a basic introduction to SAR theory, scattering, correlation between images, image formation algorithms, and image collection.

2.1 Basic SAR Theory

A synthetic aperture radar produces a complex image of the area it flies over. A common configuration for radar operation is strip-map mode. In this mode, antennas are mounted on the side of a plane which then flies over an area and images a rectangular strip of ground to one side of the plane. Figure 2.1 defines some of the terms that are used to describe SAR geometry. The azimuth direction refers to the direction of flight. As the radar flies farther, the azimuth extent of the image increases. The range direction is perpendicular to the flight path.

As the radar flies over an area, it transmits a series of modulated pulses at a carrier frequency, f_0 [14],

$$x(t) = cos(2\pi f_0 t + \theta(t)). \tag{2.1}$$

The pulse modulation, $\theta(t)$, is most often a linear frequency modulation (LFM) that results in a SAR signal that ranges from f_0 to $f_0 + \beta$, where β is the bandwidth of the system. Each pulse reflects off everything illuminated by the antenna beam. The reflections that bounce back to the radar are received and recorded as raw radar data.

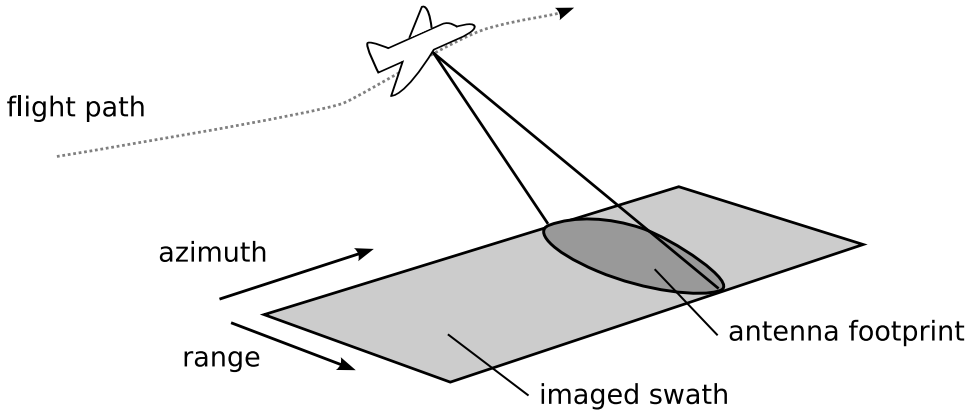

Figure 2.1: Simplified geometry for SAR image collection. Range is the distance from the radar perpendicular to the flight direction, and azimuth is the distance parallel to flight direction.

Image processing algorithms create an image from the raw radar data. The assumptions and methods used by different algorithms vary, but they all rely on knowledge of the exact phase and frequency of the transmitted signal relative to the expected and actual return signals. The raw data records the phase of the incoming reflections, and then the phase relation is retained through the imaging algorithms. A complex image results.

2.2 Scattering and Speckle

Scattering refers to the way in which an electromagnetic wave interacts with and reflects from an object. SAR creates images based on how the radar signal scatters off the objects in the imaged area. There are many differences between a radar image and an optical image because of the different scattering characteristics of microwaves compared to light. One of these characteristics is speckle. This section relates microwave scattering to speckle in SAR images and then discusses the dependence of speckle on observation aspect angle.

2.2.1 Speckle

Speckle refers to a seemingly random distribution of pixel values across a homogeneous area in a SAR image caused by coherent scattering from a pixel. The definition of a pixel

in a SAR image, the importance of radar operating frequency, and the coherent nature of scattering are covered.

Every pixel in an accurately processed SAR image corresponds to the scattering from a single patch on the ground known as a resolution element. In processing, the raw data from several pulses is used to create a pixel—each pulse having been transmitted and received at a different location. This means that the pixel representing a resolution element depends on the scattering characteristics of the element from several different angles. Scattering from a resolution element and the resulting pixel value depend on what is in the element, the radar frequency and the imaging geometry.

Both radar frequency and the size of scatterers in a resolution element affect the scattering characteristics of the element. A common rule-of-thumb in electromagnetics is that a wave interacts most with objects whose size is on the same order as the wavelength. Common operating frequencies for SAR systems are L, C, X, and Ku-band, which have wavelengths between 20 and 2 cm. For natural terrain (fields, bushes, trees, ground, etc.) a resolution element is largely composed of scatterers (branches, leaves, rocks, etc.) whose size is on the same scale as that of the radar wavelength. The scattering from a resolution element depends largely on the relative arrangement and/or movement of these scatterers. The roughness and slope of the ground across the resolution element also contribute to the scattering.

The scattering from a resolution element is the coherent sum of the scattering from all objects within the element. Figure 2.2 illustrates the return from a resolution element consisting of several point scatterers denoted as x's. If one of these scatterers is moved a few centimeters it will return a wave at a different phase. When this modified wave is added coherently to the other scattered waves it changes the phase and/or amplitude of the total scattered wave. The dependence of the total scattered wave on small changes causes speckle.

Figure 2.3 shows the speckle pattern from a field of grass. One would expect an optical image of a field of grass taken from far away to look very uniform from pixel to pixel—about the same shade and hue of green for each pixel. But for a radar image the brightness of a pixel depends on the position of each blade of grass and the shape of the

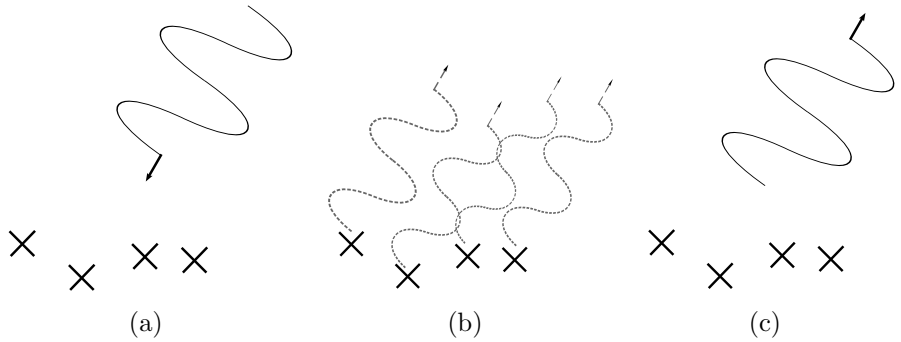

Figure 2.2: (a) Incoming wave from radar that hits a resolution element consisting of several dominant scatterers. (b) Return wave from each scatterer in the resolution element. (c) The coherent sum of the return waves becomes a single sinusoid at the same frequency as the incidence wave but with amplitude and phase that depend on the scatterers.

ground underneath the grass which vary from pixel to pixel. This results in bright and dark pixels across a field of similar looking grass.

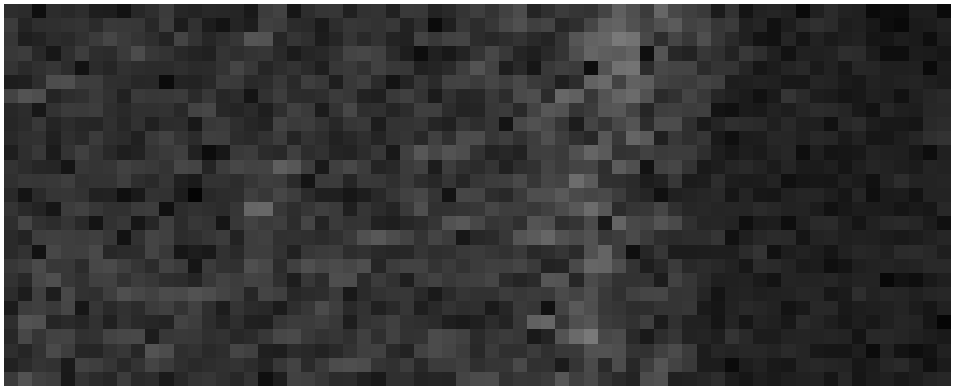

Figure 2.3: Speckle pattern over a field of grass at X-Band. On the right side of the image is a low SNR area caused by a flat body of water. The body of water shows that the noise level is below the average speckle value. The image covers 32 $\times$ 12 m at 0.5 m pixel resolution.

In creating SAR images, speckle is often seen as noise because it introduces an unpredictable high-frequency component into the image. However, speckle is repeatable; if an area is imaged twice using close to the same angle the same speckle pattern is received. When comparing two images of the same area taken with very similar flight paths, speckle is a deterministic signal which can be used to register the two images and perform CCD and interferometry [2].

2.2.2 Aspect Angle

This section further discusses how the scattering from a resolution element varies due to aspect (incidence or azimuth) angle. Figure 2.4 represents a range-cut through a resolution element. The scattering pattern represents the amplitude of the scattered wave from different look directions. At some angles the scattered fields cancel, creating nulls in the pattern; at other angles the fields combine to create lobes or peaks in the pattern. The two labeled radar look directions show how the same target can scatter very differently based only on the incidence angle of the illuminating wave.

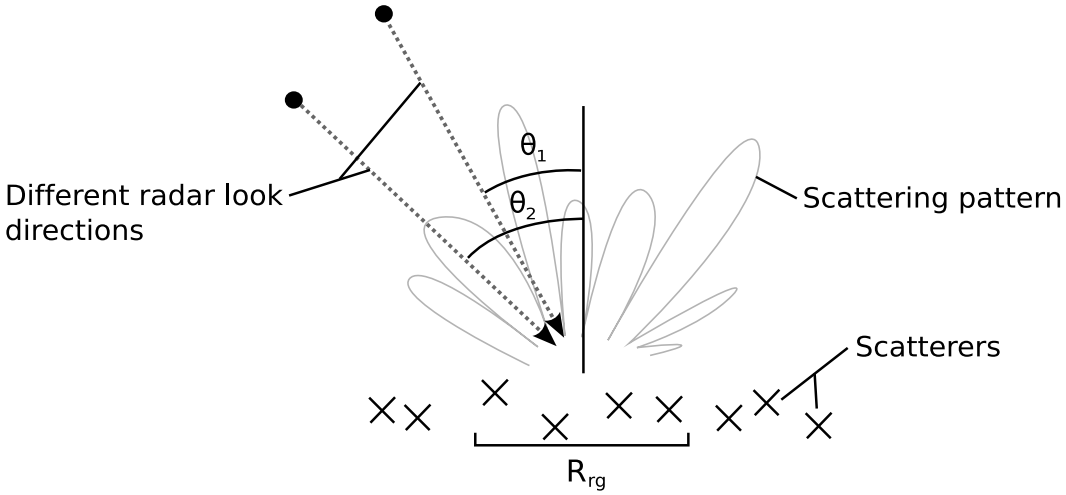

Figure 2.4: A simplified figure representing the scattering pattern of a resolution element, R_{rg}, in the range direction. The two labeled incidence angles, θ_1 and θ_2, demonstrate how the scattering amplitude quickly changes with incidence angle.

If the difference between the two angles, $\delta\theta = \theta_1 - \theta_2$, is sufficiently small then the scattering amplitude at both incidence angles is in the same lobe of the scattering pattern. The size of the lobes in the scattering pattern depends on the number of dominant scatterers in the resolution element. For a high resolution SAR, there are less scatterers in a resolution element resulting in broad lobes in the scattering pattern. Similar scattering is then observed for larger $\delta\theta$ than in a lower resolution system.

A similar scattering behavior exists in the azimuth direction. As the airplane flies past the resolution element, the element is illuminated from different angles. The total return

is the coherent addition of these returns at different angles (see Figure 2.5). The azimuth beamwidth of the antenna limits the size of the footprint and, therefore, the angles at which a resolution element is illuminated.

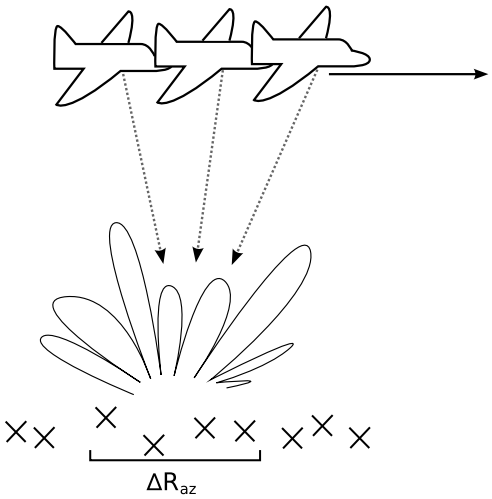

Figure 2.5: A representation of an azimuth cut of the scattering pattern of a resolution element. As the plane flies past, the radar illuminates and records the scattered field at several different times and angles.

For interferometry and CCD, similar scattering from the imaged area is required for both image collections. The speckle pattern in both images is then similar.

2.3 Correlation Between SAR Images

The similarity of the scattering and resulting speckle between two image collections is quantitatively measured as the cross-correlation. The amount of correlation is measured with the cross-correlation coefficient, γ. The range of γ is zero to one; when γ is unity the images are exactly correlated; when γ is zero there is no correlation. High correlation means that the two images are related by a highly linear relationship. This section mathematically relates differences in the imaging geometry to losses in the correlation between the images.

As described in [15] and [6] the correlation coefficient is composed of several separate sources of decorrelation. These are

$$\gamma = \gamma_{spatial}\gamma_{snr}\gamma_{volume}\gamma_{rotation}\gamma_{temporal} \tag{2.2}$$

where $\gamma_{spatial}$ is the decorrelation caused by a flight path that views the ground at a different incidence angle, γ_{snr} is caused by system noise in the radar, γ_{volume} is related to the vertical structure in a resolution element, $\gamma_{rotation}$ is caused by a rotation of a resolution element relative to the radar look angle, and $\gamma_{temporal}$ is the decorrelation caused by changes in the scene between repeat-passes.

2.3.1 Spatial Decorrelation

Spatial decorrelation is the largest obstacle to collecting coherent repeat-pass images. It is also known as baseline decorrelation because it is caused by the distance between the two antennas in interferometry, known as the baseline. For repeat-pass interferometry, spatial decorrelation restricts how far apart the two flight paths can be from each other.

Based on the geometry of the two antennas, it is possible to create an analytical model for the decorrelation between two images. This derivation follows those in [6] and [14]. Figure 2.6(a) shows the collection geometry for two images. It is assumed that the resolution element is composed of equal, evenly-distributed scatterers.

In Figure 2.6(b), the geometry has been rotated to make the following derivation clearer. A coordinate system has been assumed in which the range direction is now y and azimuth is x. The center of the resolution element is $y = 0$. The distance from the radar to a scattering point at y is

$$r_1' \approx r_1 + y \sin\theta_1. \tag{2.3}$$

This equation assumes that r_1 is large enough that the incident wave can be approximated as a plane wave by the time it reaches the scatterers.

A general expression for the value of a pixel at (x_0, y_0) is derived by convolving the return from the resolution cell by the SAR impulse response function. Using an incidence

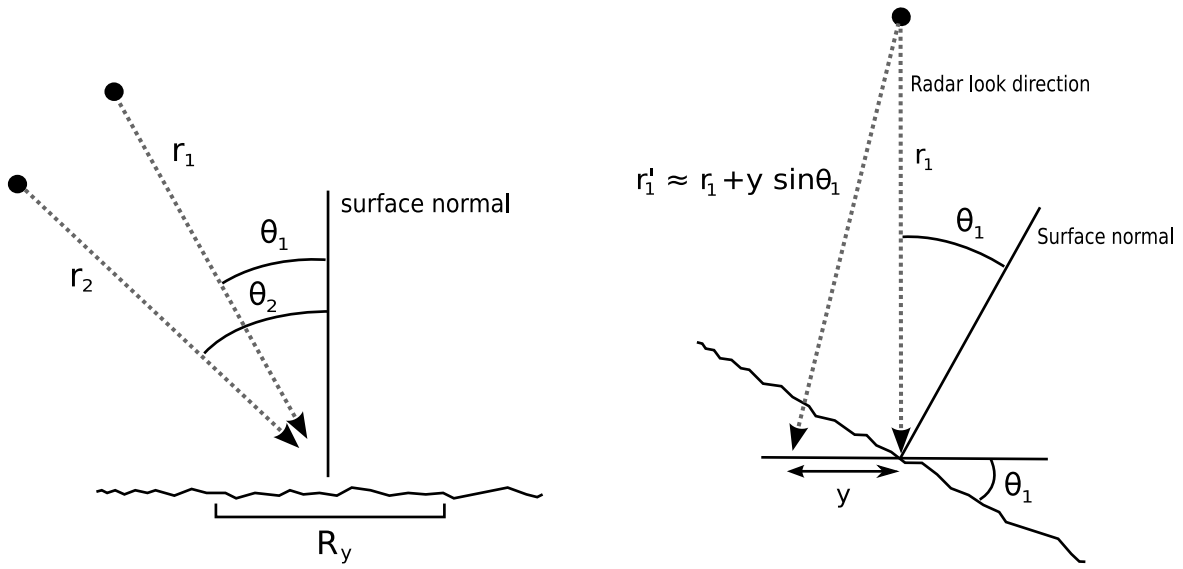

Figure 2.6: (a) Resolution element seen from two different angles representing two different flight paths. The ground is represented as a flat surface with evenly distributed roughness that causes backscatter. (b) A rotated geometry that shows the value of r_1, the distance to an arbitrary point in the resolution element, as a function of y.

plane wave given by $exp(j\frac{2\pi}{\lambda}r_1')$ the processed signal becomes

$$f_1 = \int\int \sigma(x_0 - x, y_0 - y) exp\left(-j\frac{4\pi}{\lambda}(r_1 + y\sin\theta_1)\right) W(x,y)dxdy. \tag{2.4}$$

In this equation, $\sigma(x,y)$ represents the complex backscatter of the ground; the exponential term is the additional phase caused by the integration point being a little closer to or further from the radar compared to the reference point at $y = 0$; and $W(x,y)$ is the radar impulse response, which is ideally a two-dimensional sinc function. The impulse response weights the signal from the ground close to the center of the resolution element, (x_0, y_0), so that it contributes most to the pixel.

The return from the same resolution element viewed at a different angle and distance, corresponding to a second image, is

$$f_2 = \int\int \sigma(x_0 - x, y_0 - y) exp\left(-j\frac{4\pi}{\lambda}(r_2 + y\sin\theta_2)\right) W(x,y)dxdy. \tag{2.5}$$

The only difference between Equation (2.4) and Equation (2.5) is found in the new phase term which now depends on the new range, r_2, and incidence angle, θ_2.

The cross-correlation between the two signals is

$$\begin{aligned} R_{f_1 f_2} &= \langle f_1^*(x,y), f_2(x,y) \rangle \\ &= \int\int f_1^*(x',y') f_2(x',y') \\ &= \int \sigma^*(x'-x_1, y'-y_1)\sigma(x'-x_2, y'-y_2) exp\left(-j\frac{4\pi}{\lambda}(r_1 - r_2)\right) \\ & exp\left(-j\frac{4\pi}{\lambda}(y_1 \sin\theta_1 - y_2 \sin\theta_2)\right) W^*(x_1,y_1)W(x_2,y_2)dx_1dy_1dx_2dy_2dx'dy'. \end{aligned} \tag{2.6}$$

The exponential term created by the difference between r_1 and r_2 is independent of the variables of integration so it can be removed from the integral. It changes the phase of the cross-correlation but can be ignored because it does not affect the magnitude.

By assuming that the surface, $\sigma(x,y)$, is composed of uniformly distributed scatterers that are uncorrelated, then

$$\langle \sigma^*(x,y)\sigma(x',y') \rangle = \sigma^0 \delta(x-x', y-y'). \tag{2.7}$$

This simplifies the cross-correlation to

$$\begin{aligned} R_{f_1 f_2} &= \int\int\int\int \sigma^0 \delta(x_1 - x_2, y_1 - y_2) exp\left(-j\frac{4\pi}{\lambda}(y_1 \sin\theta_1 - y_2 \sin\theta_2)\right) \\ & W^*(x_1,y_1)W(x_2,y_2)dx_1dy_1dx_2dy_2 \\ &= \int\int \sigma^0 \exp\left(-j\frac{4\pi}{\lambda} y(\sin\theta_1 - \sin\theta_2)\right) |W(x,y)|^2 . \end{aligned} \tag{2.8}$$

The equation has become a Fourier relation. Using a radar impulse response of $W(x,y) = \text{sinc}(x/R_x)\,\text{sinc}(y/R_y)$, evaluating the integral, and normalizing the result yields an expression for the decorrelation caused by viewing the target from two slightly different angles,

$$\gamma_{spatial} = 1 - \frac{2\cos\theta\,|\delta\theta|\,R_y}{\lambda}, \tag{2.9}$$

where θ is the average incidence angle, $\theta = \frac{\theta_1+\theta_2}{2}$, and $\delta\theta = \theta_1 - \theta_2$. This equation offers several insights into what causes decorrelation. Obviously, a large difference in incidence angles causes greater decorrelation. Somewhat less obvious is how the average incidence angle affects the correlation. Targets imaged closer to the radar are less correlated than those further away. This equation also shows that a high resolution system at a low frequency is less susceptible to spatial decorrelation.

It is possible to express Equation (2.9) in terms of the baseline. Given that the distance between planes is B, the equation becomes

$$\gamma_{spatial} = 1 - \frac{2\,|B|\,R_y \cos^2\theta}{\lambda r}. \tag{2.10}$$

The baseline is assumed to be in the horizontal plane.

2.3.2 SNR Decorrelation

The noise from two different receivers on a single-pass interferometer or from one receiver at two different time periods on a repeat-pass sensor is uncorrelated. Generally, the signal is significantly stronger than the noise. There are, however, several common situations where the signal is weak. One situation is when the area is a smooth surface that reflects the radar energy away with little backscatter. Bodies of water and asphalt roads are such areas. Another situation is in areas of shadow. Trees, structures and terrain prevent the radar signal from reaching parts of the image. The far range is also an area that has low SNR. This is caused by spreading losses associated with imaging far distances and because the incidence angle becomes shallower which produces less backscatter.

Areas of low SNR show up as dark areas in a SAR image. For CCD, it is not possible to detect changes in the low SNR areas, and for interferometry it is not possible to extract the terrain if the SNR is too low. In order to prevent low SNR areas from influencing CCD and interferometry, the areas are masked out and ignored.

2.3.3 Volume Decorrelation

The volume correlation coefficient, γ_{volume}, is related to the amount of volume scattering that takes place in a resolution element. If there are trees or large bushes volume scattering is more significant. As shown in Figure 2.7, γ_{volume} takes into account scatterers at various heights within a resolution cell. Both volume and spatial correlation are caused by a change in incidence angle, but spatial correlation only accounts for a relatively flat scattering surface.

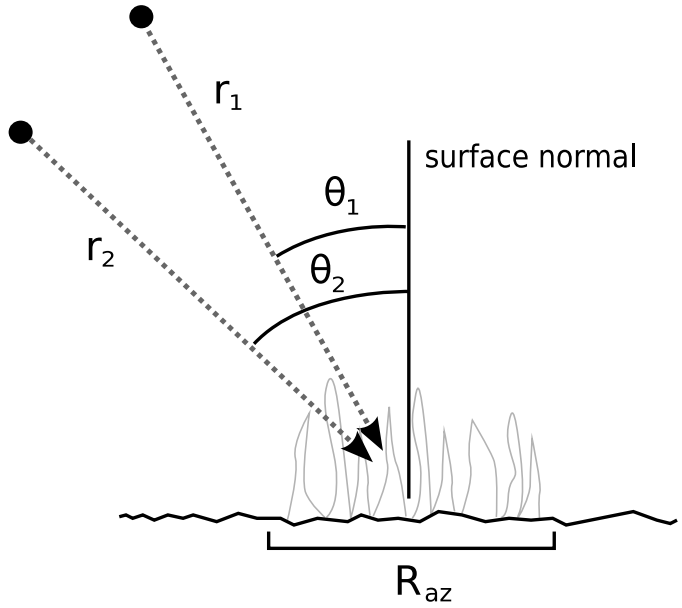

Figure 2.7: The geometry from flying over a resolution element with volume scatterers at two different incidence angles.

For many cases, volume scattering is not significant—flat terrain being one such case. Even for large volume scatters such as trees, the temporal decorrelation caused by moving leaves and branches between passes is larger than any volume decorrelation. For urban settings, volume scattering is not as significant as problems caused by layover. For this study, most images were taken over desert terrain. This results in little volume decorrelation because most of the scattering was from the ground or objects low to the ground.

2.3.4 Rotational Decorrelation

Rotational decorrelation is caused by flying past a target at a different azimuth angle as shown in Figure 2.8. For a repeat-pass collection, rotational decorrelation is present when the two flight paths are not exactly parallel. The decorrelation caused by the distance between the flights is accounted for by the spatial coefficient. Rotational decorrelation accounts for a changing distance between the target and radar over the imaging aperture.

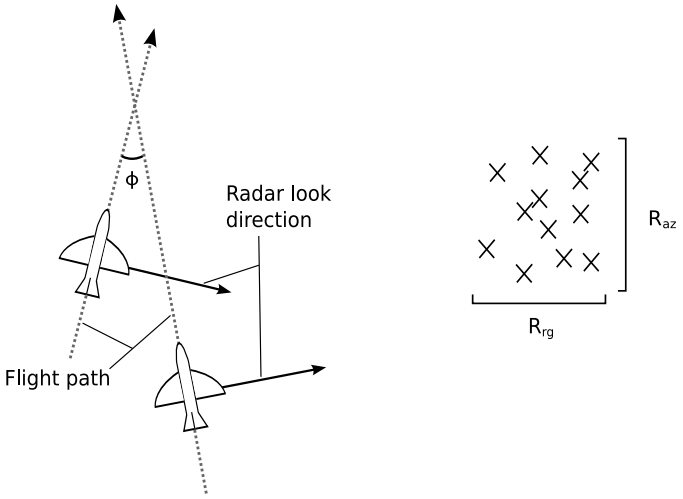

Figure 2.8: Two flights seen from above that pass the target at slightly different angles. The distance from the resolution element is changing at different rates between the two passes.

In a single-pass collection, decorrelation is caused by changing the crab angle of the plane. The crab angle is the difference between the direction the plane is pointed and the direction it is flying. This is common for small planes when flying in high winds. Note that it is not the magnitude of the crab angle but a change in crab angle that causes decorrelation.

Using a similar analysis as presented in Section 2.3.1, the rotation correlation coefficient is found by [6] to be

$$\gamma_{rotation} = 1 - \frac{2 \sin\theta \, |\delta\phi| \, R_x}{\lambda}, \tag{2.11}$$

where ϕ is the angle between the two flight tracks. This equation is verified in [6] by modeling a resolution element with several scatterers. The resolution element is then rotated

at different angles. The results show that at L-band, rotations of as much as 1° drop the correlation to 0.8 while at C-band the returns become completely uncorrelated after about 0.7°.

2.3.5 Temporal Decorrelation

Changes in the imaged area between two passes cause temporal decorrelation. Temporal changes are often caused by natural movement such as grass and leaves moving in the wind. Also, man-made changes such as walking across grass, driving over a gravel road, or digging up an area cause decorrelation. CCD often seeks to identify the latter type of changes. The natural changes introduce undesired decorrelation in the scene creating false CCD targets. For repeat-pass interferometry, both kinds of changes would be undesirable and lead to inaccuracies in topographic estimation.

Most of the sources of decorrelation listed in this section are dependent on the flight path (spatial, volume, rotational). If highly correlated images from two flights are needed the paths must be close enough so that decorrelation from geometric differences is minimal. After data has been recorded from two close flights, accurate imaging algorithms are still necessary to maintain the correlation inherent in the data.

2.4 Imaging Algorithms

SAR image formation algorithms recover the complex reflectivity of a resolution element with minimal distortion and/or artifacts. Care must be taken to ensure that not only the magnitude, but also the phase, is correct. For most of the images in this thesis the backprojection algorithm is used. This section gives a brief description of this algorithm.

Backprojection (BPJ) is a time-domain algorithm. In contrast to frequency-domain algorithms such as the range-Doppler algorithm and chirp scaling algorithm, time-domain algorithms do not make any assumptions in order to allow processing in the frequency domain. Backprojection is a tomographic reconstruction method similar to that used for CT (computed tomography) scans in the medical imaging field. The application to SAR processing

is presented in [2] and [8]. In [1], the backprojection integral is given as,

$$f(\bar{r}_0) = \int g(x', R) exp(j2kR) dx', \tag{2.12}$$

where x' is the along track variable of integration; k is the wavenumber of the carrier frequency; $R = |\bar{r}(x') - \bar{r}|$ is the distance from the plane's position, $\bar{r}$, at x' to the resolution element's position, $\bar{r}_0$. Figure 2.9 shows the SAR imaging geometry labeled with the terms from Equation 2.12.

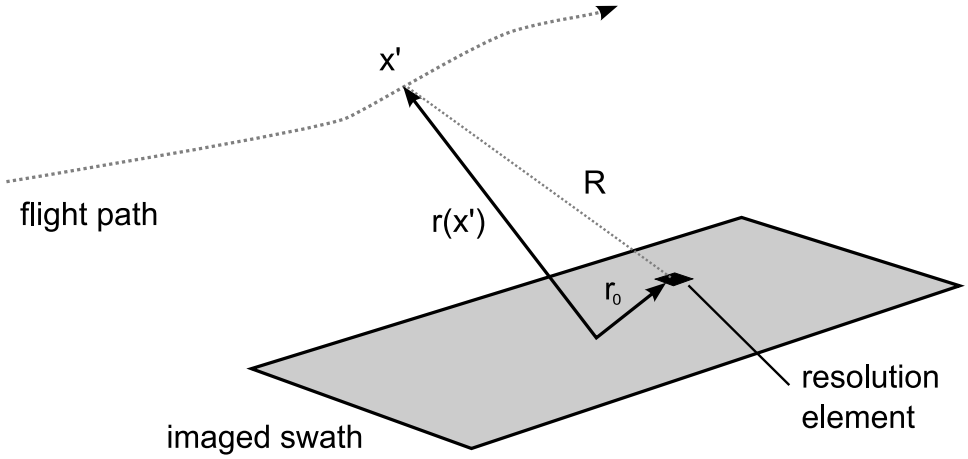

Figure 2.9: The geometry for backprojection depends on knowledge of the range, R, between the resolution element and the plane at each pulse.

The backprojection integral is implemented by summing each pulse weighted by the phase expected for a target at distance R. R changes as the plane moves between pulses. The expected phase is a matched filter that causes the returns from a target at $\bar{r}_0$ to coherently sum while the returns from all other targets incoherently sum resulting in cancellation. The matched filter is carried out for each pixel in the image. The number of computations for an image depends on the number of pixels in an image multiplied by the number of pulses within the limits of integration of the backprojection integral, also known as the aperture. Specifically, the range from each pulse to each pixel must be calculated. Figure 2.10(a) shows the aperture and ground split into the discrete sections corresponding to each pulse

and resolution element. Given that there are L pulses in the aperture and $M \times N$ pixels, the computation time is proportional to $L \times M \times N$.

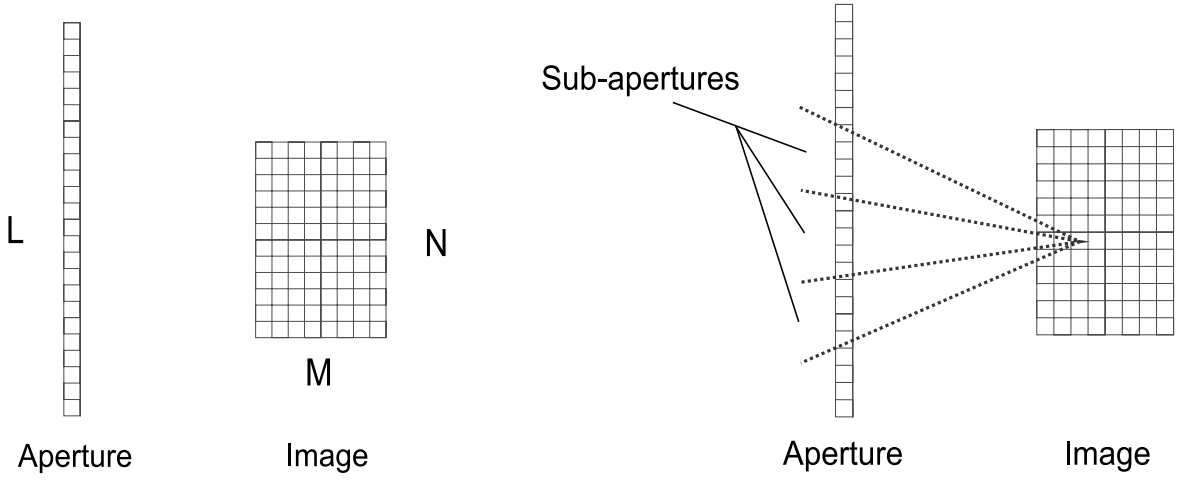

Figure 2.10: (a) Illustration of the flight line divided into aperture of L pulses and the ground divided into $M \times N$ resolution elements. Adapted from [1]. (b) The aperture divided into sub-apertures relative to one pixel in the image.

The computational burden of backprojection is much higher than that of frequency domain algorithms. However, backprojection offers several advantages. One advantage is that no assumptions are made about the flight track. Other algorithms assume a linear track and any deviation from this ideal track must be compensated for. As long as the location of the antenna at each pulse is accurately known, backprojection produces accurate images from any flight track. Another advantage is that backprojection automatically forms an image projected on the ground model. The ground can be modeled as a flat surface, or a digital elevation model (DEM) of the surface can be used. Frequency-domain algorithms initially produce an image in the slant-range which must be projected into the ground range (see Figure 2.11).

Frequency-domain algorithms also make several approximations about the antenna beamwidth and pointing direction. Specifically, the antenna must be narrow beam and mounted broadside relative to the flight direction. Backprojection is capable of making images with a wide beam antenna and with significant squint. When a wide beam is used,

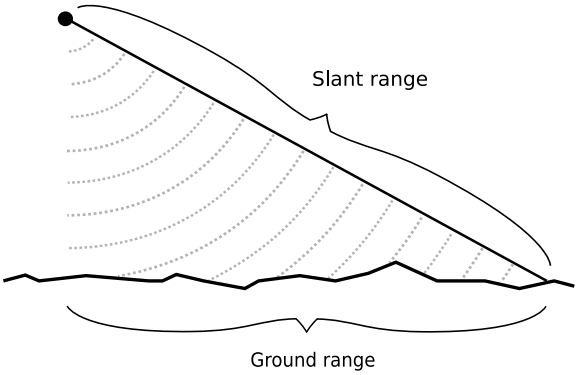

Figure 2.11: Slant range corresponds to the direct distance from the radar. Ground range is the distance from the radar measured along the surface of the earth.

the aperture can be restricted to define a narrower sub-aperture. This allows a wide beam antenna to be split into multiple sub-apertures, each sub-aperture producing an independent image of the same area. Figure 2.10(b) shows how the pulses in an aperture are split based on different ranges of angles for a given resolution element. Chapters 5 and 4 present work that uses multiple apertures derived from a wide-beam antenna.

2.5 Image Collection

The ideas of correlation between images and accurate processing discussed in the previous sections put requirements on how the images are collected. This section discusses how these requirements drive practical image collection and processing.

Both interferometry and CCD require highly-correlated images; however, identical images would not be useful for either. CCD finds changes in the phase, which correspond to changes in the imaged area, using a temporal baseline. Interferometry uses a physical baseline so that the two images have phase differences related to elevation. In both cases, a trade-off between having an appropriate baseline and maintaining correlation between images must be met.

For CCD, two images of an area that only differ by the changes of interest are desired. Ideally, the two passes over the area have the exact same flight paths—removing any spatial, rotational, or volume decorrelation. The temporal baseline must be chosen to ensure that the desired changes are detected while not allowing too many undesired changes to occur

(e.g. natural movement in the scene). For an airborne radar, the second flight path is never the same as the first flight. In practice, a tube is defined within which both flight paths should be contained to ensure correlated images. The size of the tube is chosen to reduce the spatial decorrelation between flights as given in Equation 2.9.

For interferometry, a large physical baseline gives more accurate height measurements [16]. Large baselines resulting from using two passes make it more likely that the baseline will vary, causing decorrelation between images. The ideal geometry is to have a constant baseline between passes that is a fraction of the baseline distance which causes complete decorrelation [16]. Similar to the tube required for CCD, for practical repeat-pass interferometry, often two tubes are defined around each path with a fixed baseline between the center of each tube [15]. In addition, repeat-pass interferometry requires accommodations for the changing baseline in processing to produce accurate height estimates.

Given appropriate geometries, the image must then be accurately processed. For backprojection, an accurate knowledge of the range between the antenna and each resolution element on the ground is needed. The antenna's position is obtained with highly accurate GPS and IMU (inertial measurement unit) sensors. The motion data collected by these sensors is processed to give a smooth accurate estimate of the antenna's motion during each pass. In practice, there are still small errors between the true path and the estimated path, known as residual motion errors, which reduce the correlation between images. These are further discussed in Chapter 4. The location of each resolution element is estimated from digital elevation models (DEM). DEMs provide rough elevation values which ensure the geolocation and focusing of an image. Interferometry is able to improve upon the elevation values by extracting the height difference between the actual terrain and the DEM as discussed in Chapter 5.

CHAPTER 3. REGISTRATION

It is necessary to align two SAR images for interferometry and/or coherent change detection (CCD). This process is known as registration. This chapter discusses the need for image registration. A registration algorithm is presented. Results are shown that demonstrate the performance of the algorithm in different settings and with different radar systems.

3.1 Registration and Correlation

In a registered pair of complex radar images, each pixel in one image directly corresponds to a pixel in the second image. Both pixels also directly correspond to the same area on the ground. In referencing the images, one image is designated as the master image and the second image is the slave image. When there is not direct pixel correspondence between images, the master and slave images are misaligned or misregistered. Several different types of misalignment are illustrated in Figure 3.1. In addition to pixel-by-pixel alignment, each pixel should be imaged with close to the same geometry to ensure correlation between images. When both of these requirements are met, the two images are registered and highly correlated.

Two registered images line up exactly with each other regardless of the different flight paths and different processing of the images. There are two ways this can be accomplished. First, the image could be formed by taking into account the motion of the plane and the elevation of the ground. For example, the backprojection algorithm uses flight motion data with centimeter accuracy and digital elevation models (DEMs) to create an image in which each pixel matches a specific patch of ground. Motion compensation and slant-range to ground-range mapping can yield similar results for other algorithms. However, these algorithms are only as accurate as the data on which they are based. Small errors in the GPS and motion data, known as residual motion errors, result in inaccurate images. Residual

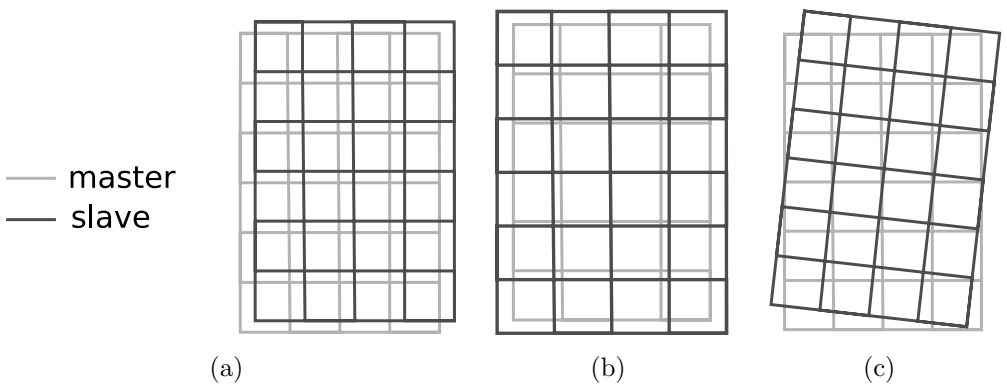

Figure 3.1: Possible registration errors between pixels in a master and slave image: (a) Translational (b) Scaling (c) Rotational. Most registration errors can be expressed as a combination of these basic errors.

motion errors cause pixel misalignment, as shown in Figure 3.1. Additionally, they cause image defocussing which results in objects blurring across pixels.

The second way to register images is with image-based registration algorithms, as presented in this chapter. These algorithms are used on fully-processed, complex images. They line up the images based on features and/or statistical properties common to the two images. Frequently, both accurate, motion-based processing and registration algorithms are used to achieve the best registration between images.

Correlation is a statistical measure of the similarity between two images. As discussed in Section 2.3, imaging from different geometries results in slightly different backscatter characteristics for the same area. When two images are highly correlated, matching pixels have almost the same complex value except for a possible slight phase difference. Interferometry and CCD are based on this phase difference and, therefore, require that the images be highly correlated.

Registration and correlation are tied together in many cases. For example, although two images can be registered with accurate processing from very different flight paths, the two images may not be highly correlated. Conversely, images can be taken from very similar flight paths, yet the correlation between them is not maximized until they are registered pixel-by-pixel. For best results, highly correlated data is first collected, then accurate processing and

image-based registration algorithms are used to simultaneously register and better correlate two images.

3.2 Registration Algorithm

The following algorithm is based on the algorithms described in [2] and [3]. It is an automatic algorithm which aligns images by maximizing the cross-correlation without any user input. A simplified outline of the method is presented below, followed by an in-depth explanation of each step in the succeeding subsections.

1. Find a global shift of the slave image that maximizes the cross-correlation between the two images. Shift the slave image by that amount.

2. Split the images into blocks. Find a shift of the slave image block that maximizes the cross-correlation for the block. Apply this shift to the slave block.

3. Warp each slave block to the master with the following steps, using only the magnitude for each pixel in the block.

   (a) Create a list of control points by dividing the block into sections and finding a subpixel shift that maximizes the cross-correlation of each section.

   (b) Use a spline warping method to find a shift for each pixel based on the shifts of the control points.

   (c) Interpolate the pixels back onto an evenly-spaced, rectangular grid.

4. Repeat step 3 using the complex value of each pixel.

3.2.1 Calculating a Global Shift Based on Cross-correlation

For this step, it is assumed that the two images are largely the same but they are shifted relative to each other, i.e., a translational misalignment. The magnitude of this shift, $(\Delta x, \Delta y)$, is found and applied to one of the images. In the ideal case, the two complex

images can be expressed using magnitude and phase components as

$$f_1(x,y) = r(x,y)e^{j\phi(x,y)}, \tag{3.1}$$

$$f_2(x,y) = f_1(x+\Delta x, y+\Delta y) = r(x+\Delta x, y+\Delta y)e^{j\phi(x+\Delta x, y+\Delta y)}, \tag{3.2}$$

where f_1 is the master image and the slave image, f_2, is a translated copy of the master image. The cross-correlation between the images is

$$R_{f_1 f_2}(s,t) = E\left(\sum_x \sum_y f_1(x,y) f_2^*(x-s, y-t)\right) \tag{3.3}$$

$$= E\left(\sum_x \sum_y f_1(x,y) f_1^*(x-s+\Delta x, y-t+\Delta y)\right) \tag{3.4}$$

$$= R_{f_1 f_1}(s-\Delta x, t-\Delta y). \tag{3.5}$$

Assuming wide-sense stationarity and uncorrelated pixels, the correlation is ideally a delta function. This is a good model over homogeneous terrain, but may be inaccurate for an urban or more structured environment. The correlation is expressed as:

$$R_{f_1 f_2}(s,t) = R_{f_1 f_1}(s-\Delta x, t-\Delta y) \approx \delta(s-\Delta x, t-\Delta y). \tag{3.6}$$

For ideal images, the cross-correlation, $R_{f_1 f_2}$, has a maximum at $(s,t) = (\Delta x, \Delta y)$ and is zero for any other shift. In practice, the value of $R_{f_1 f_2}$ is small for most shifts $(s,t) \neq (\Delta x, \Delta y)$ and has a large peak at $R_{f_1 f_2}(\Delta x, \Delta y)$. The desired shift is accurately found by locating the maximum of the cross-correlation function.

Computationally, the cross-correlation of two large images or sections of images can be calculated efficiently by taking advantage of frequency domain relations. This is done by estimating the expectation given in Equation 3.3 with the sample correlation, which is a

convolution:

$$R_{f_1 f_2}(s,t) = \sum_x \sum_y f_1(x,y) f_2^*(x-s, y-t) \tag{3.7}$$

$$= f_1(x,y) * f_2^*(-x,-y). \tag{3.8}$$

Given the Fourier relations

$$\begin{aligned} F_1(m,n) &= \mathcal{F}\{f_1(x,y)\}, \\ F_2(m,n) &= \mathcal{F}\{f_2(x,y)\}, \end{aligned} \tag{3.9}$$

then the convolution property of Fourier transforms can be used,

$$R_{f_1 f_2}(s,t) = f_1(x,y) * f_2^*(-x,-y) \tag{3.10}$$

$$= \mathcal{F}^{-1}\{F_1(m,n) F_2^*(m,n)\}. \tag{3.11}$$

Figure 3.2 shows a three-dimensional plot of the cross-correlation for two sample images. Note that there is a clear peak centered a little off the origin. The displacement from the origin corresponds to the displacement of the slave image relative to the master image. Shifting the slave image by this amount performs a rough global registration of the two images.

For large images, calculating the cross-correlation is computationally expensive. In this case, the image is broken into various subimages that are smaller than the total image but larger than the blocks. Another option is to neglect a global shift and rely on the block level shifts of the next step.

3.2.2 Splitting into Blocks and Shifting

One global alignment is usually not enough to line up two images. Because of the changing flight paths the images are misregistered in different ways at different locations in the image. To remedy this problem the images are broken into blocks of the master image, $g_1(x,y) \subset f_1(x,y)$, and blocks of the slave image, $g_2(x,y) \subset f_2(x,y)$. For each block pair

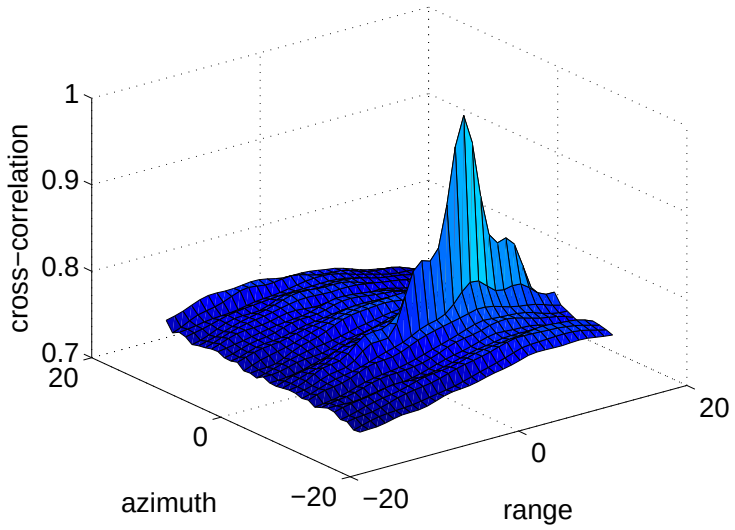

Figure 3.2: The normalized cross-correlation between two images as a function of shifting one image in two dimensions relative to the other image. The peak shows the optimal shifts to line up the two images.

a shift is found that maximizes the cross-correlation and this shift is applied to the slave block.

The size of a block must meet several different criteria. Because the cross-correlation is performed in the frequency domain, block sizes that are powers of two are efficient. The block size should be large enough to accommodate many sections so that the next step can produce a smooth warping function. At the same time, blocks should be kept small because the warping function becomes difficult to compute for many sections and must be applied to every pixel in the block. In practice, blocks on the order of 512×512 or 1024×1024 pixels are used.

3.2.3 Splitting into Sections and Finding Control Points

In the next step of the algorithm, each block is divided into sections. The following criteria should be taken into account when choosing the section size. It is important to make sure that sections do not get too small. A sufficiently large section ensures a high, tight peak in the cross-correlation function; whereas, small sections may have short peaks that

are difficult to detect in noise. A certain number of sections per block is needed to form the warping function. This requirement is discussed in Section 3.2.4.

Once an appropriate section size has been found, the next step is to determine the displacement between images. This is done using the same cross-correlation methods used for the global and block shifts. The displacements for each section are used to create a list of control points. A control point consists of a location, (x, y), in the master image and a second location, $(x', y') = (x+\Delta x, y+\Delta y)$, in the slave image. The master location, (x, y), is a pixel in the middle of the section being tested. Several additional steps beyond those used for global correlation are taken to ensure accurate displacements and to rate the strength of the control points.

The misregistration between images may not always be an integer number of pixels. Subpixel misregistrations can be obtained by interpolating the correlation matrix. Figure 3.3 shows the original correlation and an interpolated correlation map. The interpolation reveals structure in the peak of the correlation map so that a more precise displacement can be measured. Shifts of approximately 0.05 pixels can be achieved in this way [17].

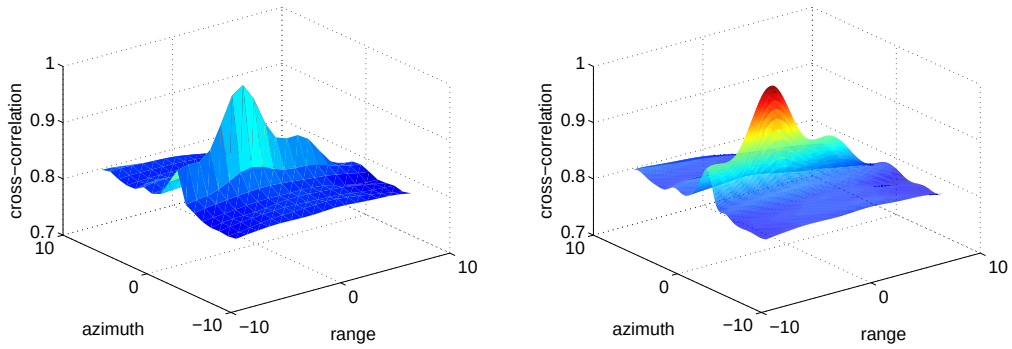

(a) (b)

Figure 3.3: (a) Correlation function formed from whole pixel shifts in range and azimuth. (b) Correlation function for subpixel shifts obtained by interpolating between values of the whole pixel correlation map.

Parts of the image with low SNR or significant differences in collection geometry may not have a single strong peak. In this case, it is best not to add the displacement of a section in this part of the image to the list of control points. In [2], a threshold is used test each

control point as given by

$$\Lambda = \frac{\max_{s,t} |R_{f_1 f_2}(s,t)|}{\sqrt{\sum_s \sum_t |R_{f_1 f_2}(s,t)|^2}}. \tag{3.12}$$

This measures the ratio of the peak to the average value of the correlation map.

In registering images for this thesis, it has been found that two threshold levels are effective. First, control points with a peak-to-rms value over Λ_1 are added to the list of control points. If less than ten control points are added to the list the top ten control points by peak-to-rms value, regardless of the first threshold, Λ_1, are chosen. Then a second threshold, Λ_2, is used. All control points in the list must be above this threshold in order to proceed with warping. Therefore, if ten points above Λ_2 cannot be found then no alignment of the images is attempted. This method of using two threshold provides flexibility in registering images which have different levels of correlation in different areas.

3.2.4 Applying a Warp

Once a list of control points has been gathered, the next step is to find displacements for every pixel. This is done by applying a thin-plate spline warp as described in [18] and [19].

Finding the displacements for each pixel is laid out in [19]. Given a list of n control points, each composed of an old coordinate and a new coordinate,

$$\{(x_i, y_i), (x_i', y_i') : i = 1...n\}, \tag{3.13}$$

it is desired to find two smooth mapping functions from the old coordinates to the new coordinates,

$$x_i' = f(x_i, y_i), \tag{3.14}$$

$$y_i' = g(x_i, y_i). \tag{3.15}$$

After f and g have been found for the control points, any other coordinates can be mapped through f and g to new coordinates.

Figure 3.4 shows how the problem applies to pixels in an image. Control points located throughout the image are shifted in the x and y directions to their new locations. Functions are found to fit the control point shifts, and then these function are used to find new coordinates for each pixel.

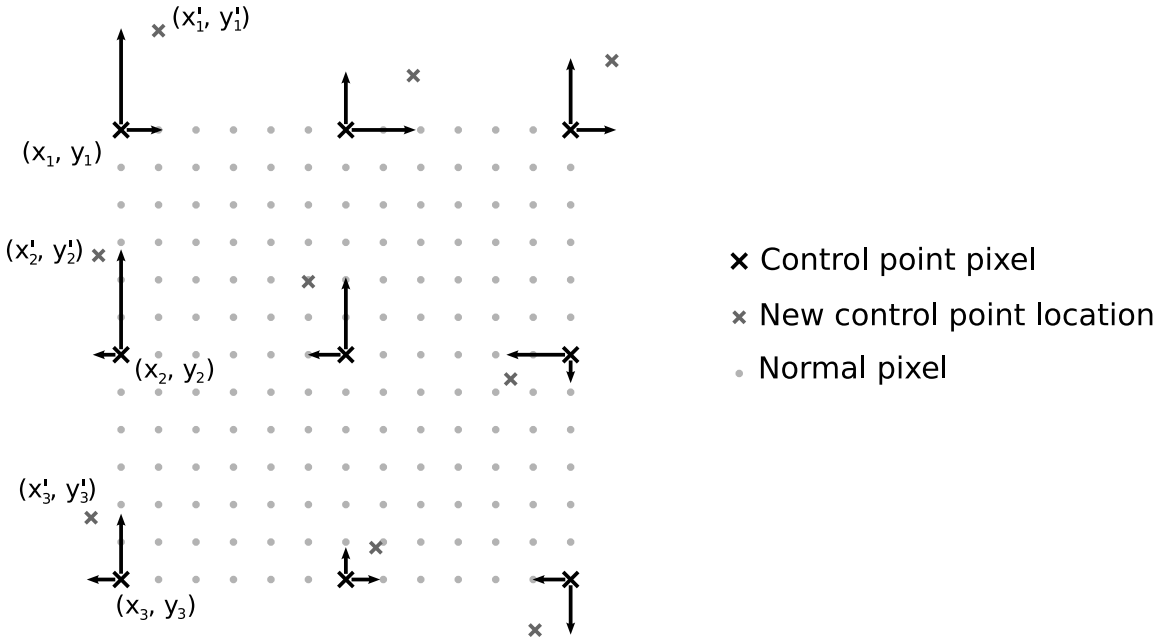

Figure 3.4: The warp algorithm finds mapping functions that map each control point (x_i, y_i) to (x'_i, y'_i). Every pixel is then mapped to a new location using these functions.

In [19], the finding smooth mapping function is solved as a surface fitting problem. The problem stated in surface fitting terms is as follows. Given a list of points in three dimensions,

$$\{(x_i, y_i, x'_i) : i = 1...n\},$$
$$\{(x_i, y_i, y'_i) : i = 1...n\}, \tag{3.16}$$

find functions f and g which define two surfaces according to,

$$x'_i = f(x_i, y_i),$$
$$y'_i = g(x_i, y_i), \tag{3.17}$$

where the surface passes through each control point. Figure 3.5 illustrates the problem expressed using surface fitting.

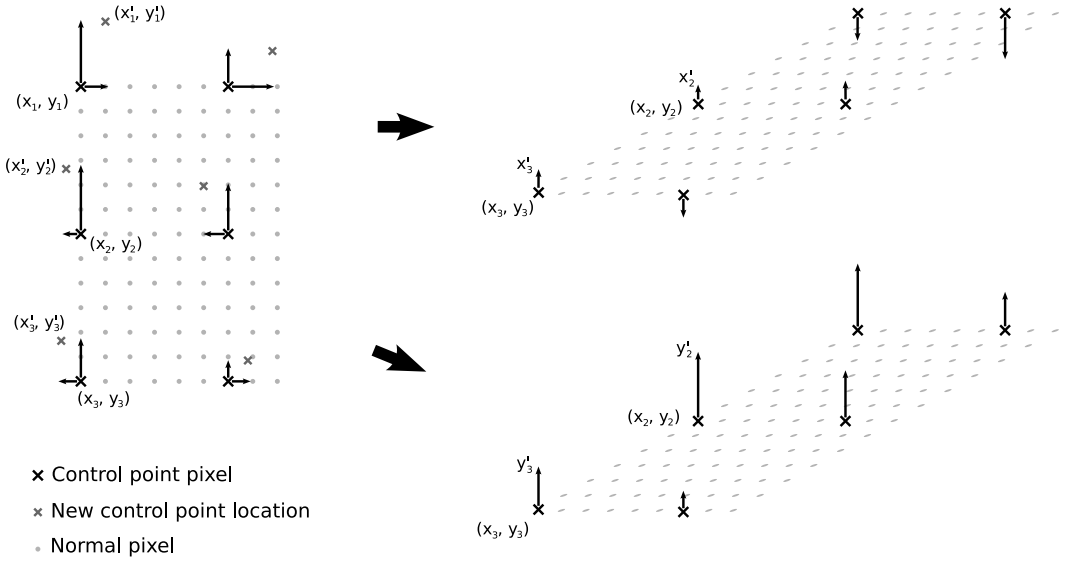

Figure 3.5: Warping problem expressed as a surface fitting problem. The new coordinates of the control points, x'_i and y'_i, form two surfaces. The surface heights in the picture are not to scale.

There are many surfaces that satisfy the surface fitting problem. Thin-plate splines is a method used to model a thin metal sheet that is pinned to heights in several places—taking into account the bending energy of the plate. There is a minimum energy constraint in this model that makes it well suited to fitting a smooth surface through several control points [18]. This surface is formed by interpolating values between the control points. The interpolant surface is

$$z(x, y) = -r^2 \ln r^2, \tag{3.18}$$

where $r = \sqrt{x^2 + y^2}$. By applying this base surface to each control point and adding each contribution to form an overall surface, the following equation is reached:

$$f(x_1, y_1) = a_0 + a_1 x_1 + a_2 y_1 + \sum_{i=1}^{n} b_i r_{1i}^2 \ln r_{1i}^2, \tag{3.19}$$

where $r_{1i} = \sqrt{(x_1 - x_i)^2 + (y_1 - y_i)^2}$ is the distance from each control point, (x_i, y_i), to the point being evaluated, (x_1, y_1).

Using the surface given by Equation (3.19) and the control points in Equation (3.17), a system of linear equations can be constructed,

$$\begin{aligned}
f(x_1, y_1) &= a_0 + a_1 x_1 + a_2 y_1 + \sum_{i=1}^{n} b_i r_{1i}^2 \ln r_{1i}^2 = 0, \\
&\vdots \\
f(x_n, y_n) &= a_0 + a_1 x_n + a_2 y_n + \sum_{i=1}^{n} b_i r_{ni}^2 \ln r_{ni}^2, \sum_{i=1}^{n} b_i = 0, \\
\sum_{i=1}^{n} x_i b_i &= 0, \\
\sum_{i=1}^{n} y_i b_i &= 0.
\end{aligned} \tag{3.20}$$

The last three equations guarantee that the weights, $\mathbf{b} = [b_1 ... b_n]^T$, and the cross-product of the weights with the control points are zero. This enforces the minimum bending energy requirement [19].

These equations can be expressed as a matrix. An intermediate step between the linear equations of Equation 3.20 and the matrix equation is

$$\begin{bmatrix} f(x_1,y_1) \\ \vdots \\ f(x_n,y_n) \\ 0 \\ 0 \\ 0 \end{bmatrix} = a_0 \begin{bmatrix} 1 \\ \vdots \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} + a_1 \begin{bmatrix} x_1 \\ \vdots \\ x_n \\ 0 \\ 0 \\ 0 \end{bmatrix} + a_2 \begin{bmatrix} y_1 \\ \vdots \\ y_n \\ 0 \\ 0 \\ 0 \end{bmatrix} + b_1 \begin{bmatrix} r_{11}\ln r_{11} \\ \vdots \\ r_{n1}\ln r_{n1} \\ 1 \\ x_1 \\ y_1 \end{bmatrix} + \ldots + b_n \begin{bmatrix} r_{1n}\ln r_{1n} \\ \vdots \\ r_{nn}\ln r_{nn} \\ 1 \\ x_n \\ y_n \end{bmatrix}, \tag{3.21}$$

which written as a full matrix equation is

$$\begin{bmatrix} f(x_1,y_1) \\ f(x_2,y_2) \\ \vdots \\ f(x_n,y_n) \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 & r_{12}^2\ln r_{12}^2 & \ldots & r_{1n}^2\ln r_{1n}^2 & 1 & x_1 & y_1 \\ r_{21}^2\ln r_{21}^2 & 0 & \ldots & r_{2n}^2\ln r_{2n}^2 & 1 & x_2 & y_2 \\ \vdots & & \ddots & & & \vdots & \\ r_{n1}^2\ln r_{n1}^2 & r_{n2}^2\ln r_{n2}^2 & \ldots & 0 & 1 & x_n & y_n \\ 1 & 1 & \ldots & 1 & 0 & 0 & 0 \\ x_1 & x_2 & \ldots & x_n & 0 & 0 & 0 \\ y_1 & y_2 & \ldots & y_n & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \\ a_0 \\ a_1 \\ a_2 \end{bmatrix}. \tag{3.22}$$

By defining vectors, $\mathbf{f}$ and $\mathbf{a}_f$, and a matrix, L, Equation 3.22 can be written as $\mathbf{f} = L\mathbf{a}_f$. The coefficients, $\mathbf{a}_f$, are now found by inverting L,

$$\mathbf{a}_f = L^{-1}\mathbf{f}. \tag{3.23}$$

In reference to the original statement of the surface fitting problem given in Equations 3.16 and 3.17, $\mathbf{f}$ is the height of the surface at each control point, i.e., $\mathbf{f} = \mathbf{x}'$ where $\mathbf{x}' = [x'_1 \ldots x'_n]^T$. The coefficients for the second surface, $g(x,y)$, are found by using the y' values for heights while the L matrix remains the same. Now, there are coefficients for each surface, i.e., $\mathbf{a}_f = L^{-1}\mathbf{f}$ and $\mathbf{a}_g = L^{-1}\mathbf{g}$.

The two sets of coefficients define the surfaces, $f(x, y)$ and $g(x, y)$, which give the new coordinates of the pixel, x' and y', respectively. The location of all pixels, (x, y), in the slave image are mapped to new locations, (x', y'). The slave image has now been stretched, squeezed, and shifted to better match the master image.

The computational cost of this transformation is largely dependent on inverting L. For a lot of control points inverting the L matrix is computationally intensive. When deciding on the number of sections per block this should be taken into account. In practice, 64 sections (dividing the block in eight along each side) works well. If each section is shifted by approximately the same amount the L matrix is highly singular. This is because the surface defined by equal shifts is a flat plane. In this case, performing spline warping has little advantage over a global subpixel shift of the block.

In summary, the warping algorithm finds the shifts for each pixel in a block based on the shifts of control point pixels. Less intensive methods of interpolation are also capable of determining shifts for every pixel. In [2], a simple bi-linear interpolation is suggested. Over scenes with flat topography simple interpolation methods perform well with much less computation. However, for scenes with significant topography; robust methods, like warping, are recommended [7].

3.2.5 Interpolating to a Regular Grid

The warping algorithm results in an image with pixels at irregular locations and spacings. An image with regularly spaced pixels is necessary to compare with the master image. To create this image of regularly spaced pixels, values must be interpolated between the irregularly spaced pixels. This resamples the warped slave image to match the master image.

A truncated sinc is used as the basis function for the interpolation. The value for every regularly spaced pixel is the sum of scaled and shifted sinc functions. This is given in the following equation [2],

$$x_m = \sum_{n=-N}^{N} x_n \operatorname{sinc}\left(\frac{t_m - t_n}{T}\right). \tag{3.24}$$

In this equation, x_m is a new pixel at mT', where T' is the new sampling interval. The original values are at $t_n = nT$, where T is the former sampling interval. From sampling theory, given that the signal is band-limited and has been sampled evenly above Nyquist, then $x_n = x(nT)$ is a complete representation of the underlying continuous signal $x(t)$. If the new sampling rate, $1/T'$, is above Nyquist it represents $x(t)$ completely.

For the problem at hand, the input sampling interval is not constant. The interval slowly varies between values slightly above or below unity ($0.99 < T < 1.01$). Going from an irregularly sampled space to a regular one can be complex and computationally intensive. In order to keep the interpolation routine simple and because the sampling interval is slowly varying, a constant sampling interval of unity is assumed. If the sampling interval in the warped image varies greatly the truncated sinc interpolation is not valid and a more complex algorithm is needed.

Figure 3.6 shows the interpolation across a line of pixels in a real image. The warped pixel values in this image are regularly spaced, and the sinc interpolation creates new values that closely follow the old values. The interpolation is applied across one dimension of the image and then the other. Each pixel in the interpolated image is on a regularly space grid that matches the sampling grid of the master image.

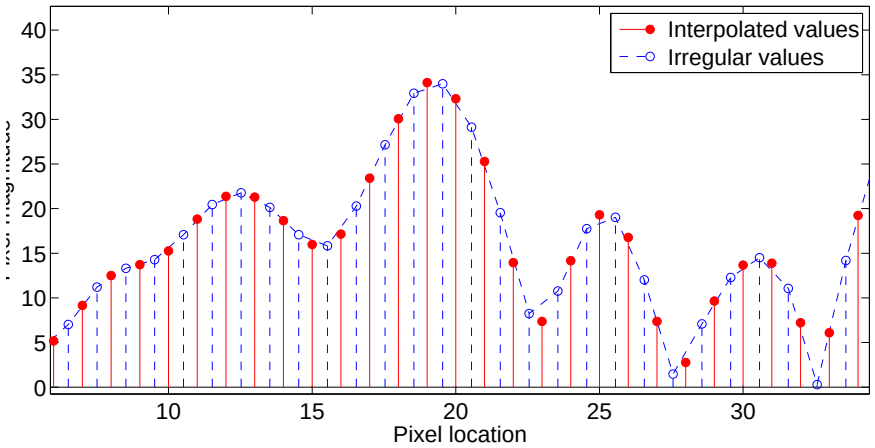

Figure 3.6: A row of pixels from a warped SAR image is shown. The dotted line follows the irregularly spaced pixels that are at noninteger steps. The interpolated values are regularly spaced. The sinc interpolant provides smoother behavior around the peaks and troughs than linear interpolation.

The slave image has now been warped and resampled to better match the master image. The maximum correlation registration algorithm compensates for small differences in collection geometry and image formation between the master and slave images so that the cross-correlation between the two images has been maximized.

3.3 Registration Results

In this section, the registration algorithm described previously is tested on different data sets. To evaluate how well the images are registered, a coherence map and phase difference are used. In Chapters 4 and 5, the coherence map and phase difference are used for CCD and interferometry. Here, they are briefly described in relation to their utility in evaluating registration.

The coherence between images is the normalized cross-correlation. The coherence map shows the coherence between corresponding regions of two images. Figure 3.7 shows two speckle dominated images and the coherence map. Areas of high correlation have values close to unity and areas of low correlation are close to zero. For the color scale commonly used for coherence maps, registration increases the coherence leading to a whiter coherence map.

If two images are collected from close to the same angle the coherent nature of SAR imaging results in a phase relation between the two complex images. The phase between images is related to the different distance that each pulse must travel to the target and back to different antennas. The phase difference, ϕ, between two images, f_1 and f_2, is given by

$$\phi(x,y) = \angle \left(f_1(x,y) f_2(x,y)^* \right). \tag{3.25}$$

The phase in a SAR image is sensitive to noise. To compensate for this the phase is estimated using a window of neighboring pixels,

$$\phi_{ML} = \angle \left(\sum_N f_1 f_2^* \right). \tag{3.26}$$

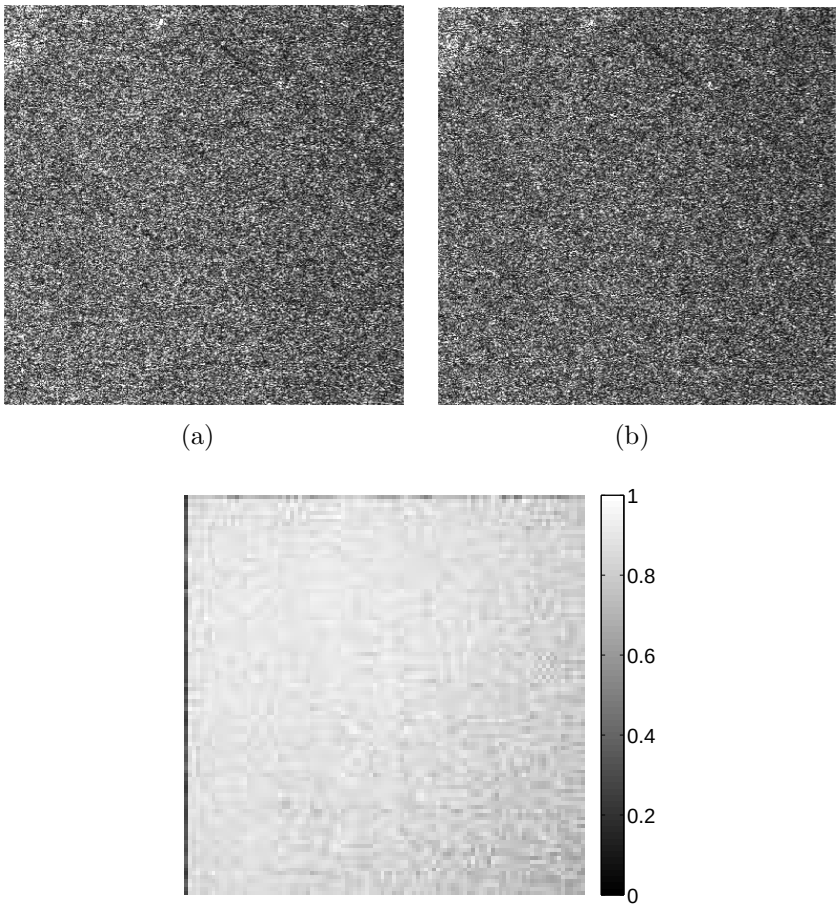

(a) (b)

(c)

Figure 3.7: An example of coherence between two images of a flat desert area. A dirt road cuts across the top right corner of the images. (a) SAR image from the initial pass. (b) Repeat-pass SAR image of the same area. (c) Coherence map. Based on the large value of the coherence map, the two images are highly correlated.

This produces a smoother phase map. Figure 3.8 shows the phase difference between two images before and after averaging. A smooth phase surface between images is desired. When the phase wraps from π to $-\pi$ the transition should be sharp. In the color map used, well registered images smoothly change shade through different levels of gray then sharply transition from white to black.

In the following examples, the collection method and other properties of the data set are first given. Then, the coherence map and phase difference are used to evaluate the level of image registration before and after the registration algorithm has been performed.

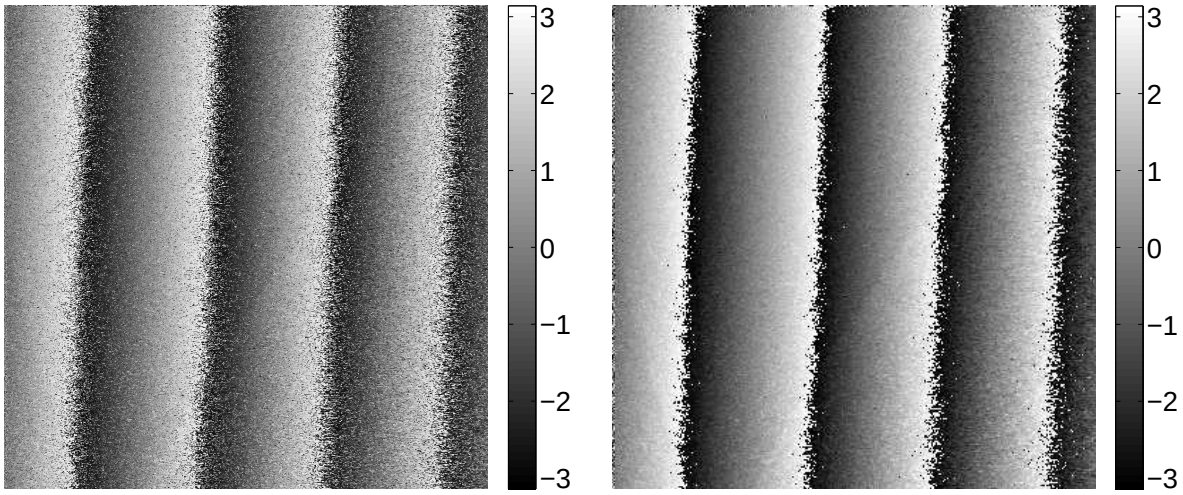

(a) (b)

Figure 3.8: (a) Phase difference between two images without averaging. The stripes are caused by the phase wrapping from π to $-\pi$. (b) Phase difference obtained by applying an averaging filter to reduce noise.

3.3.1 Single-Pass Collection

The registration algorithm is first tested on images from BYU's interferometric X-Band SAR, named YINSAR. The YINSAR performs single-pass interferometry using two antennas. This is convenient for testing registration because it is not necessary to find two different passes that are close enough to each other. The flight paths for both images are identical except for the baseline distance between the mounted positions of the two antenna on the aircraft.

Master and slave images from a portion of a YINSAR flight are shown in Figure 3.9. The scene shows a road cutting at a vertical slant across the image surrounded by bushes. Even though both images are taken from almost the same vantage point, there are slight differences between them. These are caused by differences in the two receiver channels (different receiver noise) and from the slightly different collection angle.

For interferometric flights, the horizontal spacing of the antennas on the YINSAR is about one meter which corresponds to the range resolution of the radar. Because of this, in the past, the images were registered by moving the slave image one pixel in range relative

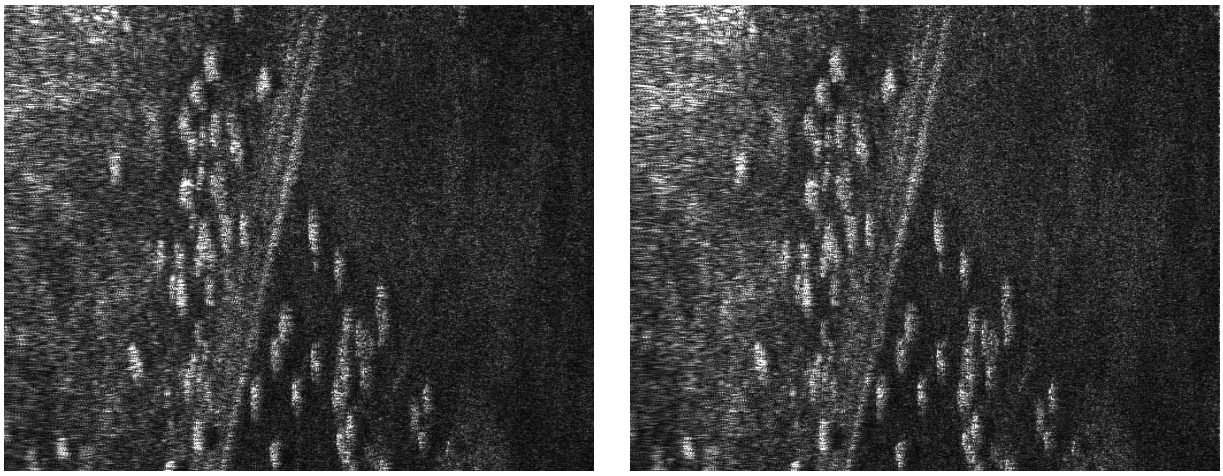

(a) (b)

Figure 3.9: (a) Master image used in single-pass interferometry example. (b) Slave image.

to the master. The results of this simple shift as compared to the maximum correlation algorithm presented in this chapter are shown in Figure 3.10.

Figure 3.10(b) shows the coherence map from registered images and Figure 3.10(a) shows the coherence map for shifted images. The maximum correlation algorithm has increased the coherence between images. The algorithm does not always perform perfectly as can be seen by the loss of coherence in the lower left corner. Overall, the average coherence across the image increased from 0.570 to 0.672. It is possible to make out bushes and the bank along the road in the coherence map. These objects and areas correspond to the bright spots in the SAR image. The darker parts of the images are dominated by uncorrelated noise, while the areas of higher SNR remain coherent.

By comparing the shifted phase difference image of Figure 3.10(c) and the registered phase image of Figure 3.10(d) improvements in registration can be seen. The registered phase image shows clearer phase wrapping in the upper half of the image. Also, throughout the rest of the image the phase is less noisy than before.

The loss of coherence in top half of Figure 3.10(a) is likely caused by a change in the roll, pitch, or yaw of the airplane. This motion changes the position of the antennas relative to each other and the ground even though they are still mounted in the same locations on the aircraft. The transition from coherent and less coherent sections is in the azimuth dimension

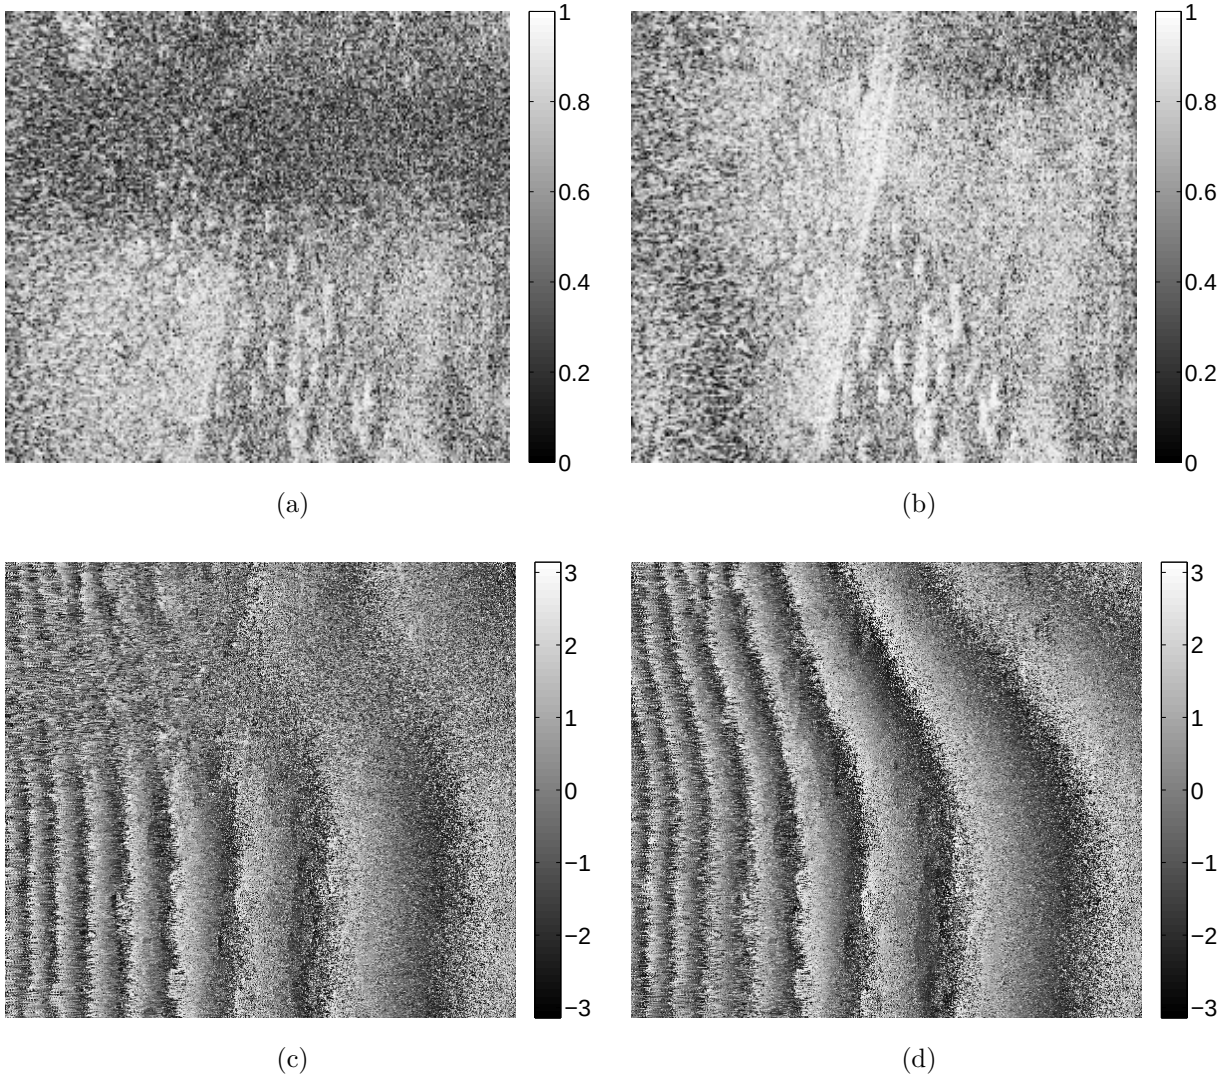

Figure 3.10: Coherence and phase images corresponding to the images in Figure 3.9. (a) Coherence map formed by shifting the slave image one pixel in range. (b) Coherence map after the registration algorithm. (c) Phase difference from simple shift. (d) Phase difference after using the registration algorithm.

which corresponds to the airplane's motion with time. This motion can be compensated for in processing if it is accurately known; however, roll, pitch, and yaw are not always accurately measured. Image-based registration, without knowledge of the plane's motion, can be used to compensate for inaccurate motion measurements in single-pass interferometry.

3.3.2 Repeat-Pass Collection

This example is from an L-Band (1.6 GHz) repeat-pass collection. The interval between collections is about one hour. The passes are chosen because the flight paths are within 10 meters of each other for almost one kilometer. The images are processed to 0.5×0.5 m resolution using detailed motion data and the backprojection routine. A flat surface is used as the ground model in forming the image. The registration performance at each step of the maximum correlation algorithm is shown.

In this example, an area of desert terrain with a small mountainous region and several roads is imaged as shown in Figure 3.11. Figure 3.11(a) is the speckled image that is used for registration and coherence estimation, and Figure 3.11(b) is the multi-looked image with reduced speckle. The plane flew along the right side of the image resulting in the right side of the mountains being illuminated and the left side being in shadows. The image also has lower SNR in the far range consistent with free-space spreading loss.

The first step in the registration algorithm is to find a two-dimensional shift of the slave image that maximizes the cross-correlation between the two images. The coherence map that results from this shift is shown in Figure 3.12(a). Over most of the flat terrain there is high coherence. However, there is lower coherence in the mountainous region (see A1) caused by assuming a flat earth in image creation. The azimuth fluctuations in coherence (see A2) are likely a result of inaccurate aircraft motion data. Both of these areas are improved by further image registration.

The phase difference image, Figure 3.12(b), shows similar problems to the coherence image. In areas of high coherence the phase wrapping is smooth with sharp changes from π to $-\pi$, but in areas of low coherence the phase is noisy (see B1). Another feature of note is a small distortion in the phase which corresponds to a small hill in the SAR images (see B2).

The second registration step is to break the image into blocks and perform a shift on each block. The results of this step are shown in Figure 3.13. At block edges there is often a discontinuity in the phase image (see D1) which causes lines of low coherence in the coherence map (see C1). In C3 and D3, the shift has been miscalculated causing low coherence and a random phase between images. Often when a block is shifted, the shift lines

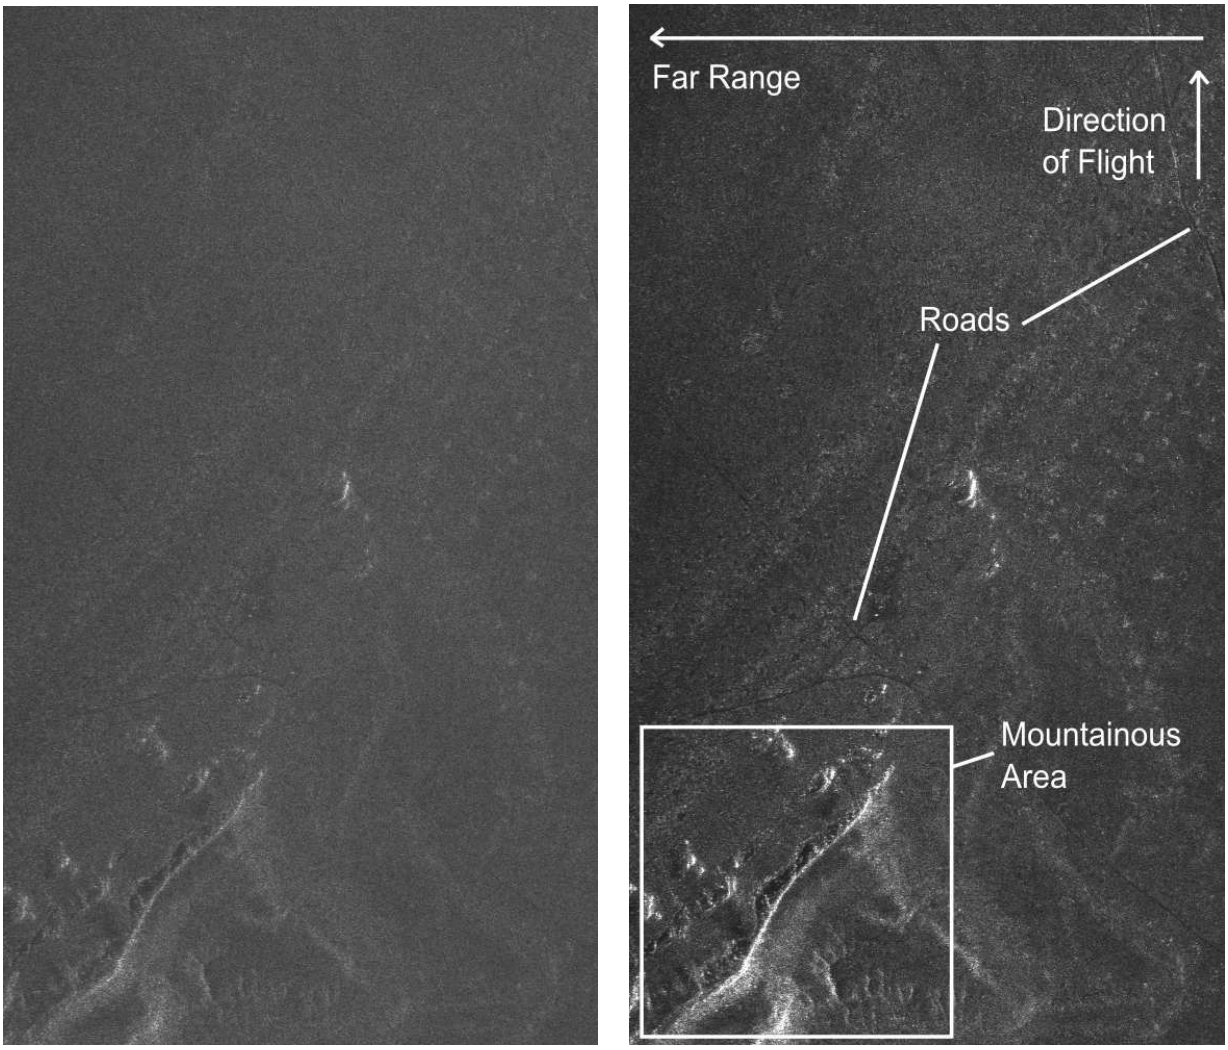

(a) (b)

Figure 3.11: Images of desert area used for the repeat-pass registration example. (a) Image with speckle as used for registration and coherence estimation. (b) Image that has been multi-looked to reduce speckle.

up part of the block better than the other part. For example, the blocks on the right edge of the image (see C2) have been shifted to cause higher coherence on the right half but much lower coherence on the left half. To register an entire block, shifts of different magnitudes are needed within the block.

The last registration step is to apply a warp to the slave image within each block. The results are shown in Figure 3.14. Warping fixes several of the problems with block shifting. The coherence has been restored to the block that lost coherence. There are less lines at

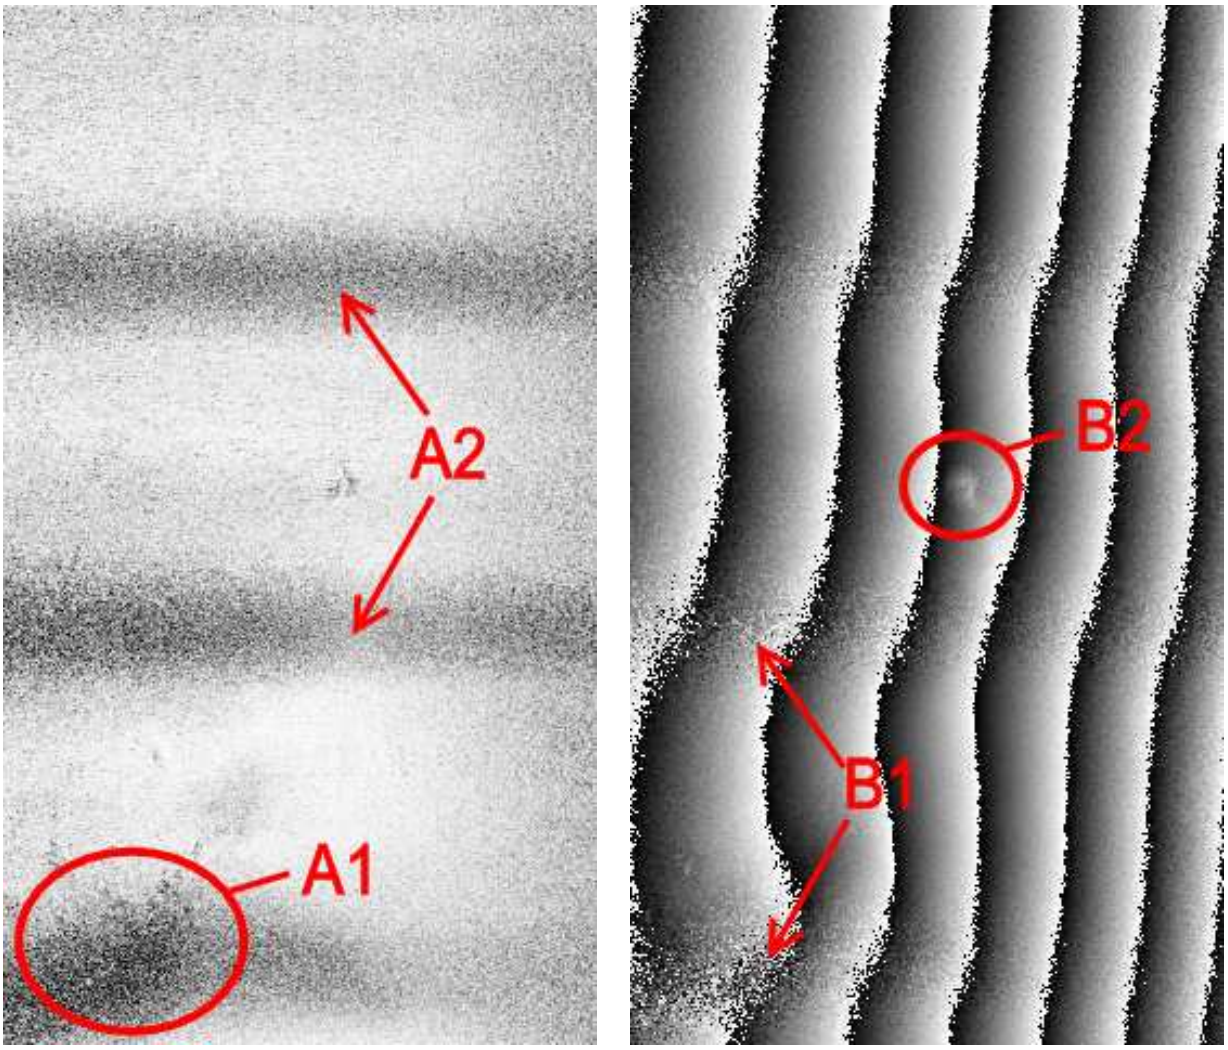

(a) (b)

Figure 3.12: The results of registering the repeat-pass example images by shifting the slave image relative to the master image are shown in (a) the coherence map and (b) the phase difference.

block edges in the coherence map and phase difference. The blocks on the right edge have high coherence across the block. A few lines remain from the block shifts (see E1 and F1), and there is a line of slightly lower coherence along the right edge (see E2). These areas show that minor improvements could still be made to the algorithm.

A visualization of the warping performed on the slave image during registration is shown in Figure 3.15. In Figure 3.15, the shifts are displayed as a vector field, and as separate

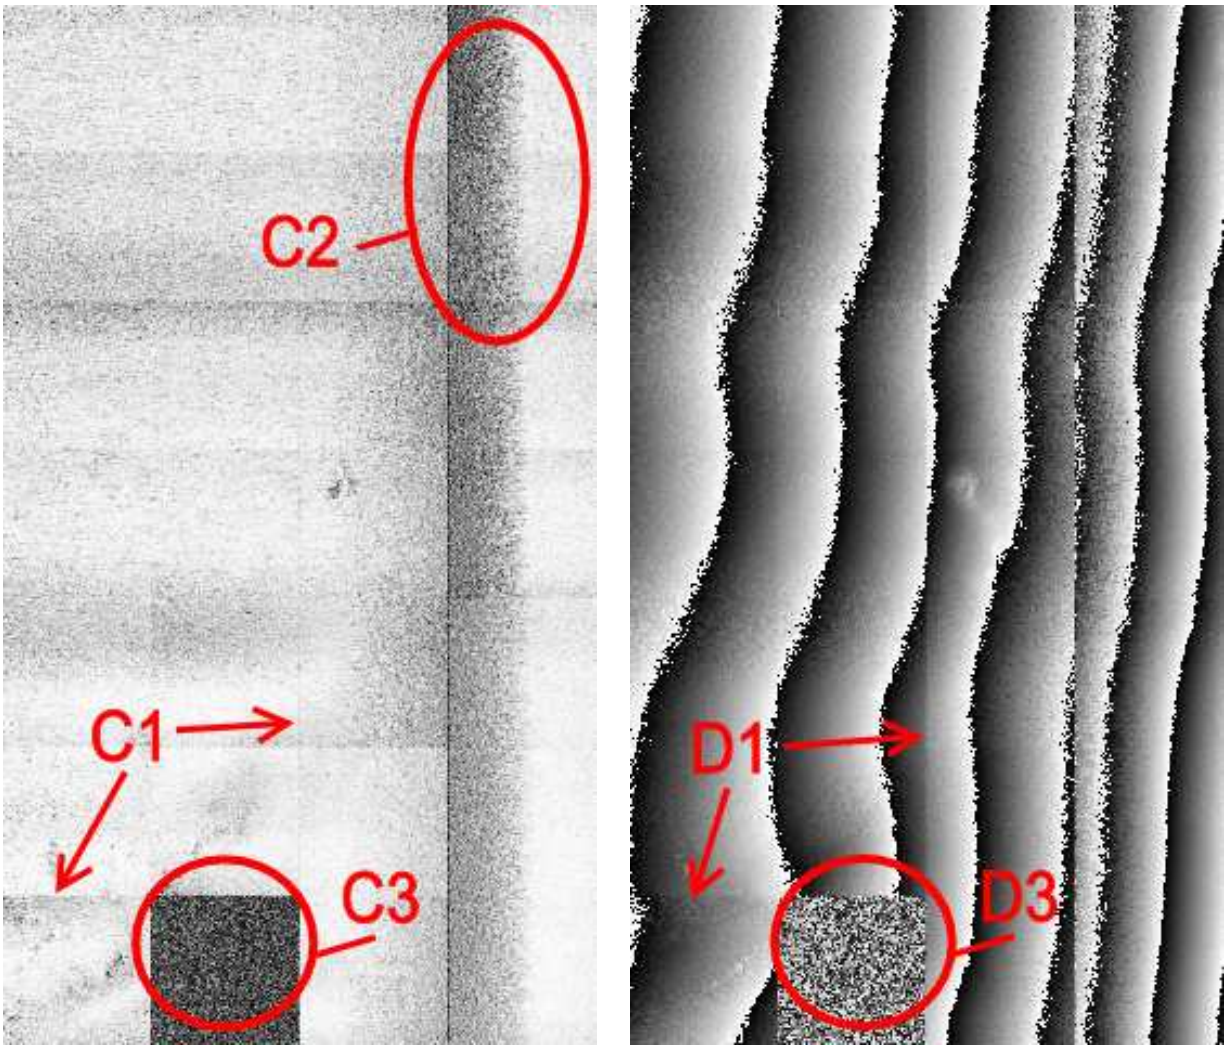

(a) (b)

Figure 3.13: (a) Coherence map and (b) phase difference after performing block shifts to register the images.

components in the down and right directions. A relatively smooth surface of shifts is shown in the images, though there are still a few sharp lines corresponding to block edges.

Several slight modifications to the maximum correlation algorithm are made in processing this example. The blocks are chosen to be 512×512 pixels but overlap with each other by 50 pixels. Only the warping performed on the center 412 pixels is used in the final registration. This reduces the discontinuities at block edges. In step three of the registration algorithm (i.e., warping), the complex value for each pixel is used. Step four of the algorithm

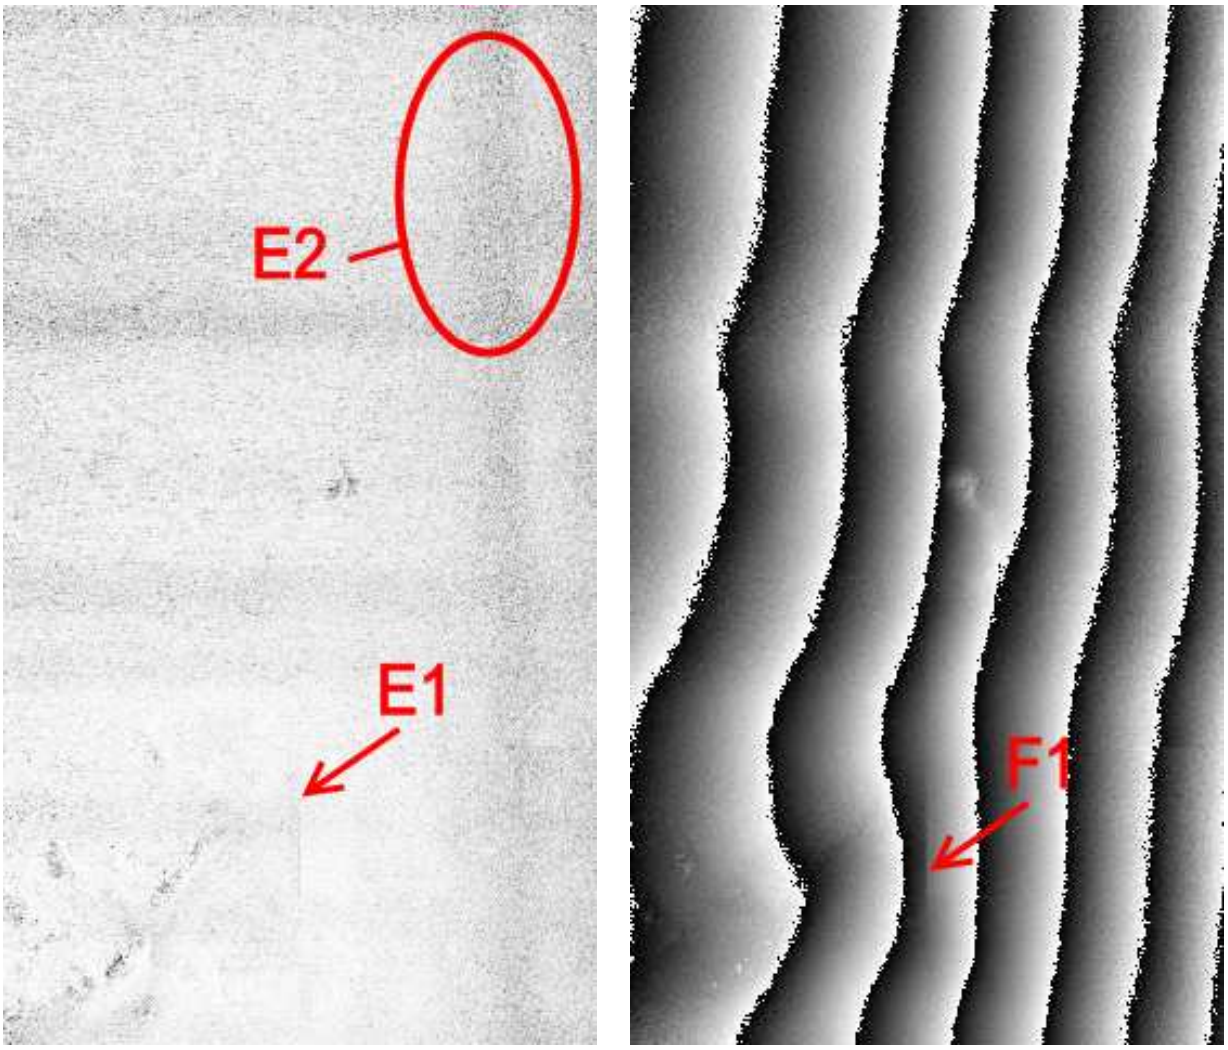

(a) (b)

Figure 3.14: (a) Coherence map and (b) phase difference after performing block warping to register the images.

is not done. This yields slightly better results than sequentially warping with magnitude and then complex pixel values.

The registration performance is quantified in Table 3.1 by calculating the average coherence across the coherence map for each step of the algorithm. As seen in the coherence maps, each registration step increases the overall coherence.

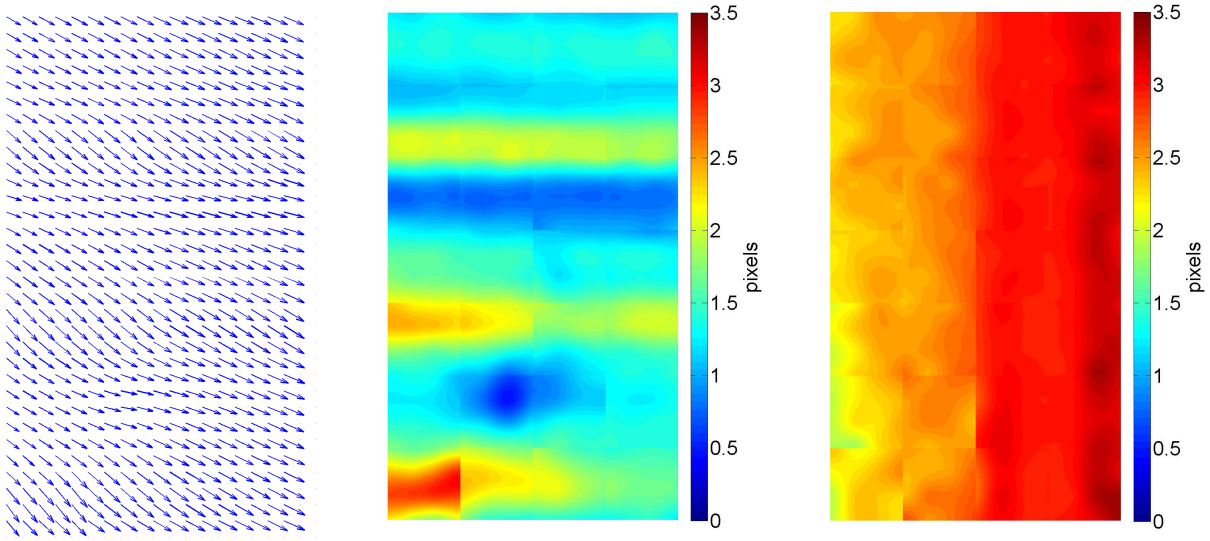

(a) (b) (c)

Figure 3.15: The shifts performed on the slave image for repeat-pass example. (a) Shifts displayed as a vector field. (b) Magnitude of downward shift. (c) Magnitude of shift to the right.

Table 3.1: Coherence after each step of the registration algorithm for the repeat-pass example

Step	**Avg. Coherence**
Before registration	0.2114
Global image shift	0.7943
Block shifts	0.8199
Block warping	0.8979

3.4 Chapter Summary

This chapter discusses the need for registration in SAR images. The maximum correlation registration algorithm is presented. Several practical observations about the implementation of the algorithm are given (e.g., setting threshold levels, efficient block and section sizes, etc.). The algorithm is used to register a single-pass pair of images and a repeat-pass pair of images.

CHAPTER 4. COHERENT CHANGE DETECTION

Coherent change detection (CCD) relies on careful data collection, precise image formation, and accurate estimation. This chapter defines the statistical estimation problem behind coherence calculations and how an accurate coherence value is reached. Practical steps to increase the coherence between images after they have been collected are given. A new method that combines the coherence estimated from multiple sub-apertures is presented. Several examples of coherent change detection, demonstrating simulated and actual results, are given.

4.1 Coherence Estimation

The coherence, γ, between images, f and g, is defined as the magnitude of the normalized complex cross-correlation. This is given as

$$\gamma = \left| \frac{E[fg^*]}{\sqrt{E[f^2]E[g^2]}} \right|. \tag{4.1}$$

In practice, the expected values in Equation 4.1 are estimated as the ensemble average over a window of pixels,

$$\gamma = \left| \frac{\sum_N fg^*}{\sqrt{\sum_N f^2 \sum_N g^2}} \right|. \tag{4.2}$$

For the ensemble average to be a consistent and unbiased estimate of the expected value, several requirements must be met. All pixels within the window (N) should be independent and be drawn from the same statistical distribution [20]. In SAR images, pixels that are spaced at or above the native range and azimuth resolutions, as determined by the antenna beam and bandwidth, are independent. However, images are often processed at a pixel

spacing smaller than the system resolution. In this case, the window of pixels should be enlarged to include an appropriate amount of independent pixels. If this is not done, the coherence will be biased toward unity.

In general, the extent of the pixel window should match the size of the changes expected to be observed in CCD. If the physical area covered by the pixel window is too large, small areas of low coherence may be averaged with areas of high coherence and, therefore, missed. In order to detect small localized changes, a high resolution system is required to provide a sufficient amount of independent pixels over a small physical area.

When calculating the coherence map for an image, the averaged interferogram is found at the same time. This is because the phase of the cross-correlation is the interferometric phase,

$$\phi = \angle \frac{\sum_N fg^*}{\sqrt{\sum_N f^2 \sum_N g^2}} \tag{4.3}$$

$$= \angle \sum_N fg^*. \tag{4.4}$$

The phase between SAR images is very sensitive to noise which is reduced by averaging over a window. The window that generates a good estimate of the interferometric phase is the same window that produces a good estimate of the coherence. Thus, the complex cross-correlation contains the information necessary for both interferometry and CCD.

4.2 Coherent Processing

In processing an image pair, several steps have been found to increase the coherence between images. These steps increase the CCD capability of a fixed-antenna, strip-map SAR operating at a constant pulse repitition frequency (PRF). However, systems capable of steering the antenna and adjusting the PRF based on real-time motion measurements do not need many of the steps mentioned in this section. For processing, the backprojection algorithm is used, but similar steps can be adapted to frequency domain algorithms.

Antenna Squint The motion of the aircraft causes the beam of the antenna to be squinted forward or backward relative to broadside, even for a nominally linear track. During

processing, antenna squint results in some parts of the image being created with a wider range of angles than other parts of the image. When images from two passes are compared, there is often a large difference in the angles used to form a pixel in the initial image versus the pixel that correspond to the same physical area in the repeat-pass image. Processing can compensate for this difference in angle by filtering the data based on angle of arrival. A common linear track is used to define a range of arrival angles, known as a sub-aperture. Any data collected from angles outside of the sub-aperture is not used to form the image. Figure 4.1 gives an example of how a restricted aperture compensates for squint during the flight.

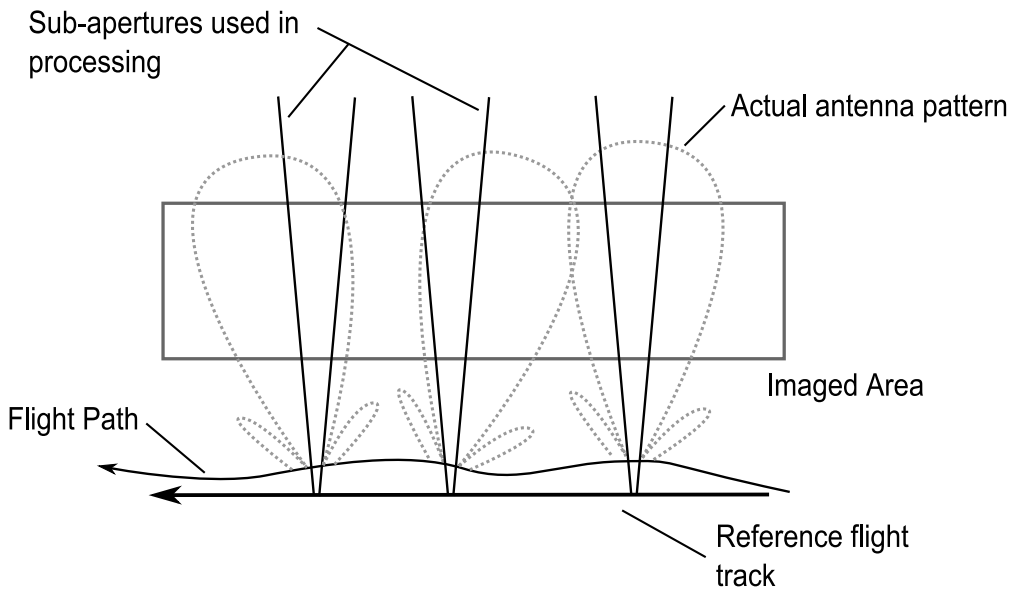

Figure 4.1: During a normal collection the motion of the plane causes the antenna beam to be squinted forward or backward at different times. Fixed apertures, defined relative to a linear reference track, restrict the azimuth range of angles.

When aperture filtering is implemented for backprojection, a filter is applied at each pixel. Instead of restricting which part of the antenna beam is processed, each pixel restricts which pulses are used to create the pixel. Each pulse is geometrically checked to see if it is within the pixel's aperture. Figure 4.2 shows the problem from this perspective.

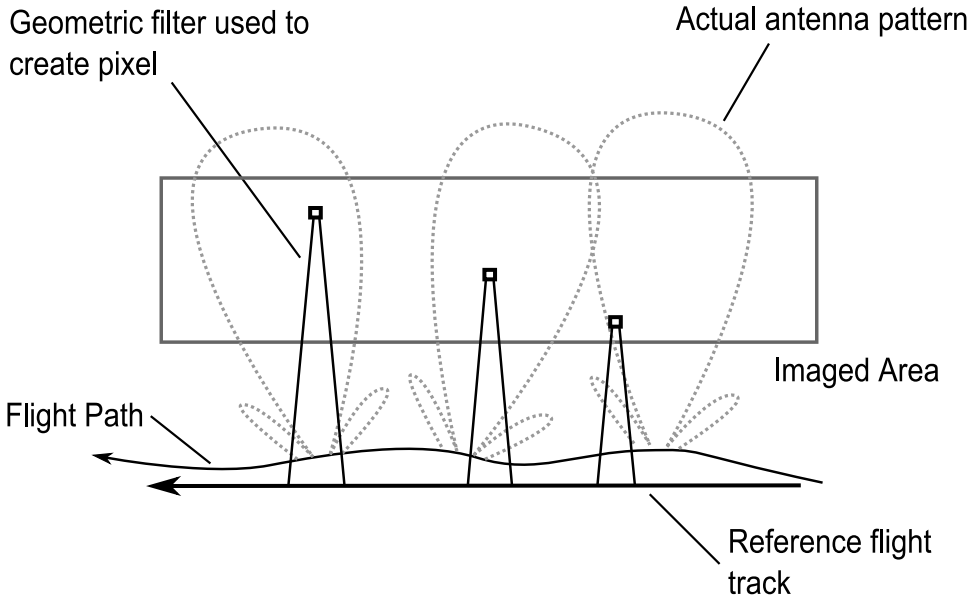

Figure 4.2: For backprojection processing, a geometric filter is used that defines a range of angles for each pixel.

Restricting the aperture works well for an antenna with a wide azimuth mainlobe. This is because when the antenna is squinted by a few degrees, the weighting of the aperture by the antenna beam is not significantly different than when the antenna is pointed broadside. However, for a narrow beam antenna, the mainlobe can be squinted such that it is outside of the image-forming sub-aperture. When an image is created, parts of the image will drop out because only weak returns come in from the sidelobes.

Antenna Pattern A further improvement in processing can be made by compensating for the antenna pattern. This is especially applicable in conjunction with processing a fixed broadside aperture as discussed above. The main lobe of the antenna can be weighted such that there is an effectively constant gain across the beam. Processing a few degrees off of the beam center because of squint is then equivalent to having processed the center of the beam.

For a narrow beam antenna this method does not work. Data collected a few degrees off of the beam center is in the sidelobes of the antenna pattern. The radar data received from the sidelobes is attenuated by the antenna pattern resulting in a lower SNR inherent to the returns. Changing the weighting of this data increases the noise

in the image. The better approach for narrow beam collection is to use a gimbaled antenna. All data is then collected at the correct broadside angle relative to some predefined path and no compensation in processing in necessary.

PRF/Velocity Ideally, radar pulses are transmitted at regularly spaced locations. Some radars feed velocity information from a motion sensor into the radar electronics to ensure regular space pulses. However, simpler radars pulse at a regular temporal interval, known as the pulse repetition frequency (PRF). Constant sampling in time leads to more pulses per distance traveled being collected when the radar is traveling slower and less pulses when the radar is traveling faster. A basic compensation is to weight the pulses in an aperture depending on the velocity of the aircraft relative to a constant reference velocity. A more elaborate compensation would be to interpolate the raw radar pulses to regularly spaced samples. Given that SAR processing relies on neighboring pulses being correlated, a smooth interpolation of new pulses is possible.

These steps make it possible to recover the correlation between images that would have otherwise been lost with normal processing. Because all steps are performed after the collection of images, they provide a way to increase coherence without increasing the complexity and cost of the radar hardware.

4.3 CCD Using Multiple Sub-apertures

A wide beam antenna gathers data over a broad range of angles. In processing the azimuth beam is often split into multiple sub-apertures and separate images are created for each section. Given that the sub-apertures don't overlap, each image is statistically independent of the other images because it represents a target's backscatter from a different range of angles. After image formation, the sub-aperture images can be power averaged to create a higher contrast, low speckle image.

The method presented in this section uses sub-apertures in a similar way to increase the coherence between two images. First, the method processes two flights using k sub-apertures for each flight. Then the coherence between corresponding sub-apertures from the reference and repeat-pass flights is estimated to create k coherence maps. For each pixel

location, (x, y); the maximum coherence value among sub-apertures is chosen to represent that location in a combined coherence map as given by

$$\gamma(x, y) = \max_k \gamma_k(x, y). \tag{4.5}$$

When many sub-apertures are used, the method results in a high average coherence for a stationary scene.

As an example of multi-aperture CCD, consider a pair of L-Band collections taken with a widebeam antenna. The coherence maps for nine of twelve sub-apertures processed from a 60° beam are shown in Figure 4.3. Each aperture is 6° wide and the apertures overlap by a small amount. As previously, the coherence map is scaled from zero to unity with dark areas having low coherence and white areas high coherence. The combined coherence map, as found by taking the maximum coherence across sub-apertures, is shown in Figure 4.4(a). All the areas of low coherence found in the sub-apertures have been removed from the multi-aperture coherence map.

Coherence estimation with a wide beam could also be achieved by processing the image to a high resolution and then finding the coherence over a very large window. However, the current method has several benefits over creating higher resolution images. One advantage is the reduced processing time. As explained in Section 2.4, the computational time of the backprojection algorithm is driven by finding the range from each pulse to each pixel. Increasing the azimuth resolution by a factor of a requires a larger aperture, aL, and more pixels in one dimension, aN, resulting in a computational time proportional to $a^2 L \times M \times N$. In creating sub-apertures images the range from each pulse to each pixel is calculate once and then reused for every sub-aperture. A small overhead is added but the computational time remains proportional to $L \times M \times N$.

Another advantage is that residual motion errors (motion in the flight track not recorded by the GPS) do not accumulate over a large aperture. Instead, the motion errors are processed at different times for each sub-aperture. Any given location on the ground is imaged from different angles and, therefore, from a different part of the flight path for each sub-aperture. Areas of low coherence caused by motion errors in a section of the flight path

will shift through the sub-apertures as the error section is processed from different angles. Movement of the low coherence area is demonstrated in the sub-aperture coherence maps in Figure 4.3.

Another advantage of the sub-aperture CCD method is that it can replace motion compensation and registration routines. Areas of low coherence caused by residual motion errors can be remedied by using a motion compensation or autofocus routine [5], [21]. Also, registration algorithms can correct areas of low coherence as shown with the correlation-based registration routine presented in Chapter 3 and in other work [4]. However, with sub-apertures these areas of low coherence are ignored because other sub-apertures have good coherence for the same area. If registration is performed in addition to multi-aperture coherence estimation, as shown in Figure 4.4(b), there is little improvement in the coherence. This demonstrates that the multi-aperture method is sufficient to overcome low coherence caused by motion errors.

In addition to accommodating for flight errors, multi-aperture CCD overcomes areas of low SNR caused by shadowing. In areas shadowed by the radar, there is low coherence because the returns are dominated by random noise instead of correlated signal. Any changes in these areas can not be detected. In fact, areas of low SNR common to both images are often masked so they are not marked as false changes. With multiple sub-apertures an area in shadow in one aperture is often visible in another. Figure 4.5 shows the coherence map for a mountainous region for a single aperture and for five combined apertures. In the coherence image, ridge lines and other high SNR areas are visible as white, high-coherence areas. Areas of low coherence corresponding to slopes facing away from the radar are dark (the radar flew diagonally across the top left corner of the image). With multi-aperture CCD, many of these areas of low coherence are removed.

One disadvantage of combining multiple sub-apertures based on the maximum value is the loss of sensitivity at different angles. Changes that are more apparent at a particulare range of angles represented by one sub-aperture will be missed when they are combined with coherence maps from other sub-apertures. If each sub-aperture was autofocused, registered, and used independently, they provide more possibilities to detect the change; whereas, by combining sub-apertures this information is lost. But for changes in the scene that affect the

scattering properties of a target from many angles, combining multiple sub-apertures works well.

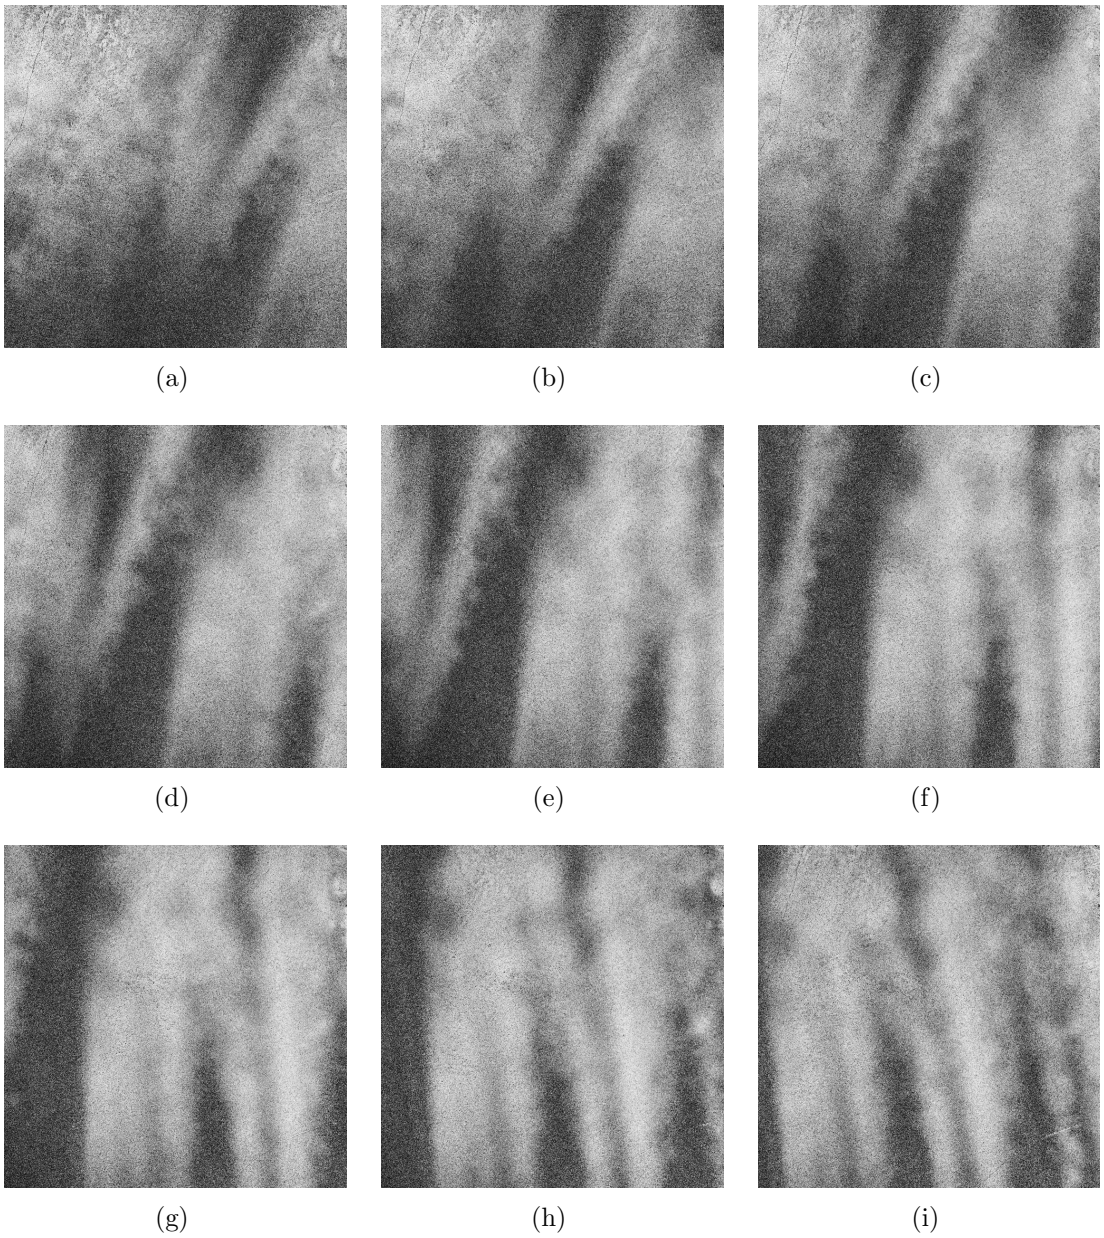

(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

Figure 4.3: Coherence maps for nine sub-apertures, each covering 6°. Errors in the flight path result in a loss of coherence at different areas in each coherence map.

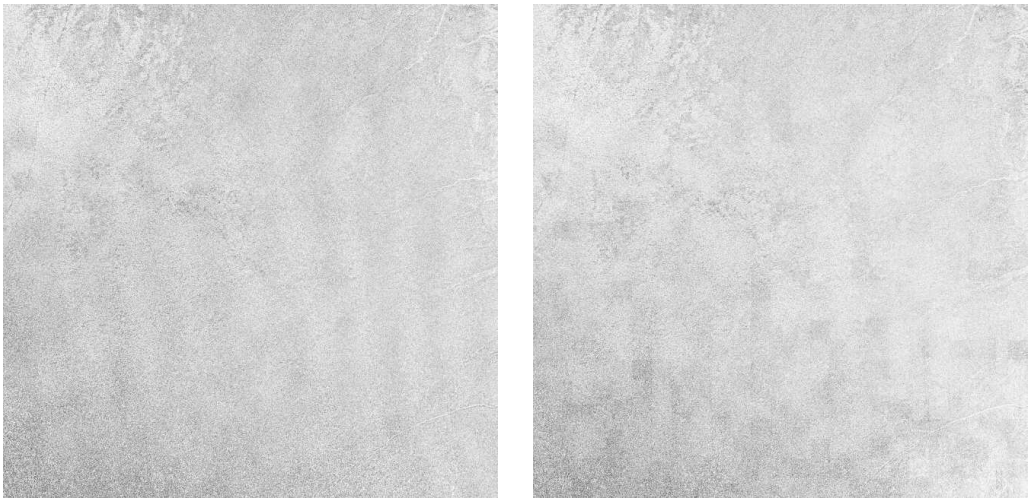

(a) (b)

Figure 4.4: (a) Multi-aperture coherence map created by taking the maximum value at each pixel between sub-aperture coherence maps. (b) Multi-aperture coherence map created after first registering each sub-aperture. The block structure is an artifact of the registration process.

4.4 L-Band Simulation

This section presents a basic simulation of automated coherent change detection. The same set of sub-apertures as shown in Figure 4.3 is used. A target is introduced into one of the images by changing the phase in portions of the image. A change of phase would go unnoticed in magnitude images; therefore, the detection of phase changes illustrates the sensitivity of CCD. Automated detection routines are used to find the probability of detection and the probability of false alarm.

The images used in this example were taken over a desert area at L-Band (1.6 GHz) with a bandwidth of 130 MHz. A 60° azimuth antenna beam was used to collect the data. In processing, the beam is split into 12 sub-apertures, each 6° wide. Each image has a 0.5 x 0.5 m pixel resolution with a native resolution of about 1 m in azimuth and range. A ring shaped target is created by adding an extra random phase to pixels in each sub-aperture image generated from the second flight. The random phase is uniformly distributed from 0 to 2π and is added to a ring that is five pixels wide (2.5 m). Such a target could represent the disturbances left along the path of a large animal or vehicle.

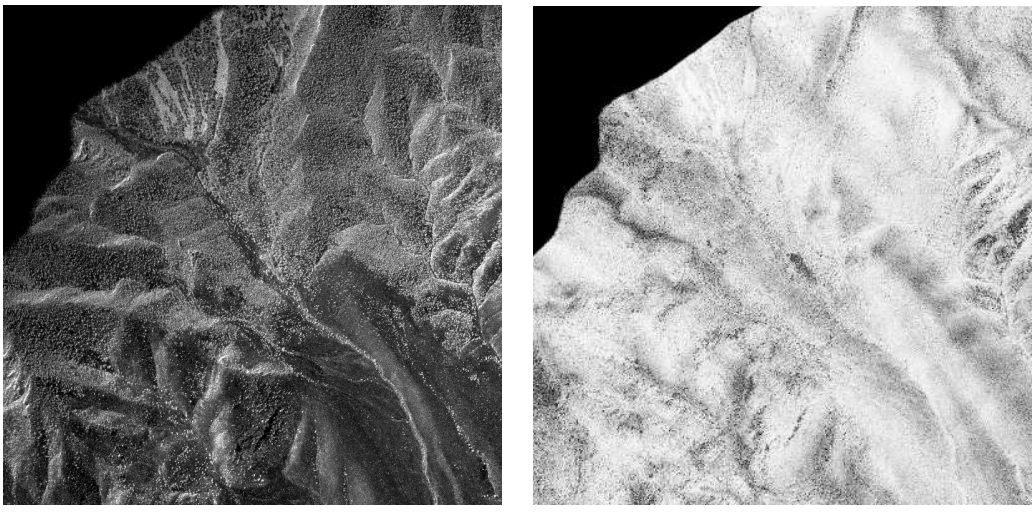

(a)

(b)

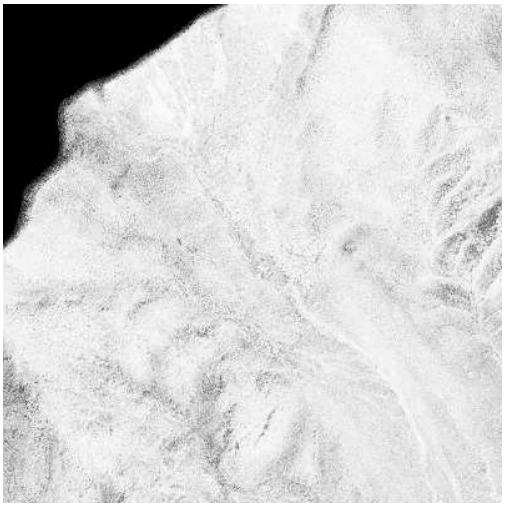

(c)

Figure 4.5: (a) Radar image of mountainous area showing bright slopes and shadows caused by terrain. (b) Coherence map from a single sub-aperture. Many of the low coherence areas correspond to slopes facing away from the radar. (c) Coherence map create by combining sub-apertures. The coherence has increased in many areas of previous shadow.

Two automatic detection methods are used to demonstrate CCD performance. The first method is a basic threshold detector; every pixel in the coherence map is compared to a fixed threshold value, and pixels below the threshold are marked as changes. When large areas of the image have low coherence this method produces many false detections. Another method for detecting changes is the cell-average detector in which a pixel is tested by comparing its value to the average value of a window of pixels surrounding the test pixel.

Often a buffering ring of pixels is placed between the pixel under test and the reference pixels. If the difference between the test and reference pixels is large enough, a change is marked. This method prevents false detections in areas of low coherence and enables slight changes in areas of high coherence to be detected.

The performance of the detectors is measured in terms of probability of detection (P_D) and probability of false alarm (P_{FA}). A good detector maximizes the P_D while minimizing P_{FA}. In application, there are specific requirements for the performance of the detector, and a trade-off between maximizing the P_D and minimizing P_{FA} must be met. A visual tool for analyzing trade-offs is the receiver operating characteristic (ROC), which plots the P_D verses the P_{FA}. A ROC based on actual data is created by incrementing the thresholds used in the detector. At each step, the number of detections and false alarms is measured to estimate the P_D and P_{FA}. Each measurement fills in a part of the plot until a complete curve is made.

Figure 4.7 shows the performance of the cell-average and threshold detectors. The cell-average detector performs much better than the threshold detector because the cell-averaged detector is able to detect small decrease in coherence without generating false detections in the lower coherence area (e.g., the lower left corner of Figure 4.6).

Figure 4.8 shows how the multi-aperture method increases CCD performance. The performance is measured using the cell-averaged detector on combined coherence maps created from increasing amounts of sub-apertures. As more apertures are used to create the coherence map, the performance of the detector increases. The multi-aperture ROC also shows that adding additional apertures helps more in the beginning than later on. This is because after the large areas of low coherence have been removed by combining several sub-apertures, additional apertures have less of an effect.

While this example does not show actual CCD it does demonstrate several promising results. High coherence between repeat-passes at L-band is achieved using multi-aperture CCD. Changes in the phase introduced into one image are detected in the coherence map between images. However, care must be taken in attempting to extend these results to other environments. The flat desert landscape is less sensitive to spatial decorrelation than a mountainous area or an urban environment. Areas with thicker vegetation are more sus-

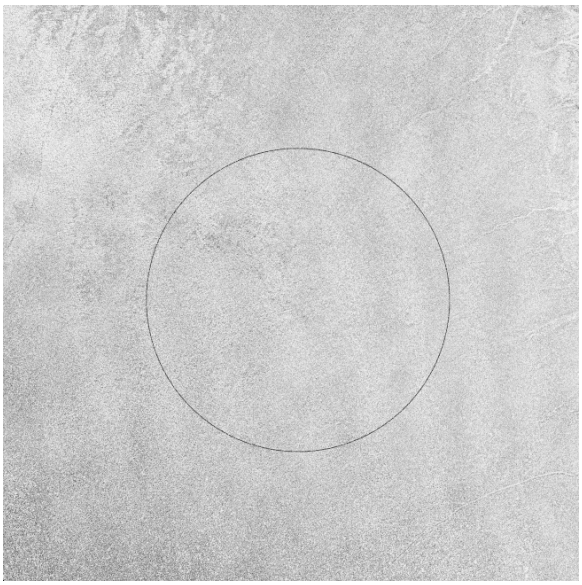

Figure 4.6: Multi-aperture coherence map with a simulated circular target added.

ceptible to temporal decorrelation. The next example illustrates CCD in a more challenging environment.

4.5 X-Band Example

Several passes were taken over a rural farming area at X-Band. While no specific ground verification of changes is available, there are several visible differences between passes. Specifically, a tractor plowed through a field and a irrigation pipe was rotated around a pivot (see Figure 4.9). The change in the roughness of the soil caused by plowing is enough to brighten the magnitude image of the field. Though the changes shown are detectable by noncoherent change detection, the principles necessary for coherent change detection still apply; specifically, there is a decrease in coherence between images as the result of man-made changes.

The images for this example are processed by dividing the 10° azimuth beam into eleven, overlapping 1° apertures. The native resolution of each sub-aperture image is 0.9 x 0.6 m (azimuth x range), while the pixels are spaced 0.5 m in both directions. The sub-apertures are relative to a common linear track as described in Section 4.2. Also the pulses have been resampled to a constant distance between pulses. This careful processing

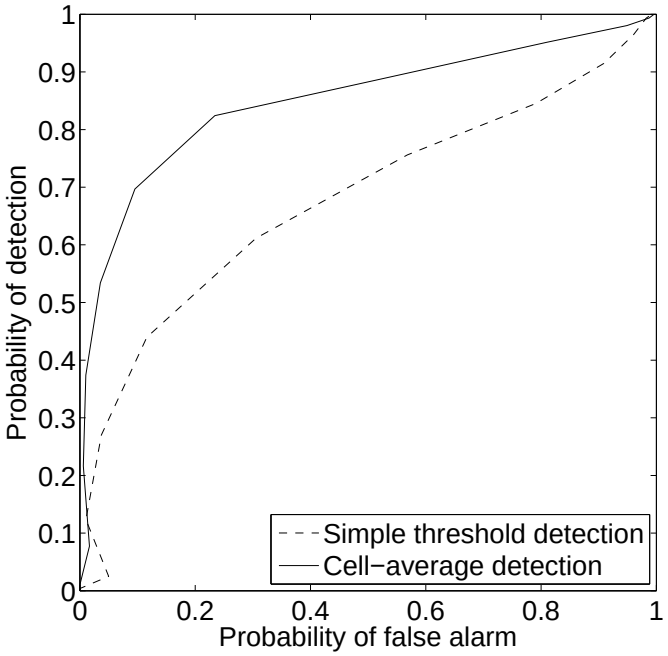

Figure 4.7: Performance of cell-average and threshold detectors on multi-aperture coherence maps.

is necessary to preserve the correlation between repeat-passes because it is more difficult to obtain correlated images at X-band.

CCD at X-band is difficult because many of the geometric sources of decorrelation depend on the wavelength as discussed in Section 2.3. The two flight lines at X-band must be much closer than at L-band (generally a 10 m tube is suggested [15]). However, X-band is also sensitive to smaller changes in the scene because of the smaller wavelength. A general rule is that a change of 0.2 times the wavelength will cause complete decorrelation [3]. X-band is therefore sensitive to changes on the order of several millimeters versus several centimeters for L-band.

Coherence maps for several sub-apertures are shown in Figures 4.10(a)-(c). For the forward-looking sub-aperture, there are times when one or both of the antenna are squinted backward so that no valid images are generated. This results in the blank parts of the image. There is also banding in the azimuth direction. This banding is likely caused by differences squint between passes that result from different antenna patterns being applied

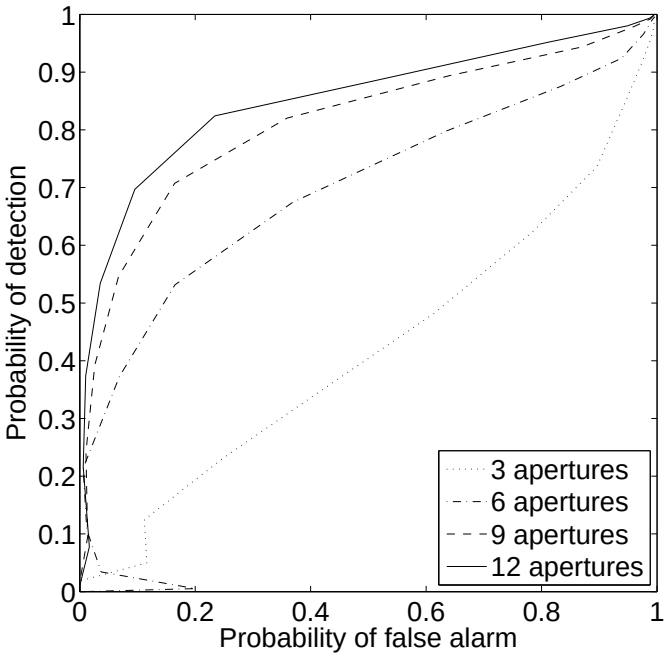

Figure 4.8: Performance of detectors using different numbers of apertures. Cell-averaged detection is used for each example.

to the incoming pulses. The low-coherence bands may also be the result of residual motion errors.

In the multi-aperture coherence map shown in Figure 4.10(d), many of the low coherence areas are removed. The changes caused to the plowed field are clearly visible in the coherence map as a arc-shaped, low-coherence area. The part of the field that was plowed in both images remains coherent between passes. The change in the sprinkler line is more difficult to see because it is in a low coherence part of the image. There are two slightly darker lines corresponding to the two sprinkler positions between passes.

The difficulty of CCD in an urban area is also demonstrated with this example. There are many pockets of low coherence corresponding to flat surfaces, and shadows. Many of these areas could be compensated for by masking out the low SNR areas common to both images, though this is not done here. Also of interest are the patches of low and high coherence that correspond to different fields in the magnitude images. Taller crops in one field relative to the other likely cause lower coherence areas because of increased volume decorrelation and more movement causing temporal decorrelation.

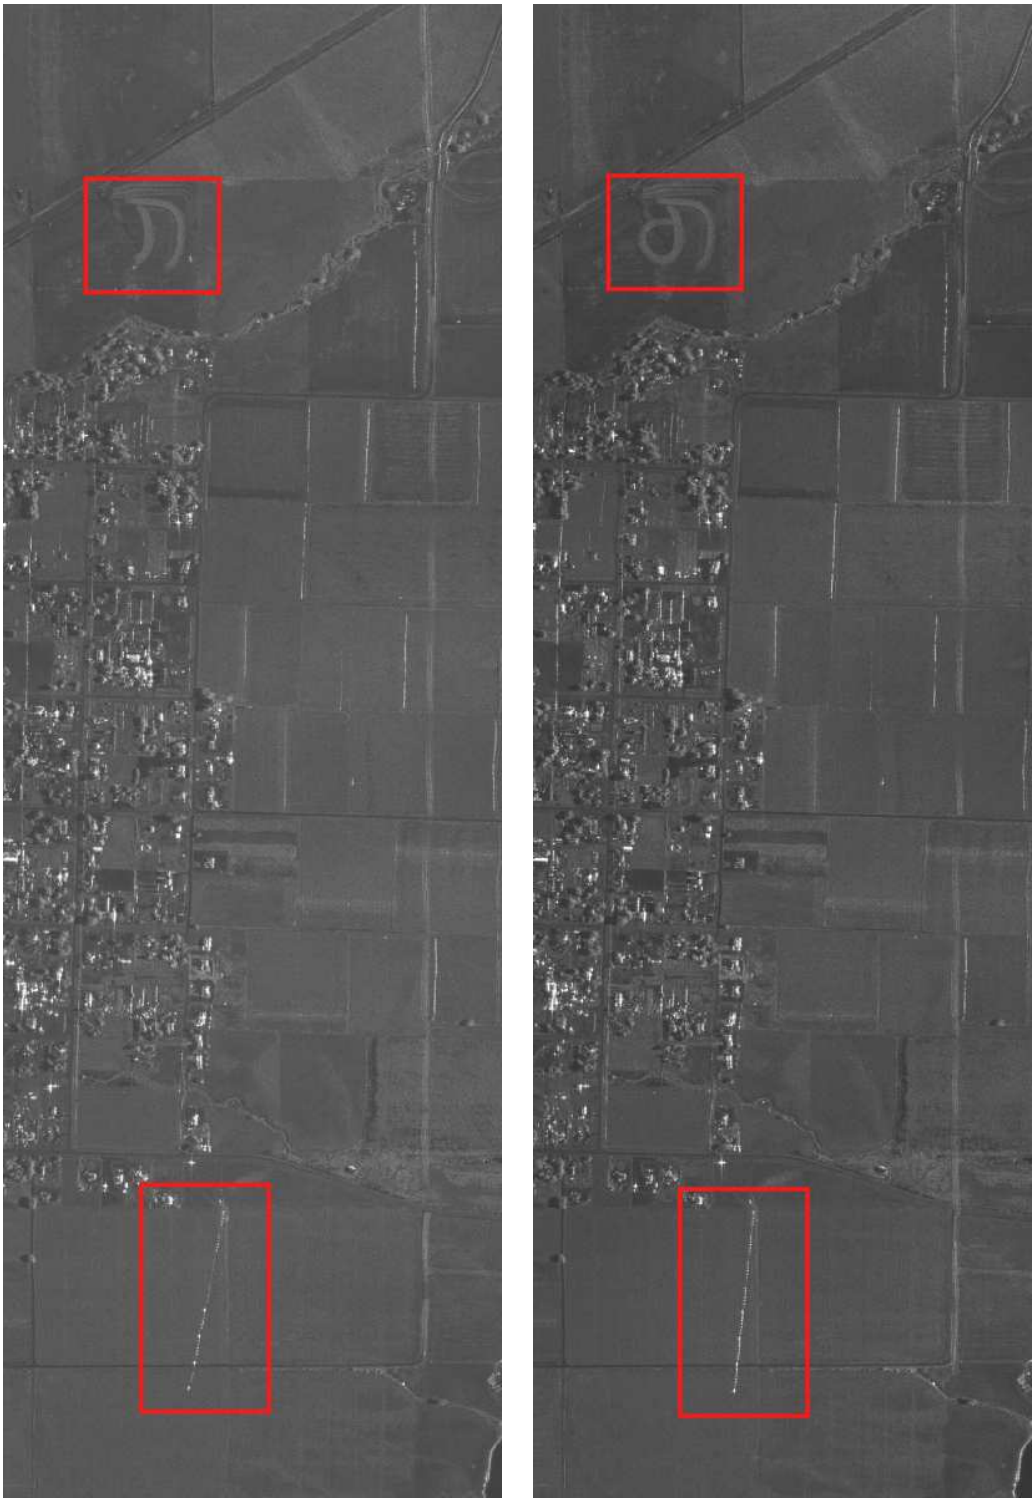

(a) (b)

Figure 4.9: (a) Master image (b) Slave image. Note the changes caused by the field being plowed at the top of the image and the rotation of a line of sprinkler pipes at the bottom of the image.

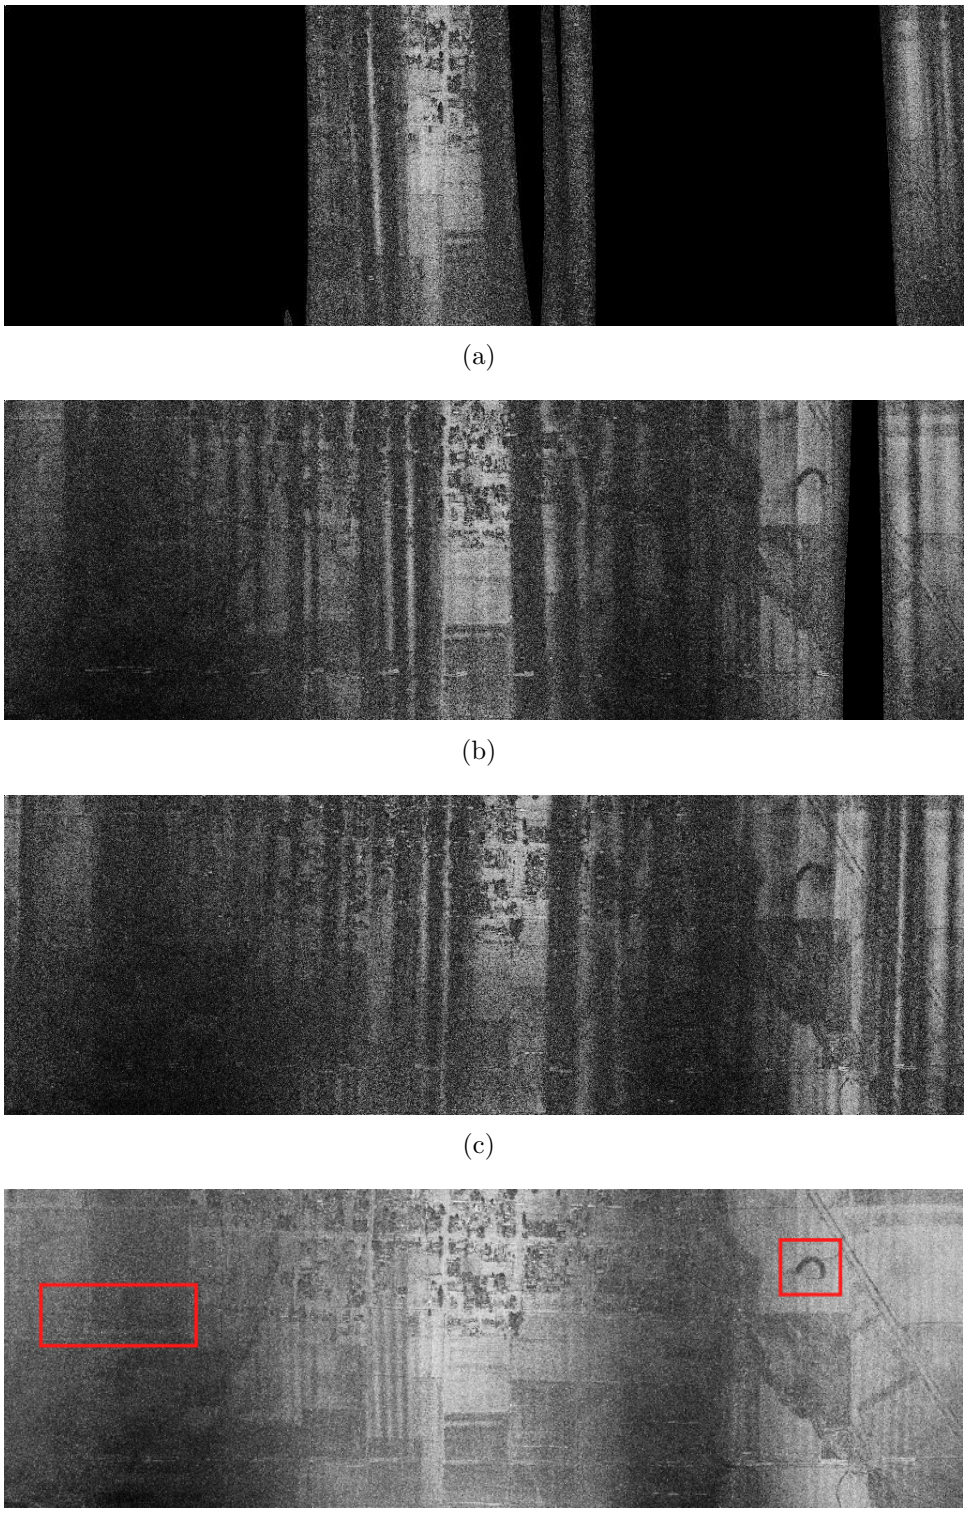

Figure 4.10: The coherence maps generated from (a) a forward-looking sub-aperture (b) the broadside sub-aperture and (c) a backward-looking sub-aperture. (d) is the coherence map generating by combining all sub-apertures.

The results of this X-band collection show the advantage of multi-aperture CCD. If a single, narrow-beam fixed antenna had been used the results would be similar to the broadside sub-aperture shown in Figure 4.10(b). However, to further increase the correlation between passes on a consistent basis, a gimbaled antenna and more accurate navigation control may be necessary.

4.6 Chapter Summary

This chapter shows several processing techniques that increase the coherence between repeat-pass images after they have been collected. Specifically, the multi-aperture CCD method is an innovative approach to utilizing a widebeam antenna to increase the coherence between images. High coherence between image collections has been shown at L-Band. In X-Band, a loss of coherence caused by man-made changes to a scene was shown.

Based on these examples, future development of CCD at L-Band shows promise because of the high coherence between images while being less sensitive to flight path differences. However, CCD at L-Band is less sensitive to small changes in a scene and may not be adequate for some targets of interest. For CCD at X-Band, the coherence between images needs to be improved by better processing and image registration or an improved ability to fly identical flights paths.

CHAPTER 5. INTERFEROMETRY

Interferometry uses the phase difference between two images to find the height of an area based on the collection geometry. This chapter outlines the steps necessary to form an interferogram. Differences in generating interferograms for imaged formed with backprojection are noted. A method of combining height estimates from multiple apertures is presented and applied in a repeat-pass interferometry example.

5.1 Generating a Repeat-pass Interferogram

The geometric differences between two flights are shown in Figure 5.1. The positions of two antennas are shown at an equivalent azimuth location corresponding to when the vector from airplane to pixel is orthogonal to the direction of flight, i.e., the pixel is at broadside. The ranges from each antenna to an arbitrary point on the ground are labeled as R_1 and R_2. Also, the distance between the two antennas, the baseline, is labeled as B. A mathematical model for interferometry is given as follows. Corresponding pixels in two SAR images, $f_1(x, y)$ and $f_2(x, y)$, are formed by

$$f_1(x,y) = W(x,y) * \left[\sigma(x,y) exp(-j\frac{4\pi}{\lambda} R_1(x,y))\right], \tag{5.1}$$

$$f_2(x,y) = W(x,y) * \left[\sigma(x,y) exp(-j\frac{4\pi}{\lambda} R_2(x,y))\right], \tag{5.2}$$

where $W(x, y)$ is the SAR impulse response (generally a two-dimensional sinc function), $\sigma(x, y)$ is the complex backscatter of the ground, $*$ is the convolution operator, and $R_1(x, y)$ and $R_2(x, y)$ are the ranges from each antenna to the point on the ground, (x, y). This model assumes the images are taken from close enough geometries that the ground backscatter is the same for both images.

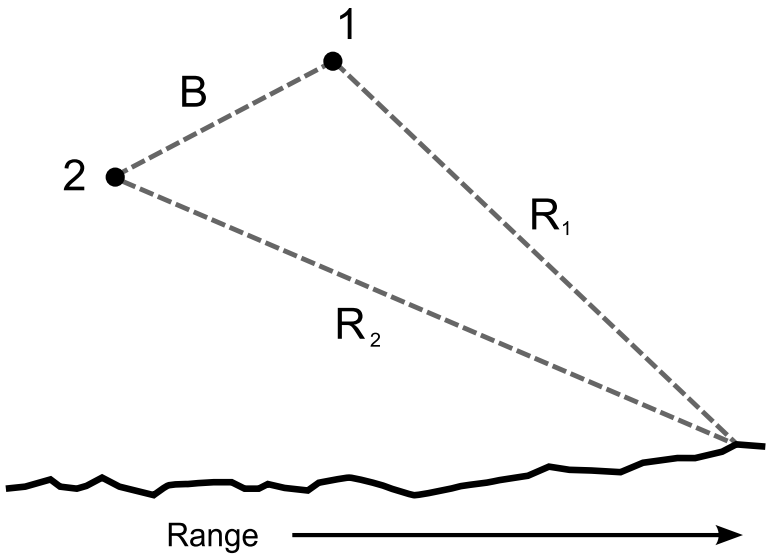

Figure 5.1: Basic geometry for interferometry. The position of antenna 1 and 2 are labeled, and the ranges to an arbitrary pixel on the ground is shown.

The interferometric phase between two pixels is generated by multiplying one image by the complex conjugate of the other and taking the phase of the complex result. Ideally, all phase terms cancel except for the range dependent phase terms. Mathematically, this is expressed as:

$$\begin{aligned}\phi(x,y) &= \angle\left(f_1(x,y)f_2(x,y)^*\right) \\ &= \frac{4\pi}{\lambda}\left(R_1(x,y) - R_2(x,y)\right) \\ &= \frac{4\pi}{\lambda}\delta R(x,y). \end{aligned} \tag{5.3}$$

Equation 5.3 gives the relation between the differential range, $\delta R(x,y)$, and the interferometric phase, $\phi(x,y)$. The differential range is related to the height by a linear equation,

$$\delta h(x,y) = a(x,y)\delta R(x,y). \tag{5.4}$$

Substituting Equation 5.3 for the differential range then gives,

$$\begin{aligned}\delta h(x,y) &= a(x,y)\delta R(x,y) \\ &= a(x,y)\frac{\lambda}{4\pi}\phi(x,y) \\ &= \lambda^*(x,y)\phi(x,y). \end{aligned} \tag{5.5}$$

Equation 5.5 is an expression for the height, $\delta h(x,y)$, at each pixel in the image as a function of the interferometric phase, $\phi(x,y)$, and the interferometric scale factor, $\lambda^*(x,y)$. The following sections discuss of these terms.

5.1.1 Differential Height

The height derived in interferometry depends on assumptions made in forming the complex images. When no assumptions are made about the height of the terrain being imaged, the differential phase increases as a function of range as shown in Figure 5.2. This increase in the differential range results in a range dependent phase ramp known as the flat-earth phase.

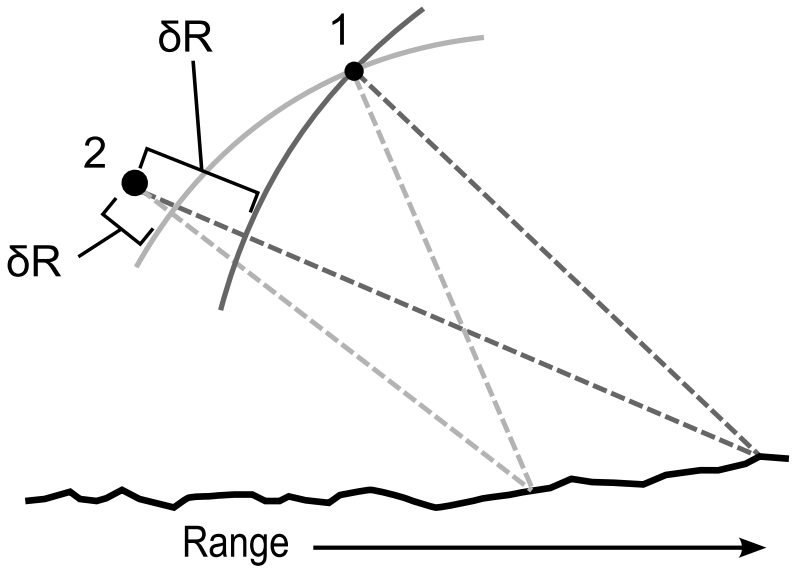

Figure 5.2: The differential range, δR, increases as areas farther from the antennas are imaged.

A example of the simulated interferometric phase over an area is shown in Figure 5.3. The phase in the image is displayed modulo 2π. The vertical bands are caused by resticting the ramping flat-earth phase with the modulo 2π. Small aberations in the bands represent a portion of the interferometric phase dependent on the terrain. By assuming that the earth is flat and using measurements of the antennas' position, the flat-earth phase is estimated (see Figure 5.4). Removing the estimated flat-earth phase from the interferometric phase results in phases that are a direct function of the terrain height as shown in Figure 5.5. When the interferometric phase is scaled, the resulting differential heights show the difference between the actual terrain and the assumed flat earth.

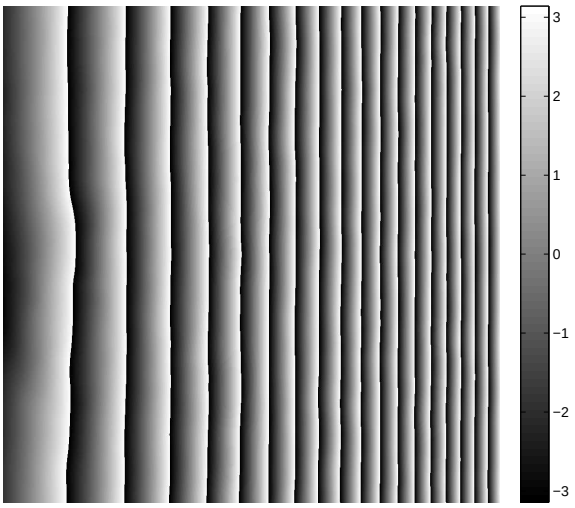

Figure 5.3: Simulated interferogram showing the flat-earth phase and small aberations corresponding to terrain.

For image formation with backprojection, the range between the aircraft and each pixel is computed using a digital elevation model (DEM) and motion measurements of the aircraft position. Assuming a DEM instead of a flat surface changes the interferometric

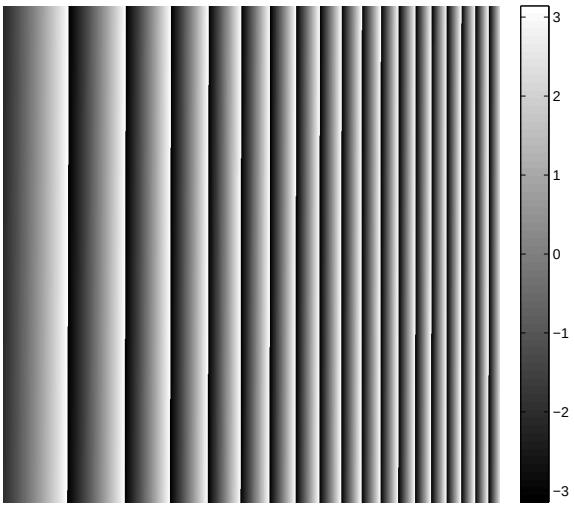

Figure 5.4: The phase due to a flat earth used to remove the flat-earth phase from an interferogram.

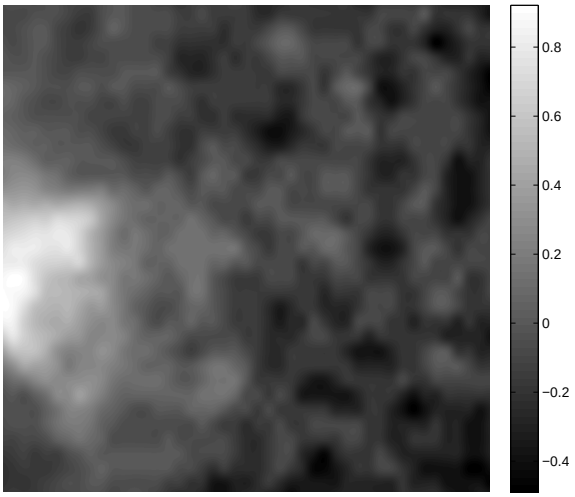

Figure 5.5: The interferogram obtained by removing the flat-earth phase to leave only the phase caused by the terrain.

phase given in Equation 5.3 to

$$\begin{aligned}\phi(x,y) &= \angle\left(f_1(x,y)f_2(x,y)^*\right) \\ &= -\frac{4\pi}{\lambda}\left((R_1(x,y)-\hat{R}_1(x,y))-(R_2(x,y)-\hat{R}_2(x,y))\right) \\ &= -\frac{4\pi}{\lambda}\delta R(x,y). \end{aligned} \tag{5.6}$$

The computed ranges based on the DEM are $\hat{R}_1(x, y)$ and $\hat{R}_2(x, y)$. This equation leads to a differential height that is the difference between the actual terrain and the DEM used to process the two images. If an ideal DEM were used backprojection would completely compensate for the expected phase and there would be no phase difference between images. There is also no flat-earth phase in backprojection images because the assumed DEM accounts for the increase in differential range in the range direction.

5.1.2 Interferometric Phase

Obtaining an accurate height estimate requires an accurate measurement of the interferometric phase. Errors in the phase between images are caused by noise and wrapping of the phase. This section discusses the steps necessary to reduce these problems.

Phase noise in SAR images is dominated by receiver noise. Reciever noise is reduced by averaging the phase over several pixels. As noted in Section 4.1, the averaged interferometric phase is the phase of the cross-correlation between images. As an example consider Figure 5.6 which shows the original phase difference between images and the phase after the pixels have been averaged.

The discontinuities in Figure 5.6 are caused by the phase wrapping between values of $-\pi$ and π. A continous phase surface is made by adding multiples of 2π to areas separated by the discontinuities. Accurate phase unwrapping in two-dimensions is covered extensively in [22]. Figure 5.7 shows the average interferogram before and after it has been unwrapped.

Care must be taken in choosing the appropriate method and order of smoothing and unwrapping the phase. If a filter is applied to the detected phase between images before it has been unwrapped, the discontinuities will be smeared. Phase unwrapping algorithms require sharp jumps in phase to operate correctly and will fail if the jumps have been smoothed. However, if the averaging filter is used on complex data, before detecting the phase, then the discontinuities will remain sharp. After the phase has been unwrapped, further filtering can be performed without causing errors.

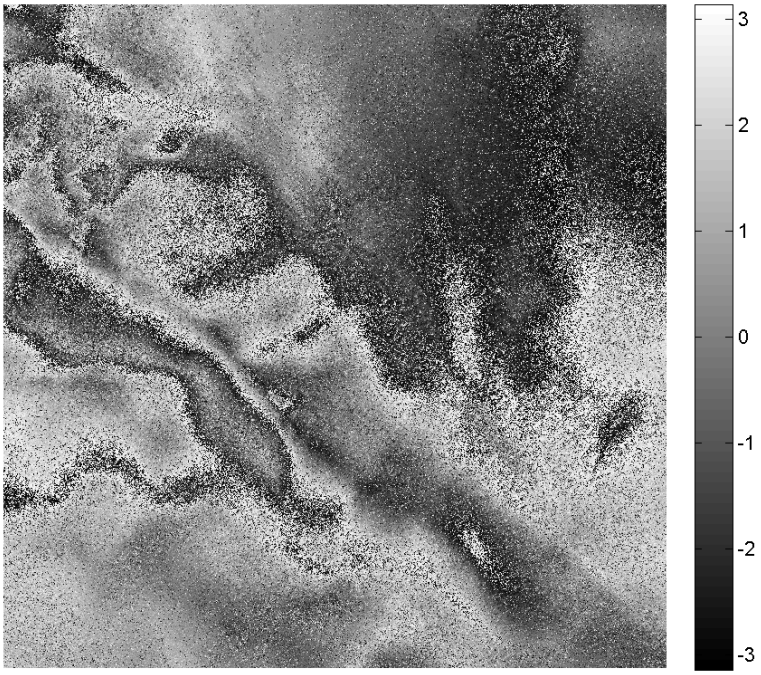

Figure 5.6: Phase difference between images without averaging or unwrapping.

5.1.3 Interferometric Scale Factor

The interferometric scale factor converts the phase to a height value. The scale factor, $\lambda^*(x, y)$, depends on the geometry between antenna positions and the pixel being imaged. In Figure 5.8, several angles and distances used to create an expression for the scale factor are labeled. In the literature, several expressions for the interferometric scale factor are given. In [16], the scale factor is given as

$$\lambda^*(x, y) = \frac{\lambda R_1}{2\pi B}[\sin(\alpha) - \cos(\alpha)\tan(\alpha - \theta)], \tag{5.7}$$

and in [2] the scale factor is

$$\lambda^*(x, y) = \frac{\lambda}{4\pi} cos(\psi_1) \frac{1}{\psi_2 - \psi_1}. \tag{5.8}$$

Both of these equations make different trigonometric approximations and assumptions about the registration between images. Both equations have been used to produce height estimates

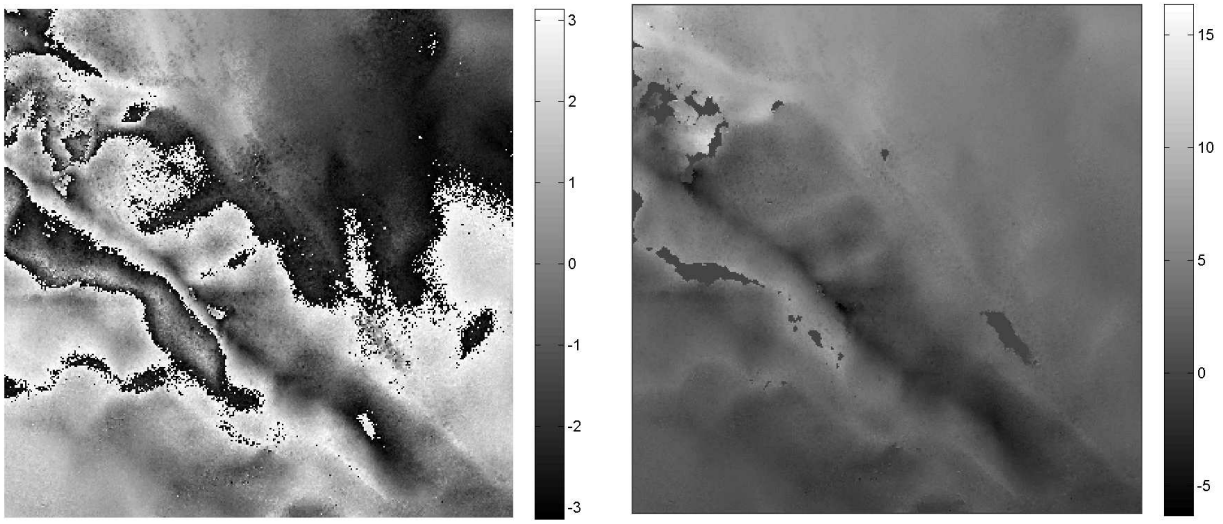

(a) (b)

Figure 5.7: (a) Averaged interferogram before phase unwrapping. (b) Unwrapped interferogram.

that are roughly equevalent. In order to better compare the performance between these scale factors, a careful collection with calibrated targets is needed.

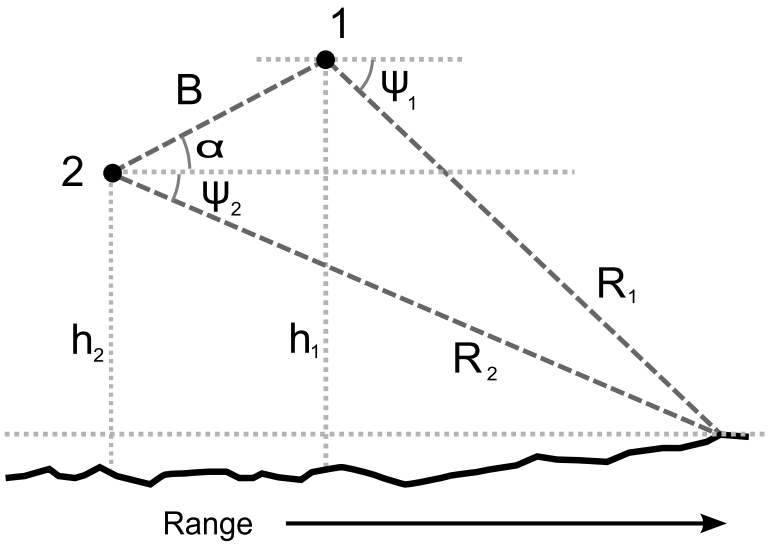

Figure 5.8: Various angles and heights used to express the interferometric scale factor.

The angles and distances used to calculate the scale factor are relative to an assumed pixel height and measured aircraft position. Once an improved height estimate is obtained

it can be used to find a more accurate scale factor and reprocess better images. However, errors in these measurements bias the scale factor. Also, incorrect positions result in images with phase and registration errors.

For repeat-pass interferometry, the baseline changes throughout the flight. In calculating the scale factor for each pixel, the baseline is the distance between flight paths when the paths are broadside to the pixel. Figure 5.9 shows how broadside is defined relative to a straight reference path. Over the synthetic aperture the baseline varies, but it is assumed the variation is small. This assumption is valid when a narrow beam or sub-aperture is used to form the pixel.

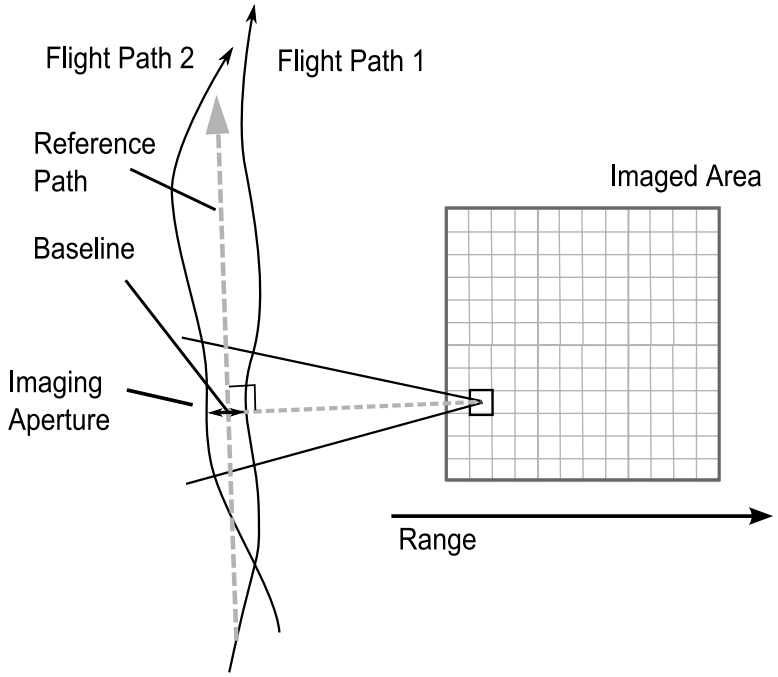

Figure 5.9: Changing interferometric baseline between two flights. The baseline for a pixel is defined at the broadside location relative to the reference path.

In the next section, non-broadside apertures are used to form interferograms. Complex images corresponding to squinted apertures are created by restricting the range of angles used to form each pixel as discussed in Section 4.2. Instead of using the baseline at broadside, a baseline is chosen that is in the middle of the sub-aperture range of angles.

5.2 Interferometry with Multiple Sub-apertures

I have developed a new method for interferometry based on the multi-aperture CCD method described in Section 4.3. In the following I present this method and several new variants. Multi-aperture interferometry splits a wide antenna beam into multiple sub-apertures. The interferometric phase between images for each sub-aperture is filtered, unwrapped, and scaled as outlined in Section 5.1. The different height estimates are then combined to provide a more accurate height map than would be possible from a single aperture.

There are several ways to combine the height estimates from sub-apertures. Each method has different advantages and drawbacks as outlined below.

Average Height For this method, all height estimates are averaged to reach an estimate for the true height. Given K sub-apertures, the following average is performed for each pixel,

$$\delta h(x,y) = \frac{1}{K}\sum_k \delta h_k(x,y). \tag{5.9}$$

By combining independent estimates of the height, errors associated with zero-mean, gaussian noise (thermal noise in the receiver) are reduced.

This method has several disadvantages. One arises from weighting the height estimates from each sub-aperture equally. Often the sub-apertures further from broadside have a lower SNR resulting in noisy height estimates. Biases introduced in the calculation of the interferometric scale factor are often common to multiple apertures and, therefore, do not average away. Also, the assymptotic behaviour of the scale factor results in sharp spikes in the height estimates which do not average away.

Coherence Weighted The height estimates from each aperture are averaged using a weighting that depends on the coherence of the pixel. For each pixel the height estimate is

$$\delta h(x,y) = \frac{1}{K}\sum_k \gamma_k(x,y)\delta h_k(x,y). \tag{5.10}$$

Often low coherence areas are associated with low SNR sections of the image or residual motion errors in the recorded flight track. Areas of low SNR, as found in apertures far from broadside or in areas of shadow, produce noisy estimates of the height. Residual motion errors result in a bias in the height estimate because the scale factor calculated from the recorded flight track is not correct. In both cases, reducing the contributions of these areas results in a more accurate estimate of the height.

Weighting by the coherence removes several of the disadvantages of an equally weighted average. However, it does not completely remove biases in the scale factor or the assymptotic behavior of the scale factor. In fact, areas where the flight tracks align, causing a large scale factor, are often the most coherent areas because there is no baseline decorrelation.

Minimum Scale Factor In order to remove the assymptotic behavior of the scale factor, a third method is suggested. The height estimate that has the smallest scale factor is chosen as the correct height, as given by

$$l = \arg\min_k \lambda_k^*(x, y), \tag{5.11}$$

$$\delta h(x, y) = \delta h_l(x, y), \tag{5.12}$$

where the aperture, l, with the minimum scale may be different at each pixel, (x, y). This method produces a height estimate with relatively few discontinuities and no large assymptotic areas.

This method still has several disadvantages. Because only one height estimate is used there is no reduction in height errors through averaging. At times, areas of low coherence have the smallest scale factor, which results in poor height estimates. The minimum scale factor may still be biased as well.

In the next section, the minimum scale factor method is used because it produces large areas without discontinuities. The discontinuities are along straight lines creating strips of smooth terrain. A combined method that performs a coherence-weighted average over multiple, low-scale-factor regions was attempted but led to irregulare lines of discontinuities.

A more accurate method of calcualting the scale facter would likely reduce the discontinuities, making a combined method optimal.

5.3 Multi-aperture Interferometry Example

Repeat passes were flown over a mountainous area at L-Band. From these flights, two highly coherent passes where chosen and used in this example. There is no set baseline between the passes. The paths often cross in latitude and longitude as well as in altitude, resulting in varying baselines and incidence angles. A magnitude image of the scene is shown in Figure 5.10. The plane flew south-west across the top-left corner of the image, illuminating a diagonal swath. For this example, height maps are generated from the marked area.

A 20° antenna beam was used to collect the data. The beam is processed into five sub-apertures, each 6° wide. The height maps for four apertures are shown in Figure 5.11. Portions of the DEM are visible in the unimaged area and other parts that are masked out. The additional heights add greater detail in portions of the height map. Specifically, the peaks of the hills are sharper and the valleys are deeper. Also, many little ravines, which are too small to show up in the coarse DEM, are now visible. Figure 5.12 shows the portion of the terrain in greater detail along with the original DEM. The sharper peaks, deeper valleys, and ravines are consistent with the features seen in the optical image of the same area. The height maps have been rotated and lit to match the aspect angle and lighting of the optical image.

At times the flight paths align to cause large values for the scale factor. The interferometric heights at these times are masked out and only the DEM is shown. Because the same section of the flight path is processed at different times for each height map, the masked areas are at different locations in each sub-aperture. The masked areas have similar shapes which slide over between apertures corresponding to the different times that the same flight track is processed. One of these sliding areas is noted with arrows in Figure 5.11.

Several challenges to interferometry are visible in Figure 5.12(a). The boxed area is masked out because it would not unwrap correctly. The interferometric phase in this region is noisy because the terrain is very rocky leading to layover during image formation. The circled region highlights a bias in the range direction. The height is much less than the

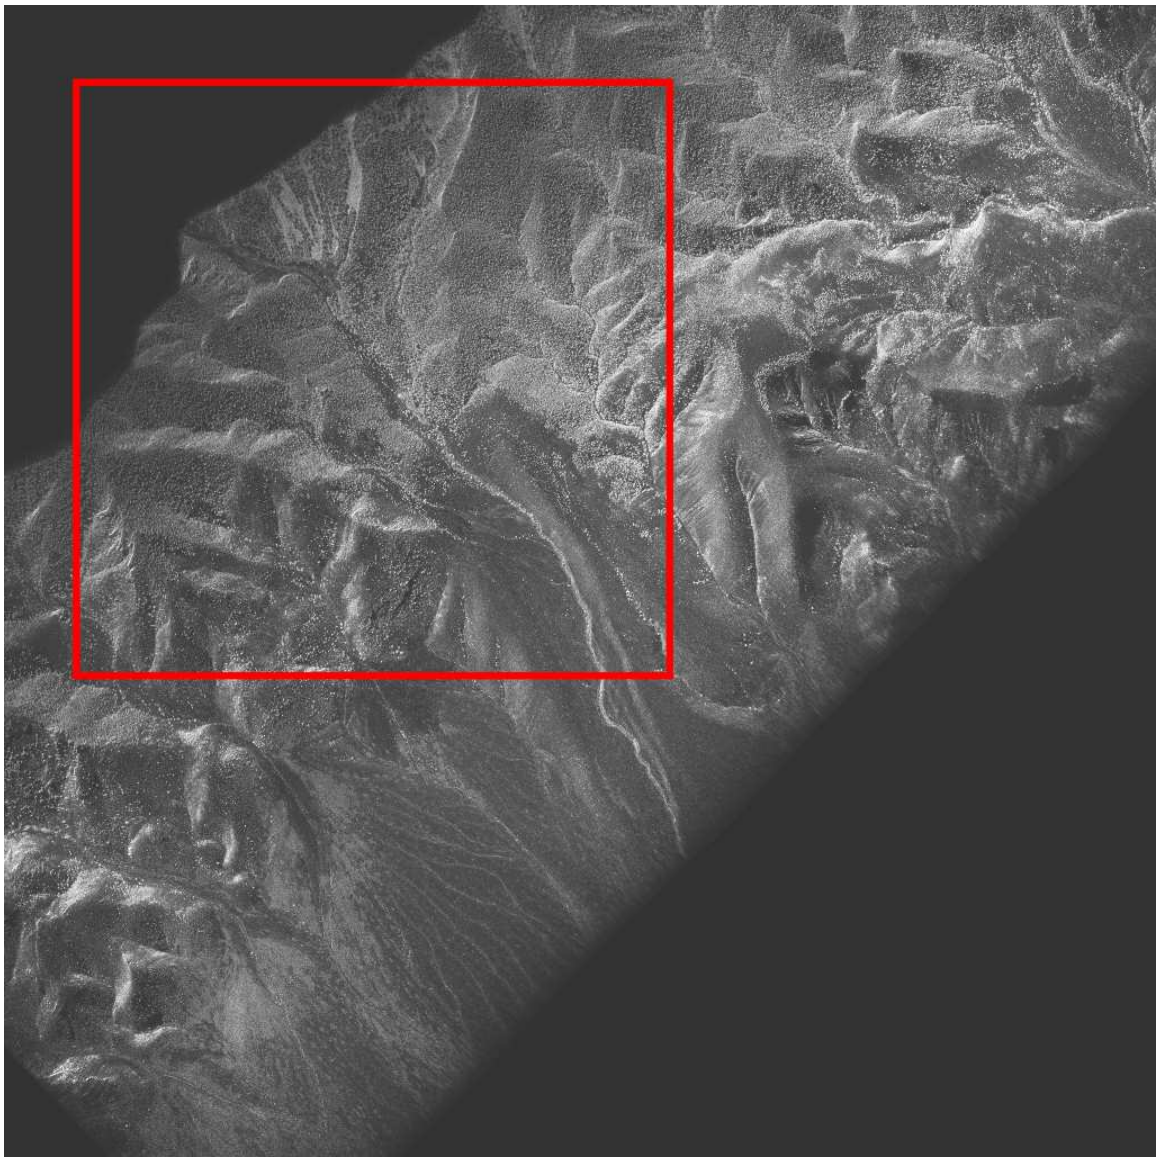

Figure 5.10: Radar image of area used for repeat-pass interferometry example. The plane flew south-west across the upper-left corner of the image.

DEM in this portion of the near range. In the far range, the new heights are larger than the DEM. The bias is likely caused by an inaccuracy in calculating the interferometric scale factor. It may be that Equation 5.5 makes an approximation that is valid for the small baselines of single-pass interferometry but breaks down for the larger baselines associated with repeat-pass interferometry.

Figure 5.12(a) also shows that the DEM has a patch-like pattern caused by the interpolation method. The DEM is generated by interpolating heights for every pixel based

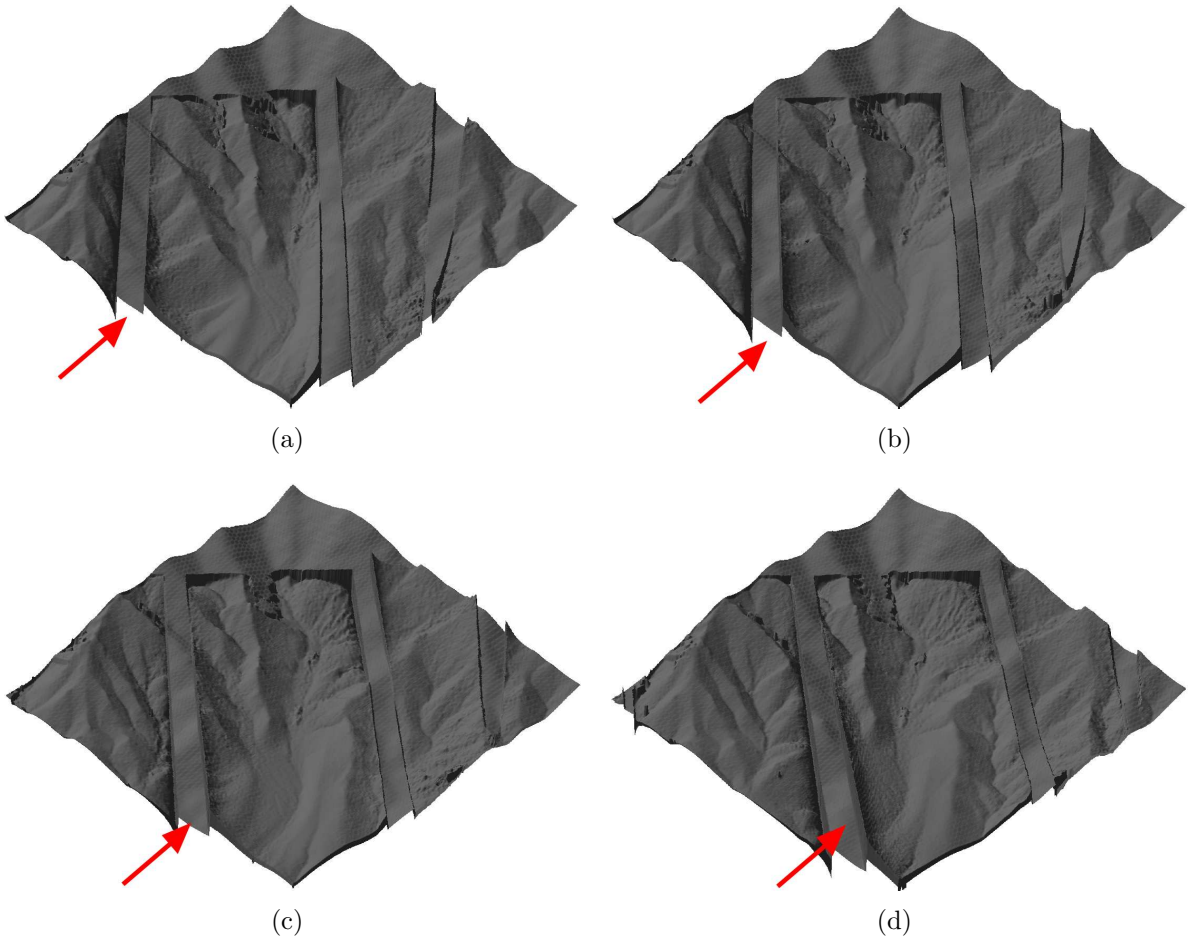

Figure 5.11: Height maps generated by repeat-pass interferometry for four of five sub-apertures. (a) is a forward-looking aperture, (b) is the broad-side aperture, and (c) and (d) are backward-looking apertures. Highlighted are the masked out areas of invalid interferometric heights as discussed in the text.

on coarsly spaced postings. When the interferometric heights are accurate they remove this patch-like pattern from the surface leaving only the finer texture of the actual ground.

Figure 5.13(b) shows the combined height map. The minimum scale factor method is used to combine the height maps from each aperture. There are now no areas masked-out because of aligned flight paths. Using the aperture with the minimum scale factor creates small discontinuities in the height map but leaves large continuous patches.

(a)

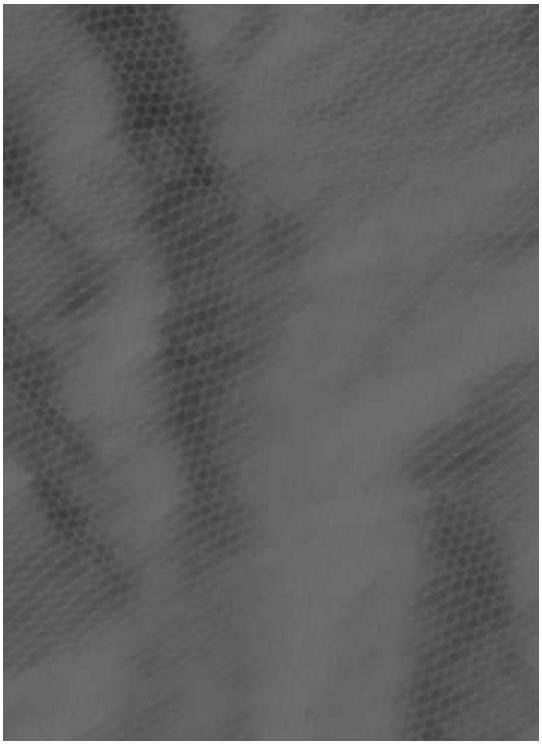

(b)

(c)

Figure 5.12: (a) Portion of one height map enlarged to show detail. (b) The DEM for the same area. (c) Picture taken from the same aspect angle as the height maps.

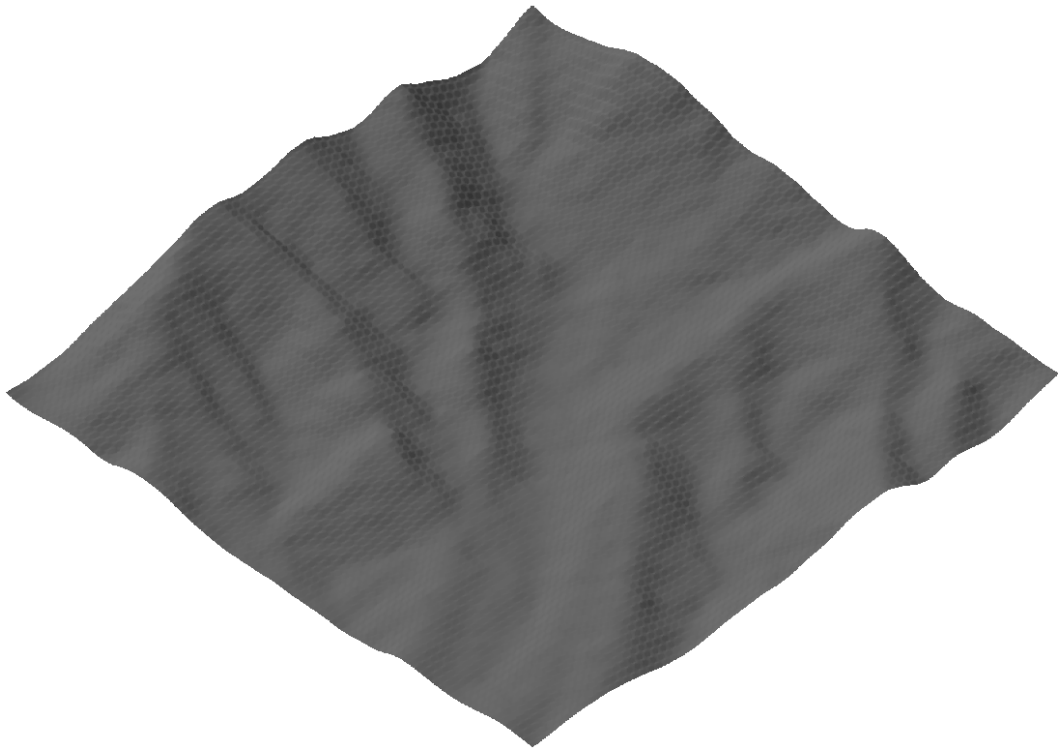

(a)

(b)

Figure 5.13: (a) Height map of the DEM used in processing. (b) Height map generated by combining the interferograms from five sub-apertures.

5.4 Chapter Summary

In this chapter, it is shown that the interferometric height generated from backprojection images is the difference between the true terrain and the DEM used in processing. Previous methods of unwrapping and scaling the interferometric phase are applied to repeat-pass images resulting in height estimates that add sharpness and detail to the DEM.

A new method of using multi-apertures to form an interferogram is presented. This method is capable of replacing sections of invalid heights caused by crossing flight paths with valid heights. Though it is not implemented, the possibility of reducing noise in the height estimate by averaging independent height estimates generated from different sub-apertures is also noted.

Some biases are seen in the interferometric heights. These may be caused by the inaccurate measurements of the flight paths or invalid assumptions in deriving the interferometric scale factor. Future work could develop a more accurate scale factor. Interferometric simulations and, eventually, a test with highly calibrated targets would be necessary to properly measure improvements of multi-aperture interferometry over single aperture interferometry.

CHAPTER 6. CONCLUSION

This thesis covers work done to implement coherent change detection (CCD) and repeat-pass interferometry for the NuSAR system. In the process, a great deal about accurate SAR processing has been learned and innovative work in CCD and interferometry has been accomplished.

Because SAR is a coherent imaging system, it is sensitive to small changes in phase caused by changes in the range from the radar to the target. CCD uses phase differences to detect small changes in a scene and interferometry uses phase differences to estimate heights in a scene. Both CCD and interferometry require two complex images which have been collected and processed in a such a way as to preserve an accurate phase relation and correlation between the images.

Images can be created incorrectly due to errors in the position of the plane or unknown terrain on the ground. Registration algorithms shift, stretch, and align two formed images to reduce these errors. The maximum correlation method is a common way of warping one image to another so that the correlation between images is maximized.

CCD is performed on two images of the same area collected from very similar flight lines. In processing, several steps are taken to ensure the matching images are coherent. Backprojection processing is especially adept at compensating for different collection geometries during image formation. Multi-aperture CCD uses several apertures to increase the coherence in a scene.

For interferometry, knowledge of the imaging geometry combined with the phase relation between two images is used to find the heights of an area. Multiple apertures are used to produce independent estimates of the heights which can be combined in different ways. For repeat-pass interferometry, portions of the flight path produce areas of invalid heights in

single-aperture interferograms. Combining height estimates from multiple apertures removes these areas.

6.1 Contributions

A number of new contributions to the public literature on SAR processing, CCD, and interferometry are made by this thesis. These are specifically highlighted as follows:

Coherent Processing A collection of steps to increase the coherence between images is given: (a) use backprojection processing to create apertures relative to a common reference track, (b) resample the raw data to adjust for differences in velocity, and (c) account for the antenna pattern. These steps specifically increase the coherence for fixed antenna SAR systems, which require less cost and complexity than systems with gimbal-mounted antennas.

Multi-Aperture CCD A new method that uses multiple sub-apertures to perform CCD is presented. Multi-aperture CCD is able to compensate for low coherence sections likely caused by errors in the recorded flight track. Also, the overall coherence of a scene is increased. As with the coherent processing steps, multi-aperture CCD reduces the hardware requirements for the SAR sensor.

Backprojection Interferometry It is shown that interferometry with backprojection yields the difference between the elevation model used and the true elevation.

Multi-Aperture Interferometry This thesis discusses several ways of combining independent height estimates from multiple apertures into a single improved height map. An example is shown of how multiple apertures are combined to improve the coverage of a repeat-pass interferometric flight.

6.2 Future Work

This thesis suggests several new avenues of research into CCD and interferometry. Work in these areas adds new capabilities to SAR systems beyond imaging. Further work

offers the possibility of increasing the usefulness of fixed antenna SAR systems through processing instead of through hardware modifications.

One area of future work is in backprojection processing. Backprojection processing relies on knowing the range from the aircraft to the target at each pulse. It is shown that an error in the location of the target (the elevation assumed in processing) leads to an interferometric phase between images. A more complete model that traces the phase in backprojection processing could be developed. Such a model would account for errors caused by inaccurate assumptions of the ground elevation and the plane's position. Understanding gained from this model would help in developing registration and autofocus techniques tailored to backprojection.

Further research into multi-aperture CCD would result in improved coherence maps ready for autodetection routines. The coherence could be increased by using higher accuracy motion data to lessen residual motion errors. Robust registration algorithms would also remove the effects of motion errors. In order to fine-tune coherence estimation to specific targets, a change detection data set (multiple repeat-pass flights with well documented ground changes) is necessary.

The usefulness of multi-aperture interferometry relies on developing a better interferometric scale factor. A scale factor that accommodates for a varying baseline and squinted geometry is necessary. Multi-aperture interferometry applied to a single pass would also be useful. Because single-pass interferometry has a fixed baseline, correlation between images is very high and there are no unusable portions of the flight when the antennas directly align. Finally, when combining interferograms from multiple apertures, a focus on improving the height estimation quality with independent samples would be desirable.

REFERENCES

[1] L. Ulander, H. Hellsten, and G. Stenstrom, "Synthetic-aperture radar processing using fast factorized back-projection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 3, pp. 760–776, July 2003. viii, 2, 18, 19

[2] C. V. Jakowatz, Jr., D. E. Wahl, P. H. Eichel, D. C. Ghiglia, and P. A. Thompson, *Spotlight-Mode Synthetic Aperure Radar: A Signal Processing Approach.* Kluwer Academic Publishers, 1996. 2, 8, 18, 25, 29, 35, 73

[3] M. Preiss and N. J. S. Stacy, "Coherent change detection: Theoretical description and experimental results," Defence Science and Technology Organisation, Tech. Rep., 2006. 2, 25, 61

[4] R. Scheiber and A. Moreira, "Coregistration of interferometric SAR images using spectral diversity," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 38, no. 5, pp. 2179–2191, Sep 2000. 2, 55

[5] K. de Macedo, R. Scheiber, and A. Moreira, "An autofocus approach for residual motion errors with application to airborne repeat-pass SAR interferometry," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 46, no. 10, pp. 3151–3162, Oct. 2008. 2, 55

[6] H. Zebker and J. Villasenor, "Decorrelation in interferometric radar echoes," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 30, no. 5, pp. 950–959, Sep 1992. 2, 11, 16

[7] A. W. Doerry, "Collecting and processing data for high quality CCD images," Sandia National Laboratories, Tech. Rep., 2007. 2, 35

[8] W. G. Carrara, R. S. Goodman, and R. M. Majewski, *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms.* Artech House, 1995. 2, 18

[9] I. G. Cumming and F. H. Wong, *Digital Processing Of Synthetic Aperture Radar Data: Algorithms And Implementation.* Artech House, 2004. 2

[10] D. Stevens, I. Cumming, and A. Gray, "Options for airborne interferometric SAR motion compensation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 33, no. 2, pp. 409–420, Mar 1995. 2

[11] D. Thompson, D. Arnold, and D. Long, "YINSAR: a compact, low-cost interferometric synthetic aperture radar," *IEEE International Geoscience and Remote Sensing Symposium*, vol. 1, pp. 598–600 vol.1, 1999. 3

[12] E. C. Zaugg, D. L. Hudson, and D. G. Long, "The BYU microsSAR: A small, student-built SAR for UAV operation," in *IEEE International Geoscience and Remote Sensing Symposium*, July 2006. 3

[13] M. C. Edwards, "Design of a continuous-wave synthetic aperture radar system with analog dechirp," Master's thesis, Brigham Young University, 2009. 3

[14] M. A. Richards, *Fundamentals of Radar Signal processing.* McGraw-Hill, 2005. 5, 11

[15] S. Hensley, K. Wheeler, J. Hoffman, T. Miller, Y. Lou, R. Muellerschoen, H. Zebker, S. Madsen, and P. Rosen, "Status of a UAV SAR designed for repeat pass interferometry for deformation measurements," Jet Propulsion Laboratory, Tech. Rep., 2004, available as http://hdl.handle.net/2014/40738. 11, 21, 61

[16] A. E. Robertson, "Multi-baseline interferometric SAR for iterative height estimation," Master's thesis, Brigham Young University, 1998. 21, 73

[17] F. Li and R. Goldstein, "Studies of multibaseline spaceborne interferometric synthetic aperture radars," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 28, no. 1, pp. 88–97, Jan 1990. 29

[18] F. Bookstein, "Principal warps: thin-plate splines and the decomposition of deformations," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 11, no. 6, pp. 567–585, Jun 1989. 30, 32

[19] A. Goshtasby, "Registration of images with geometric distortions," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 26, no. 1, pp. 60–64, Jan 1988. 30, 31, 33

[20] R. M. Gray and L. D. Davisson, *An Introduction To Statistical Signal Processing.* Cambridge University Press, 2004. 49

[21] P. Prats, A. Reigber, and J. Mallorqui, "Topography accommodation during motion compensation in interferometric repeat-pass SAR images," in *IEEE International Geoscience and Remote Sensing Symposium*, vol. 1, July 2005, pp. 4 pp.–. 55

[22] D. C. Ghiglia and M. D. Pritt, *Two-Dimensional Phase Unwrapping: Theory, Algorithms, and Software.* Wiley-Interscience, 1998. 72