Height Reconstruction Techniques for Synthetic Aperture Lidar Systems [1,2]

Curtis W. Chen and Scott Hensley
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109-8099
818-354-8034
curtis.chen@jpl.nasa.gov

Abstract—We examine height-reconstruction techniques for airborne synthetic aperture lidar (SAL) sensors imaging the ground surface. Techniques related to amplitude-comparison monopulse are least sensitive to the phase instability problems associated with the short system wavelengths and would be easiest to implement (through the use of a focal-plane detector array), so we deal primarily with amplitude-based techniques. For pair-wise monopulse processing, we present an analysis of the height sensitivity of the SAL system in terms of the system parameters. We also provide simulation results in order to verify our expressions for the system performance. The well-known bias issues associated with monopulse processing imply the need for high carrier-to-noise ratios. The height-estimation accuracy of the system will also be affected by phenomena related to the vertical distribution of targets and target motions; we give an overview of these effects as well. Finally, we propose an array-based joint processing approach that may be applied instead of pair-wise monopulse processing. We show that the proposed approach represents the maximum-likelihood estimator for obtaining the target elevation angle, and hence the target height, and we demonstrate that the proposed approach significantly reduces bias-induced errors.

TABLE OF CONTENTS

1. INTRODUCTION 1
2. OVERVIEW AND SYSTEM MODEL 2
3. SENSITIVITY AND PERFORMANCE 3
4. OTHER EFFECTS 5
5. ARRAY PROCESSING 6
6. CONCLUSIONS 8
ACKNOWLEDGEMENT 8
REFERENCES 8
BIOGRAPHIES 9

1. INTRODUCTION

The data-processing techniques and acquisition modes of a synthetic aperture lidar (SAL) instrument operating at optical wavelengths are closely related to the analogous modes of a synthetic aperture radar (SAR) instrument operating at microwave frequencies. It is consequently natural to explore the applicability of SAR processing techniques to SAL sensors. In this paper, we examine the feasibility of adopting SAR height-reconstruction techniques to the context of an airborne SAL sensor for the purpose of obtaining high-resolution 3-D imagery of the ground at optical wavelengths.

Interferometric SAR (InSAR or IFSAR) techniques are commonly used to obtain very accurate digital elevation models of land surfaces. Such techniques could in principle be used in the context of SAL instruments as well, but it is unlikely that a foreseeable SAL system could feasibly meet the stringent requirements on phase accuracy and stability required for precision interferometry. In particular, because of the short wavelengths of SAL sensors, interferometric SAL measurements would be highly sensitive to errors arising from uncompensated phase-center motion, atmospheric turbulence, target decorrelation, and target motion or vibration. A SAL interferometer would also require multiple phase centers, which might be costly to implement because of the need for dual optical chains in the receive system.

A more practicable option for SAL height reconstruction might be the use of techniques related to amplitude-comparison monopulse processing. Assuming that the SAL receiver utilizes a (coherent) focal-plane detector array in elevation, the output of each detector corresponds to a different receive-beam pattern. This scenario is analogous to the microwave case of a reflector antenna with multiple offset feeds. The data streams from the SAL detector channels can therefore be processed jointly or in pairs to estimate the elevation angle from which the dominant energy for a particular range-azimuth image pixel arrives. Such processing allows the elevation angle to be estimated to a fraction of the diffraction-limited beamwidth

[1] 0-7803-8155-6/04/$17.00© 2004 IEEE
[2] IEEEAC paper #1516, Version 1, Updated October 2, 2003

λ/D, where λ is the wavelength and D is the diameter of the real receive aperture. (Note that the elevation angle estimates from the real-beam 3-D image are limited to an accuracy of about λ/D.)

The improvement in accuracy over λ/D performance is dependent upon the carrier-to-noise ratio (CNR); moreover, elevation angle estimates from such processing will exhibit a bias which becomes increasingly severe as the CNR decreases. In this paper, we give expressions for the height sensitivity (i.e., the monopulse slope) of a SAL system employing this technique, and we verify our expressions with Monte Carlo simulations. We also examine effects related to nonzero target velocities and the spatial distribution of scattering in the target scene.

We furthermore derive a different elevation-angle estimator that involves processing the amplitudes of the detector-channel outputs jointly rather than in pairs, and we show that this estimator corresponds to the maximum likelihood (ML) estimator for the elevation angle. We demonstrate the performance of the ML estimator with simulations.

2. Overview and System Model

Monopulse processing has long been used for angle determination by microwave tracking radars and other similar systems [1]. Traditionally, these techniques have been broadly separated into phase-comparison and amplitude-comparison categories. While some variants combine aspects from both, we describe here the basic distinction between simple phase- and amplitude-comparison monopulse systems.

Monopulse Overview

In a microwave phase-comparison monopulse system, multiple antennas are used to simultaneously receive signals from a target of interest. The antennas are pointed in the same direction and are assumed to have nearly identical far-field beam patterns, but the antennas have different phase centers. The signals received by the two antennas from a given target therefore differ in phase, with the differential phase between the two dependent upon the off-boresight angle of the target.

In the amplitude-comparison case, the two antennas have the same phase center, but their far-field beam patterns are slightly offset from one another such that the signals received from the target differ in amplitude, with the amplitude difference dependent upon the off-boresight angle. The distinction between phase- and amplitude-comparison systems therefore lies in how the phase centers and beam patterns of the antennas are arranged.

By analogy to the microwave case, the design of a coherent lidar system intended for monopulse processing must include consideration of both the phase centers and the beam patterns of the system. Specifically, for phase-comparison processing, multiple phase centers must be employed. As this would necessitate modifications to the system's optics, including duplicate optics trains that would need to remain phase coherent, phase-comparison monopulse techniques might be very difficult to implement for a SAL system. By comparison, amplitude-comparison

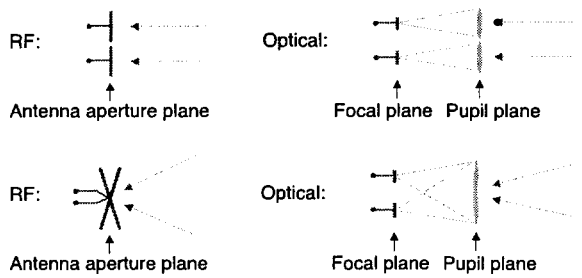

Figure 1. Schematic comparison of microwave and optical systems for phase-comparison (top) and amplitude-comparison (bottom) monopulse.

techniques might be easily implemented though the use of a focal-plane detector array and a single set of optics (see Fig. 1).

A source of confusion regarding monopulse techniques for optical systems is that in the microwave case, phase-comparison monopulse is often implemented by dividing a receive aperture into two halves which have different phase centers. The antenna aperture plane of a microwave system corresponds to the pupil plane, not the focal plane, of an optical system, however. A focal plane detector array more closely resembles a microwave reflector-antenna system with offset feeds. Either of these systems would have a single phase center and offset beam-patterns, making them suitable for amplitude- rather than phase-comparison monopulse.

It is clear that the monopulse-processing approach is determined by the design of the system hardware, as the processing technique must naturally be tailored to the manner in which the system collects data. Because of the relative difficulty of implementing phase-comparison monopulse in a SAL system, we devote the remainder of this paper to an examination of the amplitude-comparison case.

Simplified System Model

Consider the simple optical system model of Fig. 2. Let L be the length of the 1-D aperture and F be the focal length of the system. Assuming that the field in the focal plane is the Fourier transform of the field in the pupil plane, the field in the focal plane for a given point target at an angle θ_b off the optical axis is given by

$$S_f(x_f) = \int_{-L/2}^{L/2} A(x_p)\exp\left(-i\frac{2\pi}{\lambda}x_p\sin\theta_b\right)\exp\left(-i\frac{2\pi}{\lambda}x_p\sin\left(\tan^{-1}\left(\frac{x_f}{F}\right)\right)\right)\mathrm{d}x_p$$
$$\approx \int_{-L/2}^{L/2} A(x_p)\exp\left(-i\frac{2\pi}{\lambda}x_p\left(\theta_b - \frac{x_f}{F}\right)\right)\mathrm{d}x_p$$

Where x_p is the coordinate in the pupil plane, x_f is the coordinate in the focal plane, λ is the wavelength, and A is the aperture apodization function. In words, the field in the focal plane is the Fourier transform of the aperture weighting function, shifted by an amount proportional to the off-boresight angle of the target.

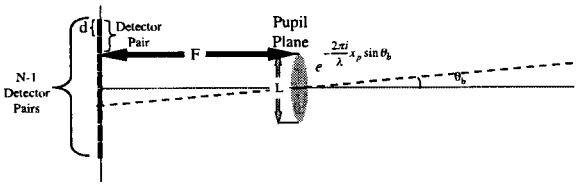

Figure 2. Diagram for 1-D system model of SAL system with focal plane-detector array.

The signal from each of the (assumed identical) detectors in the focal plane is an integral of the field S_f over the extent of the detector width d. Consider the standard amplitude-comparison monopulse processing scenario in which the sum and difference between adjacent detector channels are formed. The sum and difference signals are then given by

$$\Sigma_d(\theta_b) = \int_{-d}^{d} S_f(x_f)\mathrm{d}x_f = \underbrace{\int_{-d}^{0} S_f(x_f)\mathrm{d}x_f}_{\Sigma_B} + \underbrace{\int_{0}^{d} S_f(x_f)\mathrm{d}x_f}_{\Sigma_T}$$

$$\Delta_d(\theta_b) = \underbrace{\int_{0}^{d} S_f(x_f)\mathrm{d}x_f}_{\Sigma_T} - \underbrace{\int_{-d}^{0} S_f(x_f)\mathrm{d}x_f}_{\Sigma_B}$$

For a uniformly weighted aperture (A=1), these reduce to

$$\Sigma_d(\theta_b) = \int_{-d}^{d} \mathrm{sinc}\left(\frac{\pi L}{\lambda}\left(\theta_b - \frac{x_f}{F}\right)\right)\mathrm{d}x_f$$

$$\Delta_d(\theta_b) = \int_{0}^{d}\left[\mathrm{sinc}\left(\frac{\pi L}{\lambda}\left(\theta_b - \frac{x_f}{F}\right)\right) - \mathrm{sinc}\left(\frac{\pi L}{\lambda}\left(\theta_b + \frac{x_f}{F}\right)\right)\right]\mathrm{d}x_f$$

where the function sinc(x) is defined as sin(x)/x. Defining the monopulse discriminant μ by

$$\mu = \frac{\Delta_d}{\Sigma_d}$$

we have

$$\mu = \frac{-\frac{\lambda F}{\pi L}\left[\mathrm{Si}\left(\frac{\pi L}{\lambda}\left(\theta_b + \frac{d}{F}\right)\right) + \mathrm{Si}\left(\frac{\pi L}{\lambda}\left(\theta_b - \frac{d}{F}\right)\right) - 2\mathrm{Si}\left(\frac{\pi L \theta_b}{\lambda}\right)\right]}{\frac{\lambda F}{\pi L}\left[\mathrm{Si}\left(\frac{\pi L}{\lambda}\left(\theta_b - \frac{d}{F}\right)\right) - \mathrm{Si}\left(\frac{\pi L}{\lambda}\left(\theta_b + \frac{d}{F}\right)\right)\right]}$$

where

$$\mathrm{Si}(x) = \int_0^x \mathrm{sinc}(t)\mathrm{d}t$$

An approximate expression for the discriminant slope, k_{θ_b}, is obtained by differentiating with respect to θ_b and evaluating the resulting expression at 0:

$$k_{\theta_b} = \frac{\partial}{\partial\theta_b}\frac{\Delta_d}{\Sigma_d}\bigg|_{\theta_b=0} = \frac{\Sigma_d \frac{\partial\Delta_d}{\partial\theta_b} - \Delta_d \frac{\partial\Sigma_d}{\partial\theta_b}}{\Sigma_d^{\,2}}\bigg|_{\theta_b=0} = \frac{1}{\Sigma_d}\frac{\partial\Delta_d}{\partial\theta_b}\bigg|_{\theta_b=0}$$

The derivative of the difference channel is given by

$$\frac{\partial\Delta_d}{\partial\theta_b} = -F\left[\mathrm{sinc}\left[\frac{\pi L}{\lambda}\left(\theta_b + \frac{d}{F}\right)\right] + \mathrm{sinc}\left[\frac{\pi L}{\lambda}\left(\theta_b - \frac{d}{F}\right)\right] - 2\mathrm{sinc}\left[\frac{\pi L}{\lambda}\theta_b\right]\right]$$

Using the truncated Taylor expansions for Si(x) and sinc(x), we have

$$\mathrm{Si}(x) \approx x$$

and

$$\mathrm{sinc}(x) \approx 1 - \frac{x^2}{6}$$

so

$$\Sigma_d(0) = \frac{2\lambda F}{\pi L}\mathrm{Si}\left(\frac{\pi L d}{\lambda F}\right) \approx \frac{2\lambda F}{\pi L}\frac{\pi L d}{\lambda F} = 2d$$

and

$$\frac{\partial\Delta_d}{\partial\theta_b}\bigg|_{\theta_b=0} = -2F\mathrm{sinc}\left(\frac{\pi L d}{\lambda F}\right) + 2F \approx -2F\left(1 - \frac{1}{6}\left(\frac{\pi L d}{\lambda F}\right)^2\right) + 2F = 2F\frac{1}{6}\left(\frac{\pi L d}{\lambda F}\right)^2$$

The discriminant slope k_0 is therefore approximately given by

$$k_{\theta_b} = \frac{1}{\Sigma_d}\frac{\partial\Delta_d}{\partial\theta_b}\bigg|_{\theta_b=0} \approx \frac{1}{6}\frac{d}{F}\left(\frac{\pi L}{\lambda}\right)^2$$

The monopulse discriminant itself can thus be approximated by

$$\mu(\theta_b) \approx \frac{\pi^2}{6}\left(\frac{d/F}{\lambda/L}\right)\left(\frac{\theta_b}{\lambda/L}\right)$$

Note that if the focal-plane detectors are spaced by approximately half the diffraction-limited beamwidth λ/L, the monopulse discriminant is approximately equal to the off-boresight angle, normalized by the diffraction-limited beamwidth.

The above approximation for the monopulse discriminant is useful for providing intuition on the sensitivity of the system relative to the key system parameters. For precise estimates of the off-boresight elevation angle, and hence the target height, however, more accurate numerical analyses are necessary.

3. Sensitivity and Performance

The above approximation for the monopulse discriminant applies only to the assumed 1-D system model. Expressions for a more realistic system model are more difficult to derive analytically, however, so we now examine numerical analyses for the behavior of more complete system models.

As described above, the system response is given by the 2-D convolution of the optical system response

function (i.e., the Fourier transform of the pupil plane weighting function) with the detector response function (i.e, the size and shape of the detector). In general, this amounts to a 4-D integral since the two dimensions of the pupil plane and the two dimensions of the focal plane must be considered.

Dependence on Beam Pattern

We begin by examining the dependence of the monopulse discriminant μ on the shape of the beam pattern. To facilitate the numerical analysis and to isolate the effects of varying the beam pattern, we assume for now that the detectors are infinitely small so that they can be modeled by Dirac delta functions separated by $\lambda/2D$, where D denotes the equivalent aperture diameter. We consider beam shapes described by (1) a Gaussian function, (2) a sinc function defined as above, and (3) a jinc function defined by $\mathrm{jinc}(x)=J_1(\pi x)/2x$ where J_1 is a Bessel function of the first kind. The sum and difference channel amplitudes and the monopulse discriminant μ are shown in Fig. 3 as functions of the normalized elevation angle $\theta_b/(\lambda/D)$. The monopulse discriminant remains approximately linear over the diffraction-limited beamwidth, but the difference in slopes for the different beam shapes suggests that the beam pattern must be well characterized in order to accurately estimate the elevation angle θ_b. Moreover, departures of the monopulse discriminant from linearity could be significant depending on the application.

Dependence on Detector Size

We next relax the point-detector approximation and examine the effects of finite detector size. The sum and difference patterns and the monopulse discriminant are shown in Fig. 4 for the two cases of (1) point detectors separated by $0.6\lambda/D$ in elevation and (2) adjacent, finite-sized detectors $0.6\lambda/D$ wide in elevation and $1.2\lambda/D$ wide in azimuth. A jinc function was used for the lens response pattern in this analysis. The results indicate that the point-detector approximation is valid for a first-order analysis of the system sensitivity, but that a full analysis of the system would be required for precision applications.

Noise Performance

Expressions for the behavior of the monopulse discriminant in the presence of noise (e.g., shot noise) in the individual detector channels were derived by Kanter [2]. Let $\hat{\mu}$ be the estimate of the monopulse discriminant in the presence of noise, while μ_0 is the discriminant that would be observed in the absence of noise.
For a target of constant brightness (no target fluctuations), the estimate exhibits a bias

$$E[\hat{\mu}|\chi]=\left(1-e^{-\chi}\right)\mu_0$$

where χ is the signal-to-noise (SNR) ratio. If the target brightness fluctuates with an exponential distribution (i.e., the target has complex Gaussian statistics),

$$E[\hat{\mu}]=\int_{-\infty}^{\infty}E[\hat{\mu}|\chi]f_\chi(\chi)\mathrm{d}\chi$$

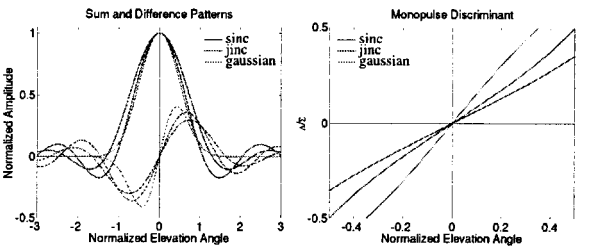

Figure 3. Sum and difference patterns (left) and monopulse discriminant (right) as functions of normalized elevation angle for various beam patterns.

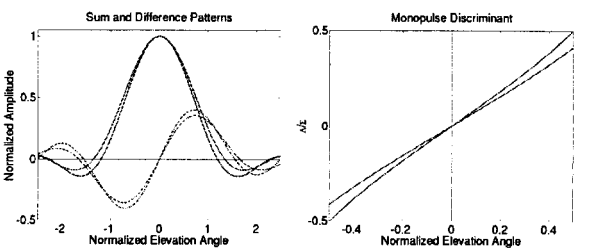

Figure 4. Sum and difference patterns (left) and monopulse discriminant (right) as functions of normalized elevation angle for the point detector approximation and for finite-sized detectors.

Let the carrier-to-noise ratio (CNR) be defined as the expected value of χ. Then,

$$E[\hat{\mu}]=\mu_0\int_0^{\infty}\left(1-e^{-\chi}\right)\left(\frac{e^{-\chi/\mathrm{CNR}}}{\mathrm{CNR}}\right)\mathrm{d}\chi$$

giving

$$E[\hat{\mu}]=\frac{\mathrm{CNR}}{\mathrm{CNR}+1}\mu_0$$

This is the well-known bias exhibited by monopulse systems. For high values of the CNR, the bias becomes negligible, but for low values of the CNR, the monopulse discriminant can be severely underestimated.

Kanter showed that the variance of the monopulse discriminant does not exist in the presence of noise because the integral defining the discriminant does not converge. An equivalent, approximate standard deviation, however, is given by

$$\sigma_{\hat{\mu}}\approx\sqrt{\frac{1-\mu_0^{\,2}}{(N_{\mathrm{looks}})(2\mathrm{CNR})}}\approx\frac{1}{\sqrt{(N_{\mathrm{looks}})(2\mathrm{CNR})}}$$

where N_{looks} is the number of independent samples ("looks") averaged together to form the estimate. Looks may be obtained by averaging adjacent image pixels together if the statistics of the scene are assumed to be locally stationary. Assume as above that the detector spacing is approximately

$0.5\lambda/D$ so that the monopulse slope is approximately unity. Then,

$$\sigma_{\theta_{el}} = \frac{\lambda}{D}\sigma_{\hat{\rho}}$$

The standard deviation of the height estimate of the target is then given by

$$\sigma_z = \rho \sin\theta_{\text{look}} \sigma_{\theta_{el}} = \rho \sin\theta_{\text{look}} \frac{\lambda}{D}\sigma_{\hat{\rho}}$$

or

$$\sigma_z = \frac{\lambda}{D} \frac{\rho \sin\theta_{\text{look}}}{\sqrt{(N_{\text{looks}})(2\text{CNR})}}$$

where ρ is the slant range to the target and θ_{look} is the angle of the look direction at the platform with respect to nadir, assuming a broadside acquisition geometry.

Note that in the above expressions, the CNR refers to the CNR of the sum channel. Other definitions are possible, however, and this definition will be revised for some of the analysis below.

Monte Carlo simulations verifying the noise performance of an amplitude-comparison monopulse system are given in Figs. 5 and 6. The computed standard deviation and mean of the monopulse-discriminant estimate are plotted as functions of the true monopulse discriminant for various CNRs. The noise is modeled by additive, white, complex Gaussian noise. The case of a non-fluctuating target is shown in Fig. 5. Variations of the estimated standard for low values of the CNR are due to the ill-behaved nature of the true standard deviation. The bias described above is evident in the less-than-unity slopes of the mean for low values of the CNR. Equivalent curves are shown in Fig. 6 for the case of a target with Gaussian statistics, assuming four looks weighted by power (the CNR refers to the single-look CNR). Because of the extra looks, the standard deviations are lower and more well behaved, yet the bias is still present. This point underscores the significance of the bias in the monopulse discriminant—averaging many measurements together can reduce the standard deviation of the estimate, but the estimate converges on the incorrect (biased) value of the discriminant. This implies that high CNRs (e.g., 10 dB or better) would be necessary for accurate height measurements.

Note that compensation for bias effects during processing would be difficult since doing so would require knowledge of the target reflectivity in order to estimate the CNR.

4. Other Effects

Distributed Target Effects

Several effects related to distributed rather than point targets are worth noting. Targets that are large compared to the wavelength will likely exhibit brightness fluctuations or speckle, so even with a high average CNR, some dark pixels may be characterized by large height

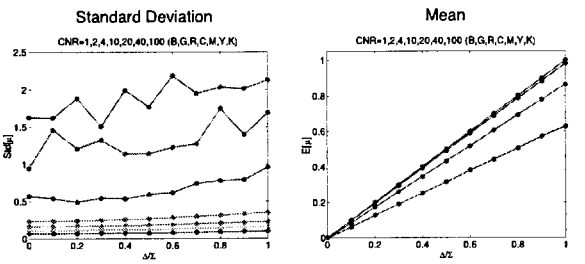

Figure 5. Results from a Monte Carlo simulation of monopulse performance for a nonfluctuating point target in the presence of noise.

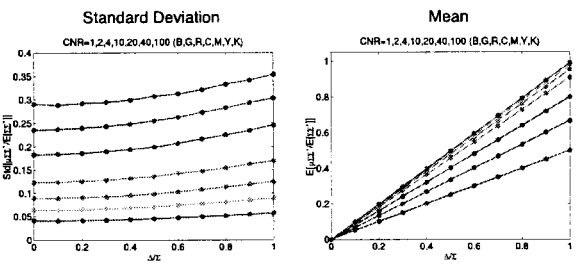

Figure 6. Results from a Monte Carlo simulation of monopulse performance for a target with Gaussian statistics in the presence of noise. Four looks are assumed.

errors. This effect can be accommodated at the expense of horizontal resolution by taking looks as described above, as long as the scene has locally stationary statistics.

While much of the above analysis assumed point targets, good height estimates of distributed scenes might still be obtained if the scatterer distribution within a resolution cell is small compared to the elevation beamwidth. That is, although a scene is comprised of scatterers at a multitude of elevation angles, only a small set of elevation angles will be contained in any given resolution cell, as illustrated in Fig. 7. However, if the vertical scatterer distribution within a cell is large compared to the elevation beamwidth, problems arise. In such cases, though, the underlying assumption that the height of the target scene can be characterized by a single value is violated and any height estimate would be somewhat suspect.

Synthetic Aperture Processing

Although most of the previous analyses have assumed a real-beam system, the same ideas are directly applicable to synthetic-aperture images. That is, synthetic aperture processing [3,4] is performed for each of the different detector channels. Monopulse processing is then performed on images from adjacent detector channels on a pixel-by-pixel basis to yield height estimates for each image pixel in the scene. Naturally, amplitude effects introduced during synthetic-aperture processing must be carefully examined for their impact on amplitude-comparison monopulse processing.

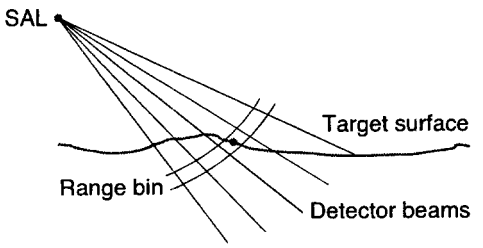

Figure 7. The angular extent in elevation of the part of the target surface contained in a single resolution cell is limited by the system resolution.

Target Velocity Effects

Nonzero target velocities will have a strong effect on the image quality because of the short wavelengths of SAL sensors relative to their microwave counterparts. In addition to their effects on the imagery however, velocity effects will also affect monopulse height estimates. The mechanisms by which these effects are manifested vary between spotlight and stripmap acquisitions, but both modes are affected.

Note that as the target velocity increases its apparent velocity will be wrapped because of the sampling (at the pulse repetition frequency or PRF) of the system. For typical SAL system parameters, the target velocity that wraps the Doppler spectrum is on the scale of millimeters per second.

In spotlight mode, in which the beam is steered to stay fixed with respect to the ground during the SAL integration time, the amplitude of a given image pixel depends on the position of the pixel in the beam. The expected amplitude ratio between pixels in adjacent detector beams thus depends on the along-track location of the pixel. That is, the monopulse discriminant slope depends on the along-track pixel position. If a target has a radial velocity component, however, it will appear in the azimuth-compressed image at a different along-track location than would be expected. Height estimation for the target will therefore be performed for the target with an incorrect inversion function that maps the observed monopulse discriminant into an erroneous elevation angle. These effects are even more pronounced if the axis of the focal plane detector array is not perfectly aligned with the elevation direction. The effects are quantified through simulation in Fig. 8. The plots may be interpreted by taking the angle-estimation error to be the vertical distance between curves of different colors and a given line style for a fixed horizontal-axis location.

In stripmap mode (see Fig. 9), all targets on the ground enter and leave the beam as the instrument passes by and therefore have the same amplitude and phase history. The amplitude of a target in the azimuth-compressed imagery therefore depends on the matched-filter parameters used during synthetic-aperture processing, most notably the Doppler centroid value. A target radial velocity component will shift the Doppler spectrum of the target, however, so the azimuth matched filter will no longer be matched for a moving target. The target's amplitude is moreover affected differentially over the detector array since the matched filter parameters for different beam positions and different detectors vary. Once again, rotations of the detector axis compound the difficulties in estimating the heights of moving targets.

5. Array Processing

Monopulse processing based on sum and difference channels and the monopulse discriminant μ implicitly assume that only two channels will be processed at a time. Therefore, for a focal-plane detector array with M channels (see Fig. 7), the processing scenario would imply that M-1 discriminant estimates μ_m would be formed from the possible pairings of adjacent detectors for each pixel. Not only would each

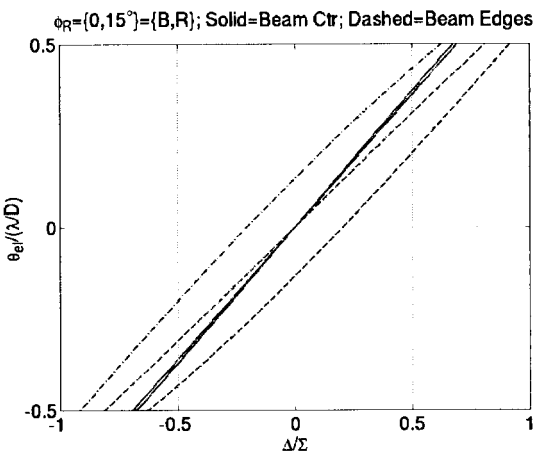

Figure 8. Elevation angle as a function of the monopulse ratio for targets at various along-track locations in a spotlight-mode SAL image for various detector-axis rotations.

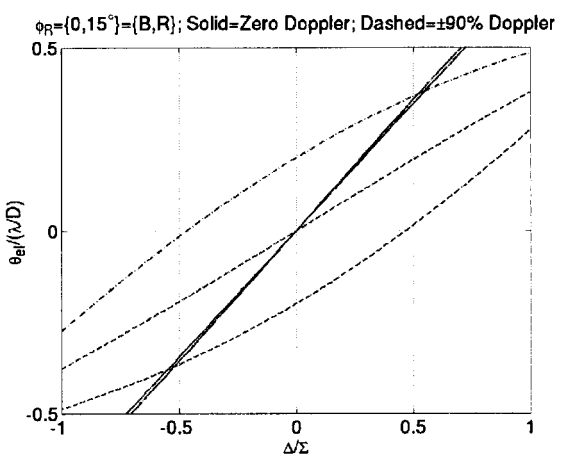

Figure 9. Elevation angle as a function of the monopulse ratio for targets with various Doppler centroid values in a stripmap-mode SAL image for various detector-axis rotations.

estimate exhibit the problems described in previous sections, but the correct estimate (i.e., the correct pair-wise output) must be chosen for the height estimate for each pixel. Although this problem is not addressed in detail here, it would presumably involve the selection of the detector pair for which the returned power is maximum. For low CNR values, however, this problem is not trivial. (Recall that each detector channel corresponds to an offset beam pattern of width λ/D; main-lobe energy returned from the target scene for a given resolution cell should therefore be isolated to a single detector pair.)

In this section we propose a different scheme for height reconstruction based on processing the amplitudes of all M detector channels jointly. Let the signal in channel m for pixel k be g_{km}. Its amplitude is given by

$$|g_{km}| = A_k h(\theta_m - \theta_{tgt}) \exp(j\phi_k) + n_{km}$$

where A_k and ϕ_k are the speckled brightness and phase of pixel k, $h(\theta)$ is the system response function (i.e., the convolution of the aperture response function and the focal-plane detector function), θ_m is the elevation boresight angle for the beam of detector channel m, θ_{tgt} is the elevation angle of the target, and n_{km} is a noise term, independent for each pixel and channel. Recall that all of the images from the different detector channels should exhibit highly correlated speckle patterns. We assume here that the speckle in each channel is identical.

The maximum likelihood (ML) estimate $\hat{\theta}_0$ for the target elevation angle is that which satisfies the following expression:

$$\hat{\theta}_0 = \arg\max\left\{ f\left(\vec{g} \middle| \theta_0, \vec{A}, \vec{\phi}\right) \right\} = \arg\max\left\{ \prod_k \prod_m f\left(g_{km} \middle| \theta_0, A_k, \phi_k\right) \right\}$$

where f denotes a joint conditional probability density function (pdf) and the vector notation is used to denote a set of variables subscripted by k or m. Assuming independent Gaussian statistics for the noise, the individual pdf for a given pixel and channel can be written as

$$f\left(g_m \middle| \theta_0, A_k, \phi_k\right) = \frac{1}{2\pi\sigma_n^2} \exp\left(\frac{1}{2\sigma_n^2} \left| g_{km} - A_k h(\theta_m - \theta_0) \exp(j\phi_k) \right|^2 \right)$$

where σ_n^2 is the noise variance. The logarithm of the joint pdf has the same maximum as the function itself, so the ML estimate can be rewritten as

$$\hat{\theta}_0 = \arg\max\left\{ \sum_k \sum_m \log\left(f\left(g_{km} \middle| \theta_0, A_k, \phi_k\right) \right) \right\}$$

where we have transformed the products into sums. Substituting the expression for the individual pdf into the double summation, and dropping scale factors which do not affect the maximization, we obtain

$$\hat{\theta}_0 = \arg\max\left\{ -\sum_k \sum_m \left| g_{km} \exp(-j\phi_k) - A_k h(\theta_m - \theta_0) \right|^2 \right\}$$

Next, we assume that $h(\theta)$ is a purely real function and expand the power term inside the double summation:

$$\hat{\theta}_0 = \arg\max\left\{ -\sum_k \sum_m \left[|g_{km}|^2 - 2\mathrm{Re}\left\{ g_{km} \exp(-j\phi_k) \right\} A_k h(\theta_m - \theta_0) + \left(A_k h(\theta_m - \theta_0) \right)^2 \right] \right\}$$

The summations can then be rearranged, and terms that do not depend upon A, θ_0, or ϕ can be dropped, giving

$$\hat{\theta}_0 = \arg\max\left\{ -\sum_k \left[-2 \sum_m \mathrm{Re}\left\{ g_{km} \exp(-j\phi_k) \right\} A_k h(\theta_m - \theta_0) + A_k^2 \sum_m \left(h(\theta_m - \theta_0) \right)^2 \right] \right\}$$

Now, note that $h(\theta)$ is band limited since the aperture function has finite support, so, assuming that $h(\theta)$ is appropriately scaled and that the detector spacing is adequately fine to Nyquist sample the function,

$$\sum_m \left(h(\theta_m - \theta_0) \right)^2 = 1$$

The ML estimator can then be written as

$$\hat{\theta}_0 = \arg\max\left\{ \sum_k \left[2A_k \mathrm{Re}\left\{ \exp(-j\phi_k) \sum_m g_{km} h(\theta_m - \theta_0) \right\} - A_k^2 \right] \right\}$$

giving

$$\hat{\theta}_0 = \arg\max\left\{ \sum_k \left[2A_k \left| \sum_m g_{km} h(\theta_m - \theta_0) \right| - A_k^2 \right] \right\}$$

We can then take the partial derivative of the term inside the outer summation and set the partial derivative to zero. This gives the ML estimate A_{kML} of A_k.

$$\frac{\partial}{\partial A_k} \left[2A_k \left| \sum_m g_{km} h(\theta_m - \theta_0) \right| - A_k^2 \right] = 2 \left| \sum_m g_{km} h(\theta_m - \theta_0) \right| - 2A_k = 0$$

$$A_{kML} = \left| \sum_m g_{km} h(\theta_m - \theta_0) \right|$$

Substituting A_{kML} into the ML estimate for θ_0 gives

$$\hat{\theta}_0 = \arg\max\left\{ \sum_k \left| \sum_m g_{km} h(\theta_m - \theta_0) \right|^2 \right\}$$

Thus, in a statistical sense, this is the best possible estimate of the elevation angle, subject, of course, to our assumptions governing the problem.

This expression for the ML estimate can be interpreted as follows. To process the data, we should convolve the set of M detector samples g for a given pixel with a sampled version of the lens response function $h(\theta)$. The location of the peak of the convolution is the ML estimate of the elevation angle. This is equivalent to matched-filtering the detector data for each pixel. For finer

precision, the data can be upsampled. If K looks are taken, the powers of the complex upsampled convolution results should be averaged across the K pixels before the maximum value over the index m is selected.

The ML estimator has some desirable properties when compared to schemes employing Δ/Σ monopulse processing. First, the ML estimator provides a way to rigorously select the detector pair to choose for the height estimate, whereas this is a separate step for Δ/Σ processing. The ML estimator also avoids most of the bias issues associated with Δ/Σ processing, though the ML estimator may still have a (less significant) bias arising from the finite extent of the detector array. The ML estimator may be somewhat more computationally intensive, however.

Results of a Monte Carlo simulation of the performance of the ML estimator are shown in Fig. 10. The plots depict the mean and standard deviation of the elevation angle estimate as a function of the true elevation angle for various numbers of looks and CNR values. Departing from the definition of CNR above, the CNR is defined here as the CNR that would be observed for a single hypothetical detector for which the target is exactly at boresight. The detector array was assumed to consist of 16 channels separated by $0.5\lambda/D$.

As expected, the standard deviation decreases as the number of looks increases and/or the CNR increases. Notably, however, the bias becomes less significant not only for an increasing CNR, but also for an increasing number of looks (contrast this with the Δ/Σ curves shown above). Moreover, the bias is a bias towards the center of the entire detector array, not the center of a particular detector pair, so errors associated with the bias in the ML scheme may be more systematic and more easily removed than those of a Δ/Σ scheme. Not unexpectedly, towards the end of the detector array (the far right of the left-hand plot), the accuracy is degraded as energy from the target falls past the end of the detector array.

An example Monte Carlo result for fixed parameter values is shown in Fig. 11. In this case, the CNR is 3 (5 dB), θ_b=1.78λ/D, and four looks are taken. Even for the relatively low CNR and number of looks, the estimate does not show a significant bias. The angular accuracy of the measurement is better than λ/D.

6. Conclusions

We have examined the performance of amplitude-comparison monopulse techniques for height estimation in the context of a SAL system. Our expressions for the height sensitivity of the system suggest that the technique is suitable for giving height estimates accurate to a fraction of the diffraction-limited beamwidth λ/D. The bias inherent in Δ/Σ processing implies that relatively high CNRs will be required for such estimates, however, even if multiple looks are taken. These bias issues can be avoided, for the most part, by employing the joint array-processing ML method derived here.

Effects such as vertical target distribution and target motions may limit the height-estimation performance of the system. In the case of target motions, however, the degradation of height-estimation accuracy may be of secondary importance since the degradation of the overall image quality may have a more dominant effect on scene interpretability.

We conclude that amplitude-comparison techniques, implemented with the use of a focal-plane detector array, are more viable for airborne SAL height estimation than phase-comparison techniques such as phase-comparison monopulse or interferometry.

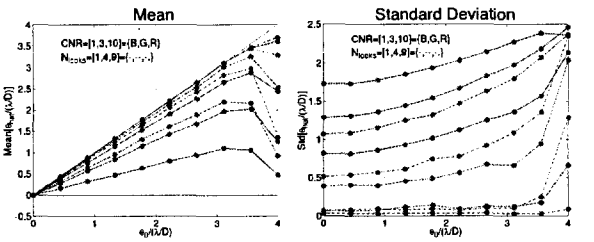

Figure 10. Results from a Monte Carlo simulation of the performance of the ML elevation-angle estimator for a target with Gaussian statistics.

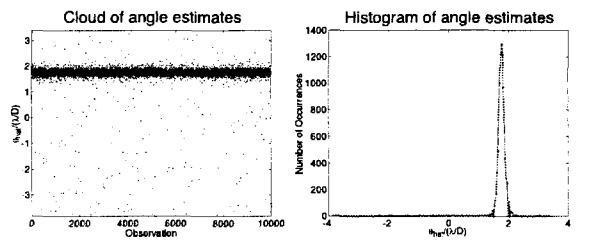

Figure 11. Example of ML elevation-angle estimator performance for a fixed target location, CNR, and number of looks.

Acknowledgement

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology and was sponsored by the Defense Advanced Research Projects Agency through an agreement with the National Aeronautics and Space Administration.

References

[1] Samuel M. Sherman, *Monopulse Principles and Techniques*, Artech House: Dedham, 1984.

[2] Irving Kanter, "The Probability Density Function of the Monopulse Ratio for N Looks at a Combination of Constant and Rayleigh Targets," *IEEE Transactions on Information Theory*, **IT23**, 643–648, September, 1977.

[3] John C. Curlander and Robert N. McDonough, *Synthetic Aperture Radar: Systems and Signal Processing*, Wiley: New York, 1991.

[4] Walter G. Carrara, Ron S. Goodman, Ronald M. Majewski, *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*, Artech House: Boston, 1995.

BIOGRAPHIES

Curtis W. Chen *is currently...*

Scott Hensley *is currently...*