SANDIA REPORT

Collecting and Processing Data for High Quality CCD Images

Armin W. Doerry

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Collecting and Processing Data for High Quality CCD Images

Armin W. Doerry
SAR Applications Department

Sandia National Laboratories
PO Box 5800
Albuquerque, NM 87185-1330

ABSTRACT

Coherent Change Detection (CCD) with Synthetic Aperture Radar (SAR) images is a technique whereby very subtle temporal changes can be discerned in a target scene. However, optimal performance requires carefully matching data collection geometries and adjusting the processing to compensate for imprecision in the collection geometries. Tolerances in the precision of the data collection are discussed, and anecdotal advice is presented for optimum CCD performance. Processing considerations are also discussed.

ACKNOWLEDGEMENTS

CONTENTS

FOREWORD

Coherent Change Detection has been demonstrated to be a very powerful tool for Intelligence, Surveillance, and Reconnaissance (ISR). Nevertheless, those of us that appreciate the power of CCD are perplexed at the seeming lack of equivalent appreciation in the larger ISR community. Investigations of this dilemma invariably yield a result that detractors 1) confuse CCD with optical change detection, and/or 2) just haven't seen good CCD products and hence question its value and/or reliability.

Further investigation often yields that those systems capable of providing good CCD results are often not operated in a manner to yield optimum CCD performance. A properly operated radar and properly processed data should routinely and reliably yield a high-quality CCD product with its attendant novel and unique ISR signatures.

How to optimally operate a radar and a strategy for best processing the resultant data into high-quality CCD products has not been significantly dealt with in the literature. This report intends to remedy this. Nevertheless, most of what is included herein is simply common sense put down on paper (or equivalent).

1 Introduction & Background

Coherent Change Detection (CCD) is a technique for observing very subtle changes between two Synthetic Aperture Radar (SAR) images. It is an Interferometric processing technique that measures the coherence between two images, and denotes 'change' where coherence is not observed, and 'no change' where coherence is observed. The rudiments of CCD processing can be found in several published sources.[1,2]

Since 'change' is denoted by lack of coherence between like regions in two images, the strategy must be to form both images with as much initial coherence as possible, and then see where in spite of our best efforts coherence cannot be achieved. Many things contribute to destroying coherence, but we want to eliminate all sources except for temporal change in the scenes being imaged themselves.

Consequently, the required strategy must have a number of variables well controlled, and processing must be adapted to mitigate the effects of residual imperfections. The entire process towards a high-quality CCD product can be divided into two principal processes.

1) Collecting the data for maximal coherence, and

2) Processing the data for maximal coherence.

It must be emphasized that the coherence calculation, that is, calculating the actual CCD product from two images is the easy part. The hard part is making sure that the two input images have the underlying characteristics to yield a quality result.

The purpose of this report is to discuss "What it takes to get good CCD results." Anyone familiar with CCD processing will see this as an exercise in stating the obvious.

2 Discussion

As previously stated, factors leading to good CCD results can be divided into two main areas, specifically 1) the collection of the data, and 2) the processing of the data. We discuss these in turn.

2.1 Collecting the Data – The Flight Path

For maximal coherence, the underlying image pair needs to be as similar as possible. Since for Sandia designed radars an image is processed from raw data that is in fact a limited region in the Fourier space of the underlying scene being imaged, it follows that the underlying raw data for the two images needs to be as similar as possible. Ideally, neglecting atmospheric variations, the respective synthetic apertures need to be collected from identical locations (start and stop positions) along an identical flight path.

What WILL NOT work are two images of the same scene collected from substantially different imaging geometries.

What WILL work are two images of the same scene collected from identical perspectives along identical flight paths.

Since achieving absolutely identical flight paths is problematic, a natural question is "How close do they need to be to be good enough?" Toward this end we examine the principal parameters describing the imaging geometry.

2.1.1 Grazing Angle

Grazing angle is the complement to the local incidence angle. For a flat (non-sloped) terrain, it is equal to the elevation angle to the radar from the target scene. For a flat earth, it is equal to the depression angle, but then the earth isn't really flat, so especially at longer ranges they will be significantly different.

The mathematical foundation for how similar the grazing angles need to be is the Van Cittert-Zernike theorem.[3] This is the same problem as "baseline decorrelation" in Interferometric SAR (IFSAR or InSAR).[4] That is, as the grazing angles differ for the same scene, a baseline is formed between the flight paths. As the baseline increases, i.e. the grazing angle difference increases, correlation between the images decreases. Correlation disappears entirely at some baseline distance known as the "critical baseline."

Consequently, from these arguments we can calculate the critical grazing angle difference that corresponds to total loss of coherence as

$$\Delta \psi_{critical} = \frac{\lambda}{2\rho_r \tan \psi}$$

where

λ = the nominal wavelength of the radar signal
ρ_r = the nominal slant-range resolution, and
ψ = the nominal grazing angle.

It seems reasonable to specify an allowable grazing angle difference to be some small fraction of this, say perhaps 10% of the critical grazing angle difference. Consequently, a limit becomes

$$\Delta\psi \leq \frac{\lambda}{20\rho_r \tan\psi}.$$

Clearly, there is greater grazing angle tolerance for finer range resolutions and for shallower grazing angles.

As an example, a Ku-band SAR operating with 0.1 meter range resolution at a 30 degree grazing angle would indicate a grazing angle difference tolerance of 0.9 degrees. This corresponds to 78 meters of out-of-plane motion at 5 km range.

2.1.2 Bearing to Target

As with grazing angle, the mathematical foundation for how similar the bearing angles need to be is the Van Cittert-Zernike theorem. This leads to the requirement that the bearing angles from the target reference location to the synthetic aperture centers needs to be less than the angle subtended by the synthetic apertures themselves (which are presumed identical for both images). That is, the synthetic apertures for the two respective images need to overlap in their aperture angle spans. Perfect overlap is required for maximum coherence. The critical bearing difference is equal to the synthetic aperture angle which is dependent on azimuth resolution., that is,

$$\Delta\theta_{critical} = \frac{\lambda}{2\rho_a \cos\psi}$$

where

ρ_a = the nominal azimuth resolution.

We are neglecting effects of, and consideration for window functions for sidelobe control. This is justified because where the synthetic aperture is extended to compensate Impulse Response (IPR) broadening, it will be substantially attenuated by the tails of the window function anyway, thereby contributing comparatively little to the coherence measurement.

As with grazing angle, it seems reasonable to specify an allowable bearing angle difference to be some small fraction of this, say perhaps 10% of the critical bearing angle difference. Consequently, a limit becomes

$$\Delta\theta \le \frac{\lambda}{20\rho_a \cos\psi}.$$

Clearly there is greater bearing angle tolerance at finer azimuth resolutions, and steeper grazing angles (contrary to tolerance for grazing angle differences).

As an example, a Ku-band SAR operating with 0.1 meter azimuth resolution at a 30 degree grazing angle would indicate a bearing angle difference tolerance of 0.6 degrees. This corresponds to 52 meters of synthetic aperture offset at 5 km range.

In cases where synthetic apertures overlap, but inadequately so, image resolution can sometimes be traded for increased coherence by a technique known as "aperture trimming" which will be discussed later.

2.1.3 Range

Here we speak of the line-of-sight or slant range between target point and radar.

To first order, matching grazing angles and bearings to target are far more important than matching ranges to target. However this is not to say that range is unimportant to forming high quality CCD products. Indeed, differences in ranges are often accompanied with differences in grazing angles.

Longer ranges will generally yield noisier images, thereby diminishing coherence via the noise. Longer ranges at shallow grazing angles are subject to greater deleterious effects due to atmospheric propagation, causing differences in spatially variant phase errors and regional illumination variations.[5]

In stripmaps, individual image patches have widths often proportional to range. Consequently different ranges will cause image patches to not properly align. This is illustrated in Figure 1. Matching ranges will assist coherence, especially in stripmaps, and especially at shorter ranges.

Mosaicked image patches

Pass 1 synthetic apertures

Pass 2 synthetic apertures

Figure 1. Stripmaps created from mosaicked spots or patches will generally have patch geometries that depend on range.

2.1.4 Squint Angle

Squint angle is the difference between the ground track of the aircraft and the bearing from the aircraft to the target reference point, as measured at the center of the synthetic aperture.

Image layover characteristics depend on the squint angle of the synthetic aperture. This is a result of a 3-D scene being projected into a 2-D image. The projection depends on the squint angle, among other things. Given identical grazing angles and bearings to target, different squint angles will nevertheless yield different layover characteristics in the images. This is most noticeable in scenes that exhibit significant (non-flat) topography, whether natural or artificial (e.g. hill sides, buildings, etc.). To first order, images of scenes with substantial topography will appear 'warped' differently.

Differences in warping can often be addressed with more sophisticated registration algorithms. If left unmitigated, the CCD rendering will typically exhibit good coherence only in some regions of the image, sometimes in bands.

As a side note, knowing the relative imaging geometries, topography can be inferred from the differences in scene warping between two images. This is the principle of Stereo SAR for topographic measurements.[6,7] In short, Stereo SAR is very related to CCD.

Another problematic manifestation is due to the fact that a pixel in a layover region will exhibit a response that is a vector sum of all points along the same arc of projection.

Should multiple points on a surface lie along the same arc, then a pixel will exhibit a response that is the superposition of all those points. Different squint angles will cause different superposition results. A pixel where a bright high object lays over in two separate images onto two different independent bright low objects cannot be guaranteed to be coherent. This is most problematic in regions with sharp topographic gradients (e.g. cliffs, building edges, etc.). This is related to the "front-porch" problem in IFSAR height measurements.[8]

Ultimately, for good CCD performance, the degree to which squint angles need to match depends on the topography of the scene being imaged, both the elevation differences within the scene and the elevation gradients within the scene. Anecdotal evidence suggests that for scenes with significant topography, even 2 degrees of squint angle difference will cause noticeable degradation in the CCD product, unless suitably compensated.

Other anecdotal evidence suggests, however, that good pilots with pilot guidance can routinely match ground tracks to within 2 degrees for straight-line flight paths. Of course, autopilot control of an aircraft can substantially improve even this.

As a final note, although we recognize that the important factor is the layover characteristics in the underlying SAR images, we also understand that identical layover characteristics from straight and level flight will be generated with either of two aircraft ground tracks 180 degrees apart. That is, in Figure 2, images will be identical whether the aircraft flies from right to left, or left to right. Consequently, for CCD purposes a ground track of X degrees with a squint angle of Y degrees is equivalent to a ground track of X+180 degrees with a squint angle of Y+180 degrees.
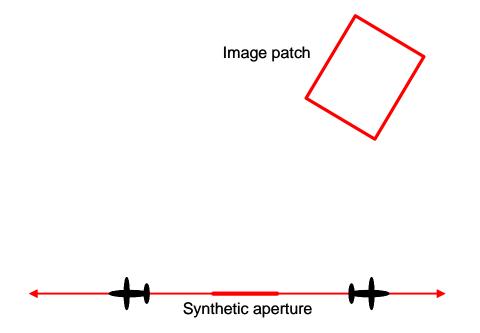

Figure 2. Equivalent layover characteristics will be generated by flight paths with 180 degree differences in aircraft ground tracks.

2.1.5 Pitch Angle

Here we speak of the pitch angle or tilt of the synthetic aperture which is distinct from that of the aircraft. It is generated from a climbing or descending flight path.

The pitch angle of the synthetic aperture affects layover in the same manner as squint angle. Consequently pitch angle differences cause the same degradations as squint angle differences.

In practice, however, this seems to be somewhat less of a problem than squint angle variations. It does seem to be somewhat more of a problem when flying circles in manned aircraft.

2.2 Processing the Data

If the data for the two SAR images to be compared were collected from identical synthetic apertures, i.e. exactly the same geometry with all the same errors and anomalies, then the coherence calculations would be trivial. In this situation image focusing would be identical and the images would be perfectly registered in all respects. For a static scene (no changes) the two images would be identical except for a floor of independent noise (observable in regions of low or no return, e.g. shadows).

In the typical event that the collection geometries and errors are not identical between the two synthetic apertures, the effects of their differences need to be mitigated to the extent possible in the processing. That is, this needs to be corrected prior to the actual coherence calculations.

The coherence calculations themselves are straight-forward and well documented in the literature.[1] We will not elaborate on the coherence calculations in this report, but rather address qualitatively what it takes for this calculation to give good results. We will remind the reader that the coherence measure is in fact the sample complex correlation coefficient between two images.

2.2.1 Focusing

Image focusing is a measure of how precisely a targets echo energy reflects its actual location in the scene. Well focused images have a target's energy well localized. Poorly focused images have a target's energy less localized, that is smeared or blurred across a larger spatial extent.

If the targets are blurred differently, then corresponding pixels in the two images contain different amounts of independent targets' signatures. Consequently, coherence is diminished. Ideally, the two images exhibit identical IPRs. Better yet, the two images should exhibit identical perfect IPRs. That is, the two images should both be well focused. In fine resolution systems, this usually requires a high-performance autofocus operation.

Anecdotal evidence suggests that many CCD results can be improved significantly by reiterating autofocus operations on the two input images as a very first step.

It should be kept in mind that focusing is not uniform in most SAR images, typically exhibiting poorer focus towards the image edges, and is especially consequential for large images at fine resolutions. As long as the misfocus remains similar, good CCD performance should be expected. However, the spatially variant misfocus is flight-path dependent, once again leading to the desire for identical imaging geometries. For large scenes, more exotic imaging algorithms might need to be employed.[9]

2.2.2 Registration

For good CCD performance, identical pixel locations need to correspond to identical scene features and spatial locations. This is the problem of image registration.

Typical image registration consists of the following steps (or their equivalent).

1) Measure displacement of corresponding regions or tie points in the two images,

2) Warp one image to match the other, i.e. minimizing the displacements in 1,

3) Repeat 1 and 2 until "good enough".

Measuring the displacement in corresponding regions is generally achieved by cross-correlating image segments between the images. Warping is done with classical interpolation or resampling techniques.

Uniform displacements between the synthetic apertures will result in grazing angle errors, bearing errors, and/or range errors. Correcting for these requires relatively simple warping functions.

Squint angle and pitch angle differences between the synthetic apertures, especially for scenes with significant topography, will require more elaborate spatially variant warping functions. Anecdotal evidence suggests this to be the case for angular differences as small as 2 degrees, although with more elaborate spatially variant warping functions good results are possible with substantially larger angular differences. Ground track differences as great as 90 degrees, albeit at 30 degrees grazing angle, have been demonstrated to achieve good coherence for 'bumpy' topographies.[7] Anecdotal evidence also suggests that matching squint angles is a greater problem for flight paths that are circles, especially under manual control of a pilot.

Even SAR data collected with high-quality motion measurement information is in practice virtually never good enough to avoid the need for registration corrections. Achieving good registration is the most difficult aspect of CCD processing.

2.2.3 Aperture Trimming

Sandia designed SAR systems typically use stretch processing of Linear FM waveforms, where the echo signals are de-ramped prior to digital sampling. Consequently the raw phase history data represent samples in the Fourier space of the scene being imaged. Other waveforms, however, can be relatively easily processed to samples in the scene's Fourier space, too.[10] A complete data set, nevertheless, represents only a small surface region in the total Fourier space of the scene. If the collection geometry for two synthetic apertures is identical, and the waveforms used are identical, then the Fourier-space surface regions will be identical as well.

SAR processing customarily projects the Fourier space data onto a 2-D plane, to allow more efficient 2-D processing instead of 3-D processing. This is justified by the normal presumption of relative flatness of the target scene surface.

Bearing angle differences, grazing angle differences, and/or waveform frequency content differences will generate respective sample surface regions in the scene's Fourier space that differ, too. Furthermore, their projections are not likely to overlap either. Furthermore yet, registering images will not necessarily register their Fourier space data regions. Any non-overlapping projection in Fourier space identifies energy that can not contribute to correlating the images. In fact, since this frequency space is inherently different between the two images, it contributes to reducing the coherence between the two images.

If the Fourier space regions for each image were trimmed to just the intersection of the two regions' projections, then coherence could be improved, albeit at the cost of a coarsening of the resolutions of the images themselves. This is usually a good trade.

Consequently, in addition to registering the images themselves, anecdotal evidence suggests that it is desirable to register their Fourier transforms, too, and trim their transforms to the intersecting subregions. This is illustrated in Figure 3.

As a final note, since window functions for sidelobe control are applied in the Fourier domain, trimming the aperture will modify the effective window function shape, thereby not only coarsening resolution, but also changing the shape of the IPR. Substantial aperture trimming may necessitate first removing any data tapering prior to trimming, and reapplying the tapering after trimming.
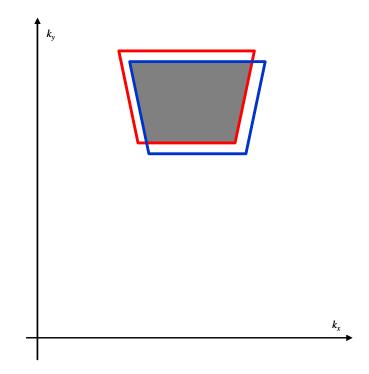
Figure 3. Fourier-space projections of the data. Data is limited to apertures. Only the overlapping portions of the apertures are useful to CCD processing.

2.2.4 Noise

CCD measures the temporal decorrelation between two images. In the absence of strong signals, uncorrelated additive noise between the two images will register as 'change'. Consequently good Signal-to-Noise Ratio (SNR) is required for good coherence measurements. For example, shadow regions, due to lack of any echo return will show poor coherence.

It is well-known that coherence is related to Signal-to-Noise Ratio (SNR) as[4]

$$\gamma = \frac{1}{1 + \dfrac{1}{SNR}} \, .$$

Consequently, solving for SNR yields the relationship

$$SNR = \frac{1}{(1/\gamma) - 1} \, .$$

To achieve $\gamma \geq 0.9$ requires $SNR \geq 9$ which we will approximate to a minimal 10 dB. That is, we need a noise floor in a SAR image at least 10 dB below the clutter field on

which we want to perform CCD for the undisturbed clutter to exhibit a coherence value of 0.9 or better.

2.2.4.1 _Additive Thermal Noise_

Because CCD's utility is principally in measuring disturbance to clutter, it is generally applied to rural (non-urban) scenes that contain corresponding terrain features, e.g. vegetation fields, dirt roads, etc. The noise in SAR images of these scenes tends to be dominated by additive thermal noise. It is well known how to predict SNR for this kind of noise based on radar parameters.[11]

Noise in a SAR image is generally referenced against an equivalent brightness in clutter reflectivity, measured as equivalent radar cross section per unit area, often standardized in units of $dBsm/m^2$ or simply dB. Values do depend on frequency, polarization, and grazing angle among other things. Tables of representative values exist in the literature.[12,13] Table 1 illustrates representative values for clutter at Ku-band.

Reliable high quality CCD results in the various clutter conditions cited would require a noise equivalent reflectivity at least 10 dB lower than the minimum reflectivity values cited. This suggests that overall, for Ku-band, a noise equivalent reflectivity should probably be specified at not much higher than −35 dB, and begin to become marginal in some cases at values above −30 dB.

Since a Ku-band SAR's range limit is often specified at allowing a noise floor to increase to −25 dB equivalent reflectivity, we recognize that CCD performance is likely to often be unsatisfactory near the published range limits of SAR.

Table 1. Ku-band clutter reflectivity values at 30 degree grazing angle, VV polarization.[12]

Clutter Type	Mean Reflectivity (dB)	Minimum Reflectivity (dB)
Soil and Rock Surfaces	−12.7	−21.2
Grasses	−13.9	−22.4
Roads	−8.5	−16.8
Dry Snow	−13.1	−19.7
Wet Snow	−16.8	−25.1

2.2.4.2 *Multiplicative Noise*

Multiplicative noise, typically manifested in the form of elevated sidelobes in the IPR, also spreads energy to unrelated pixels, similar to inadequate focusing. Multiplicative noise is specified relative to a desired IPR mainlobe response, that is, via a Multiplicative Noise Ratio (MNR). Consequently MNR is a Noise-to-Signal Ratio. Bright sidelobes generated by a target at one pixel location can mask or interfere with the true coherence measurement at another pixel location. Consequently, good CCD performance is facilitated by good MNR performance.

In a well focused image, MNR cannot generally be improved by transmitting more power or flying at nearer ranges in the manner that SNR due to additive thermal noise can be improved. MNR is highly dependent on the quality of the radar design and implementation. Consequently, a well designed radar intended for CCD applications should strive to make MNR inconsequential to CCD performance. This can be done by designing to achieve $\gamma \geq 0.99$, which requires $MNR \leq -20$ dB. Contrast this with a typical radar specification for $MNR \leq -13$ dB, which allows a coherence limit to no better than 0.95. While this doesn't sound too bad, it must be remembered that the actual multiplicative noise level depends on target scene content away from the region of interest.

For example, consider a scene with average reflectivity of −13 dB, with a small subregion at −22 dB reflectivity. We are interested in performing CCD on the small subregion. We will also ignore additive thermal noise for this example. The SAR system has a MNR of −13 dB. This establishes a multiplicative noise level of −26 dB equivalent reflectivity over the whole scene, including the subregion. In the small subregion the actual SNR becomes a very meager 4 dB, limiting coherence to no greater than 0.72. Of course this is a maximum, and the subregion average would likely be significantly less due to other coherence factors.

If, however, in our example the radar exhibited a MNR of −20 dB, then the multiplicative noise level would reduce to −33 dB equivalent reflectivity over the whole scene, including the subregion. In the small subregion the actual SNR becomes 11 dB, pushing the coherence limit to a much nicer 0.93.

2.2.4.3 *Quantization Noise*

The effects of image pixel quantization, both magnitude and phase, has been studied and reported by Thompson.[14] He found no CCD product degradation with as few as 8 bits of magnitude and 6 bits of phase in each image. As few as 3 bits of magnitude and 3 bits of phase per image were still useful for CCD given the right magnitude quantization scheme, albeit with degraded performance.

2.2.5 Image Frequency Content

Image warping for CCD involves interpolation and resampling of complex data. It is generally true that interpolation accuracy is easier for low frequency content in the data being interpolated. Quite simply, the more similar the data is for neighboring pixels, the easier it is to calculate interpolated values. Image warping is thereby facilitated by ensuring that the image spectrum is centered at 0 Hz, or DC, in both dimensions.

Images with spectrums centered at half the sampling frequency will warp with resulting phase variations and errors that manifest as fringe-like patterns in the CCD product. CCD results from images with spectrums centered at half the sampling frequency and again with spectrums moved to DC are shown in Figure 4.

As a final note, we observe that after centering the image spectrum at DC, image warping is generally easier with larger ratios of image resolution to pixel spacing, or over-sampling factors.
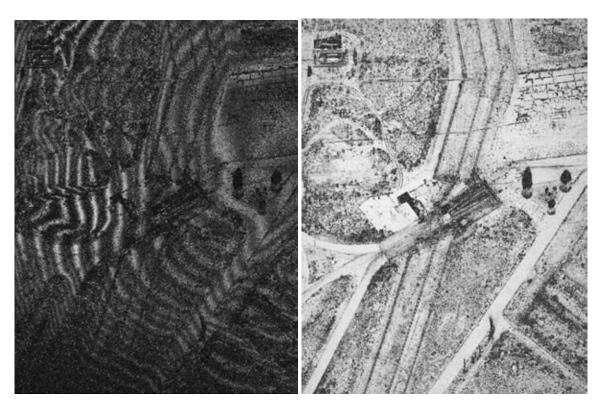

Figure 4. The left CCD product was generated from images that were warped with spectrums centered at half the sampling frequency. The right CCD product was generated from images that were warped with spectrums centered at DC.

2.3 Quality Metrics

Assuming that imaging geometries are chosen to allow good expected SNR, and the radar is designed to minimize other noise sources, the image data suitability for good CCD performance is best measured by how well image geometry parameters match each other. Anecdotal evidence suggests that the geometric parameters that warrant the most scrutiny are (in order of decreasing typical concern) bearing angle, grazing angle, and squint angle. Bearing angle and grazing angle differences should be less than 10% of critical values. Squint angle differences should be less than 2 degrees in straight and level flight with pilot guidance. Parameters outside these limits may require more sophisticated processing and perhaps result in degraded CCD performance.

CCD processing effectiveness can be inferred from a measure of the average coherence of the CCD product. If we define the interval [0,1] as the range of correlation coefficients for a pixel, with no correlation assigned to zero, and complete correlation assigned to one, then anecdotal evidence suggests that an average value over all pixels of 0.7 or greater will typically yield a high-quality and useful CCD product. This will vary somewhat due to image content, and be lower for images with large amounts of foliage and/or shadow regions. Figure 5 illustrates CCD products with average coherence values of 0.74 and 0.61 respectively.
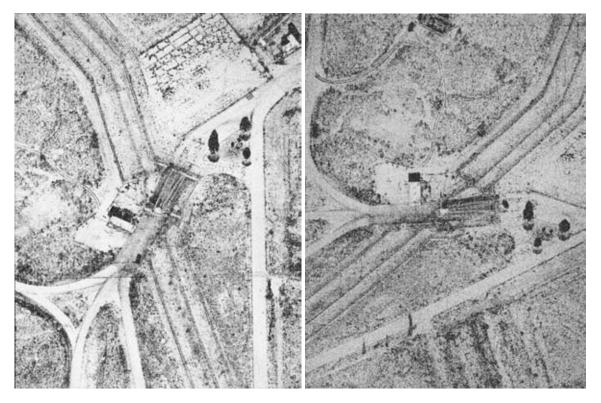

Figure 5. The CCD product on the left exhibits an average coherence of 0.74. The CCD product on the right exhibits an average coherence of 0.61. It is clearly of lesser quality although still useful.

2.4 SAR Operating Modes for CCD

The first step in generating good CCD candidate images is to collect data with as nearly identical image geometries as possible. Consequently a radar should base its collection geometry for its second pass on the geometry achieved for its first pass.

For spotlight SAR images, vehicle guidance should be based on achieving equal grazing angles and squint angles. Aperture beginning and end marks should be based on achieving equal bearings to the target scene. Small range differences are less of an issue for spotlight images.

For stripmap images, vehicle guidance should be based on achieving equal grazing angles and ranges. Aperture beginning and end marks should still be based on achieving equal bearings to the target scene. Squint angle is more easily controlled for the typical straight and level flight paths of stripmap collection geometries.

In all cases, processing the underlying SAR images should be as similar as possible, that is, with identical window functions, etc.

Of course, proper evaluation and utilization of stored SAR images requires that they be adequately annotated with complete image collection geometry and other relevant information.

3 Summary and Conclusions

The following principal conclusions should be drawn from this report.

- Matching data collection geometries is critical to forming high quality CCD products.

- Tolerances for data collection geometries is readily calculated. Formulas are given, and anecdotal observations are cited.

- Processing can to some degree compensate for mismatches in image collection geometries.

- Quality metrics for the images are described, as is a quality metric for the CCD product.

- High quality CCD products are best served by SAR operating modes that address collecting data with optimal geometries.

REFERENCES

[1] P. H. Eichel, D. C. Ghiglia, C. V. Jakowatz Jr., P. A. Thompson, D. E. Wahl, "Spotlight SAR interferometry for terrain elevation mapping and interferometric change detection", Sandia Report SAND93-2072, February 1996.

[2] C. V. Jakowatz Jr., D. E. Wahl, P. H. Eichel, D. C. Ghiglia, P. A. Thompson, *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*, ISBN 0-7923-9677-4, Kluwer Academic Publishers, 1996.

[3] J. W. Goodman, *Introduction to Fourier Optics*, McGraw-Hill, New York, 1968.

[4] H.A. Zebker, J. Villasenor, "Decorrelation in interferometric radar echoes", IEEE Transactions on Geoscience and Remote Sensing, Vol. 30, No.5, pp. 950-959, Sep 1992.

[5] Fred M. Dickey, Armin W. Doerry, Louis A. Romero, "Degrading effects of the lower atmosphere on long range airborne SAR imaging", accepted to IEE Proceedings on Radar, Sonar & Navigation.

[6] A. Doerry, "High-precision stereoscopic 3-D mapping accuracy", SPIE 2001 International Symposium on Aerospace/Defense Sensing, Simulation, and Controls, Algorithms for Synthetic Aperture Radar Imagery VIII, Vol. 4382, Orlando FL, 16 April 2001.

[7] D.A. Yocky, D.E. Wahl, C.V. Jakowatz, Jr., "Terrain elevation mapping results from airborne spotlight-mode coherent cross-track SAR stereo", IEEE Transactions on Geoscience and Remote Sensing, Vol. 42, No. 2, pp. 301- 308, Feb. 2004.

[8] D.L. Bickel, W.H. Hensley, D.A. Yocky, "The effect of scattering from buildings on interferometric SAR measurements", Proceedings of the IEEE 1997 Geoscience and Remote Sensing Symposium, IGARSS '97. Vol. 4, pp. 1545-1547, 3-8 Aug 1997.

[9] Armin W. Doerry, "Wavefront Curvature Limitations and Compensation to Polar Format Processing for Synthetic Aperture Radar Images", Sandia Report SAND2007-0046, Unlimited Release, January 2007.

[10] Armin W. Doerry, "SAR Processing with Non-Linear FM Chirp Waveforms", Sandia Report SAND2006-7729, Unlimited Release, December 2006.

[11] Armin W. Doerry, "Performance Limits for Synthetic Aperture Radar – second edition", Sandia Report SAND2006-0821, Unlimited Release, February 2006.

[12] Fawwaz T. Ulaby, M. Craig Dobson, *Handbook of Radar Scattering Statistics for Terrain*, ISBN 0-89006-336-2, Artech House, Inc., 1989.

[13] Maurice W. Long, *Radar Reflectivity of Land and Sea*, Second Edition, ISBN 0-89006-130-0, Artech House, Inc., 1983.

[14] D.G. Thompson, "Correlation and Image Compression for Limited-Bandwidth CCD", SAND 2005-3402, printed June 2005.

"If a man does his best, what else is there?"

George S. Patton

DISTRIBUTION

Unlimited Release

1	MS 1330	S. C. Holswade	5340
1	MS 1330	B. L. Burns	5340
1	MS 0519	T. J. Mirabal	5341
1	MS 1330	W. H. Hensley	5342
1	MS 1330	T. P. Bielek	5342
1	MS 1330	A. W. Doerry	5342
1	MS 1330	D. Harmony	5342
1	MS 1330	M. S. Murray	5342
1	MS 1330	S. Nance	5342
1	MS 1330	D. G. Thompson	5342
1	MS 1330	K. W. Sorensen	5345
1	MS 1330	D. F. Dubbert	5345
1	MS 1330	F. E. Heard	5345
1	MS 1330	G. R. Sloan	5345
1	MS 1330	S. M. Becker	5348
1	MS 1330	M. W. Holzrichter	5348
1	MS 1330	O. M. Jeromin	5348
1	MS 1330	D. M. Small	5348
1	MS 1330	D. C. Sprauer	5348
1	MS 1330	A. D. Sweet	5348
1	MS 0519	D. L. Bickel	5354
1	MS 0519	J. T. Cordaro	5354
1	MS 0519	J. DeLaurentis	5354
2	MS 9018	Central Technical Files	8944
2	MS 0899	Technical Library	4536