Improving Segmented Interferometric Synthetic Aperture Radar Processing Using Presumming

by: K. Clint Slatton

Final Report

Submitted to Professor Brian Evans

EE381K Multidimensional Digital Signal Processing

December 4, 1998

Abstract

Within the last decade, Interferometric Synthetic Aperture Radar (INSAR) processing has emerged as an important extension of conventional SAR because it can be used to infer topographic heights of the imaged terrain. The demand for high-resolution, large coverage-area topographic data is driving new INSAR systems to acquire data in long strips. When processing long strips of INSAR data, the motion of the sensor requires processing the data in patches, each with its own set of motion parameters. Conventional techniques for aligning these patches are not adequate for INSAR applications, where systematic height errors of roughly 2 m are visible in the topographic images. A data-smoothing procedure known as presumming can reduce these errors by rejecting aliased frequencies in the azimuth spectrum. Presumming operations were implemented in an existing INSAR processor to reduce the patch boundaries in the topographic images. It was found that the location of the presumming operations is critical to reducing patch boundaries effectively. Presumming after motion compensation smoothes features in the azimuth direction, thus making patch boundaries less distinct, but it does not specifically reduce the patch boundary errors. Implementing presumming prior to motion compensation operations may improve results.

1. Introduction

Imaging radars comprise an important class of sensors for remote sensing and reconnaissance applications because they access the microwave portion of the spectrum and are capable of day, night, and nearly all-weather operation. An imaging radar transmits pulses of frequency-modulated electromagnetic waves as it travels along a trajectory. Images of the terrain are formed after the reflected pulses are combined in a series of operations collectively referred to as a radar processor. Figure 1 depicts a radar imaging a swath on the terrain surface. The pulses are transmitted at a specified pulse repetition frequency (PRF). When the reflected pulses are summed coherently, (i.e. accounting for the varying range to a point on the ground), a long aperture antenna is synthetically created, giving rise to the term Synthetic Aperture Radar (SAR) [1].

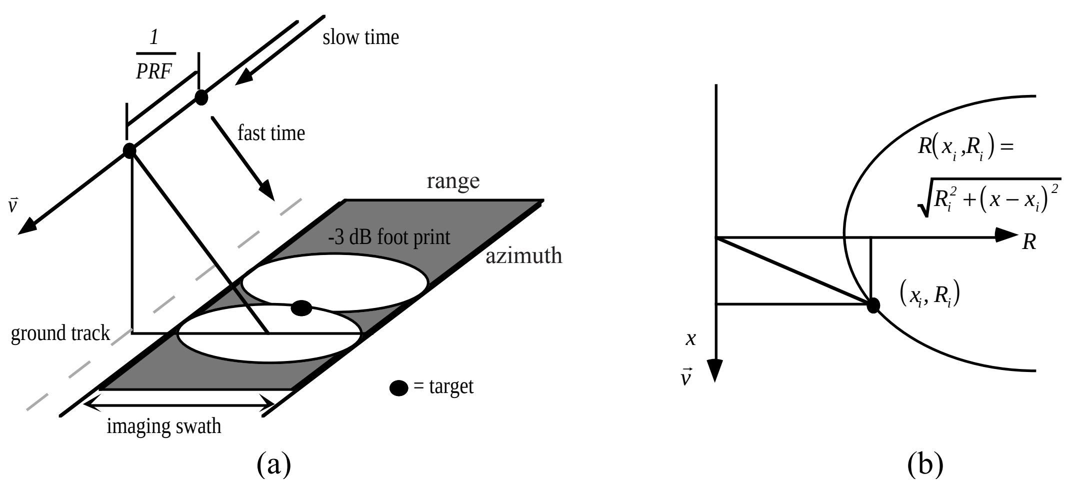

Fig. 1. (a) SAR imaging geometry and (b) a plot describing how the range to a target varies as the SAR moves along its trajectory

In the general case, SAR data form a non-separable, shift-varying 2-D process, but the time scales for the two coordinate bases of range and azimuth are sufficiently different that the signal can be approximated as two quasi-independent 1-D processes [2]. The most common method of processing SAR data is the Range-Doppler (RD) method. This method correlates the narrow-band, frequency-modulated pulse with a reference function in the range dimension. Then it correlates the Doppler-modulated result in the azimuth dimension with an azimuth reference function. Other techniques have been developed recently for processing SAR data that operate on the general 2-D signal [3]. These techniques include wavenumber-domain and chirp scaling methods, and are all variations on the idea of using a scaled inverse Fourier Transform [4] and [5]. This work examines the particular SAR processor used by the JPL Airborne SAR group, which is of the RD type. It is an early version of the JPL Integrated Processor (JPLIP), where *integrated* refers to the ability to processes polarized as well as interferometric data.

Within the last decade, Interferometric SAR (INSAR) has emerged as an important extension of conventional SAR. The contribution of the INSAR technique is that it can be used to infer topographic heights of the imaged terrain. This is possible because the terrain is imaged by two antennas (or sometimes one antenna on a repeat pass). Using in-phase/quadrature demodulation, each antenna records a complex-valued image. The relative phases from those two complex images can be related to the height of the terrain using the known imaging geometry [6].

The demand for high-resolution, large coverage-area topographic data is driving new INSAR systems to acquire data in long strips. To date, most SAR literature has discussed processing techniques in terms of their ability to minimize phase errors within a single synthetic aperture. However, the literature has not addressed phase errors associated with processing long, continuous strips of INSAR data because most current systems can only processes relatively small data sets, for which it is adequate to process the entire scene to one set of Doppler and motion parameters. When processing long strips of INSAR data, the Doppler and motion parameters vary too rapidly to use one set of values. Processing the strips in patches, each with its own set of parameters, becomes necessary. The NASA/JPL Topographic SAR (TOPSAR) instrument acquires INSAR data in these long strips.

In a post-processing operation known as deskew, the data patches output from the processor are aligned (equivalent to re-indexing the pixels) so that the final strip image appears continuous [7]. This technique is sufficient to rectify SAR magnitude images, but systematic errors are still visible in the INSAR data because of their greater sensitivity to the phase measurement. These discontinuities occur at regular intervals between the patches of processed data written to the output arrays. In a typical TOPSAR topographic image, these patch-wise height errors appear as a sawtooth function with a period of roughly 1 km in the azimuth dimension and peak-to-peak amplitudes of ± 1 m, as shown in Figure 2. These errors are within the TOPSAR design specifications and are only visible when INSAR data are collected over low-relief (flat) regions. Errors of this size are often referred to as being at the "artifact level" in the literature and are therefore rarely addressed. However, low-relief areas usually correspond to flood-prone regions, where the need for precise knowledge of the topography for accurate flood forecasting is the greatest. If INSAR data are to be successfully applied to low-relief regions, these errors must be addressed. A data smoothing procedure known as presumming can reduce these errors by rejecting potentially aliased frequencies in the azimuth spectrum leading to more accurate pixel phase values in the deskew operation.

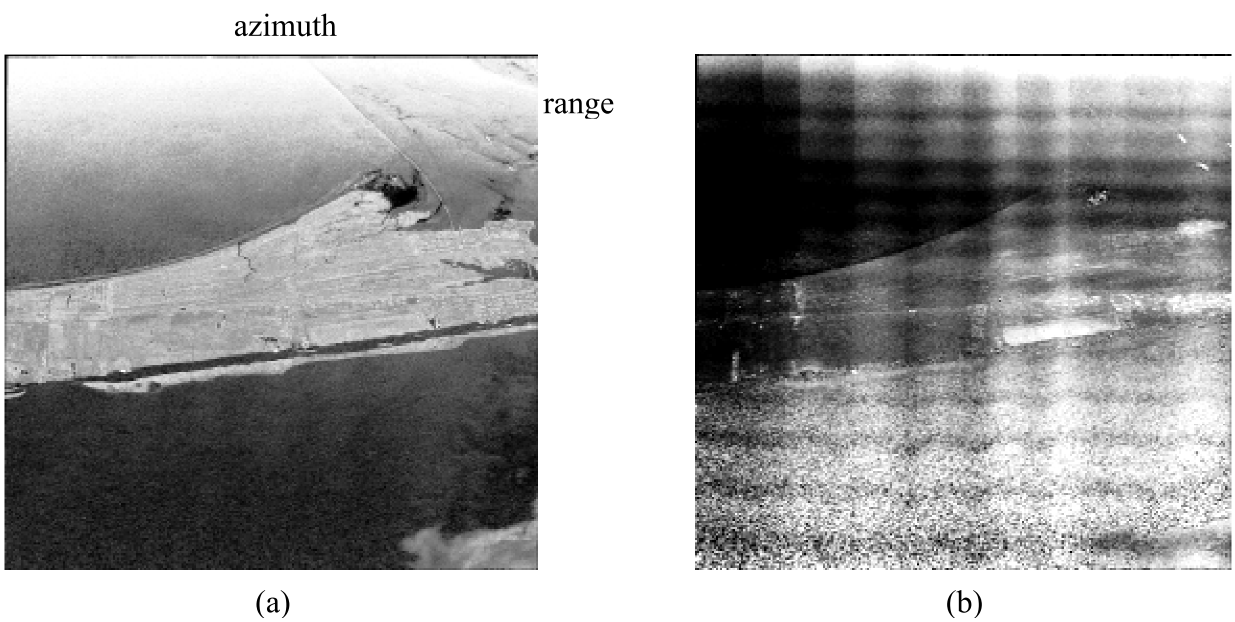

Fig. 2. 10 km x 10 km imagery acquired by TOPSAR over the coast of Texas, USA: (a) magnitude image and (b) INSAR image (topography). Patch boundaries are clearly visible in the INSAR data.

2. Nominal Processing

2.1 SAR Simulation

As mentioned in the introduction, RD processing of raw SAR data consists of 1-D match filtering (equivalent to correlating) the data in two dimensions. Each received pulse is convolved with a time-reversed, time-shifted replica of the transmitted pulse in a process known as range compression. A correction for the fact that the range to a target varies while that target is in view is made, and then the range-compressed data are match filtered in azimuth in an operation known as azimuth compression. A point-target simulator was used to study these effects. The simulator consists of a MATLAB program, and served as a useful tool to examine the signals at various stages of processing. Figure 3(a) shows the simulated SAR response to a single point target after range and azimuth compression.

2.2 Implementation

The version of the JPL Integrated Processor (JPLIP) used for this work consists of several FORTRAN-77 programs, include files, command files, and input files. The programs are hardcoded to read and write to files in a specific directory structure. Submitting a processing job is done by invoking several Unix scripts. The key steps of range and azimuth compression are done in a single "core" program. The convolutions required for the match filtering are carried out in the frequency domain, where the appropriate subsets of the data arrays and the reference functions are retained so that the Fast Fourier Transform (FFT) calculations are equivalent to linear convolutions.

Several post-processing programs subsequently deskew the data and form the topographic images. Producing SAR magnitude imagery and a topographic image for a 10 km by 3 km scene requires roughly 2 hours on an HP-9000 Unix workstation. Issues regarding segmenting the data in azimuth depend on azimuth compression, so the presumming operations were implemented in the azimuth compression subroutine.

3. Presumming

Desired range resolutions do not typically allow degradation in that dimension. However, it is common in SAR systems that the pulse repetition frequency oversamples the target relative to the desired azimuth resolution. The logical place for data reduction is therefore in the azimuth dimension. Azimuth data volume can be reduced by processing only a portion of the bandwidth of the azimuth spectrum. This practice, referred to as presumming, involves lowpass filtering the azimuth spectrum and then downsampling by an integer factor [6] and [8]. The lowpass filter serves to bandlimit the signal so that there is no aliasing when the spectrum is expanded through downsampling. Figures 3(b) show the effects of presumming on the point response shown in Figures 3(a). Presumming reduced the azimuth sidelobes of the magnitude spectrum and slowed the variations in the phase spectrum.

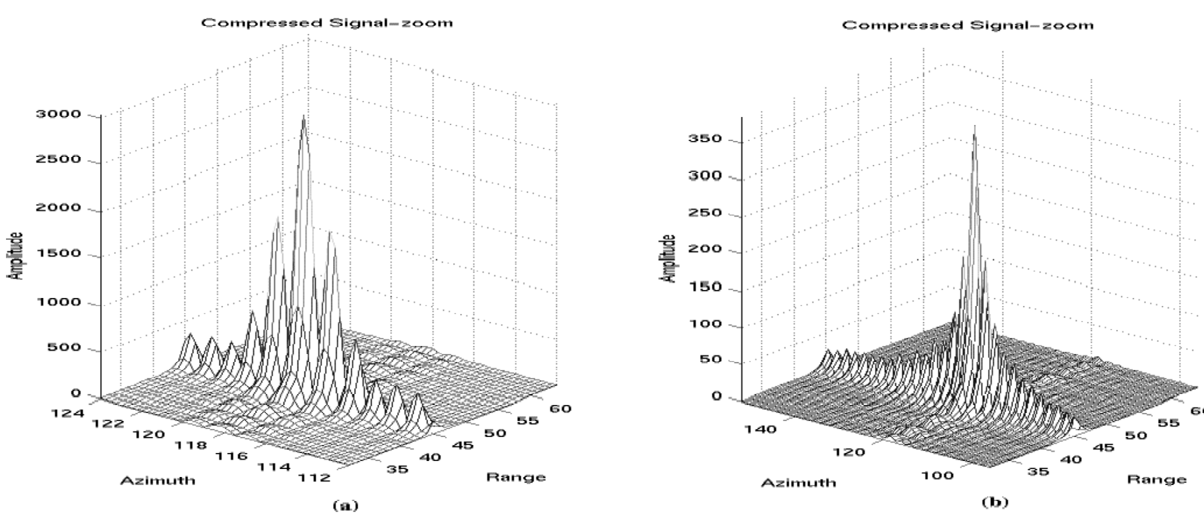

Fig. 3. Point simulator output: (a) nominal and (b) presummed

Most references to presumming in the literature discuss it as a means for reducing the computation time and storage requirements in a SAR processor [6], [8], and [9]. For these applications, presumming should occur inside the azimuth compression subroutine, just prior to convolving with the azimuth reference function. Presumming has also been used for achieving a desired data spacing relative to the motion and navigation data [10]. For that application, it was implied that presumming should be done after range compression, but before motion compensation or azimuth compression. The contribution of this work is to study presumming for reducing patch boundaries. The presumming operations were implemented after motion compensation and just prior to azimuth compression. This was done for two reasons: (1) this is where most of the literature implies presumming should be done and (2) presumming before motion compensation would have required many new arrays with different dimensions than the nominal arrays. Because many of the array sizes in the JPLIP are passed to the various programs independently, all new arrays would have to be hardcoded in several different places in the code.

4. Results

Figure 4 shows the JPLIP output after presumming. The smoothing in azimuth did reduce the patch boundaries slightly, but the basic misalignment error remained. The lowpass filter consisted of a real-valued exponential function in the frequency domain. This choice was made because it was simple to implement and ensured that the filtering did not alter the phase of the complex-valued data, which is critical to INSAR processing. For this example, a downsampling factor of 1 was used. The lowpass filtering was carried out directly on the azimuth spectrum so that the only additional computation was computing the frequency-domain filter values. Because it was desired to work with the existing data arrays, downsampling was achieved by inserting zeros for every multiple of the downsampling factor in the azimuth response (in effect, upsampling a downsampled sequence). Since zero insertion introduces high frequency content into a signal, the result had to be transformed to the frequency domain, filtered by another lowpass filter, and then transformed back [11]. The two additional 2-D FFTs accounted for roughly 50 additional multiplies per output sample and roughly 30 minutes of additional run time [12].

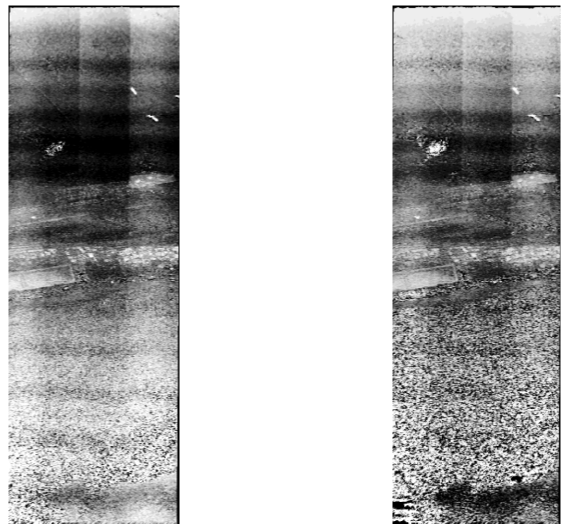

(a) (b)

Fig. 4. 10 km x 3 km subsets of the INSAR data in Figure 2: (a) original result and (b) presummed result.

When presumming is used for data reduction, downsampling factors of 10 to 1 or greater are common, but for this work a factor of 2 to 1 was used [9]. The nominal azimuth resolution after azimuth compression is roughly 3 m in TOPSAR data. However, the JPL deliverable products are averaged down to a resolution of roughly 10 m through a process known as multilooking. Therefore downsampling factors of up to 3 to 1, combined with lowpass filtering, could be used with negligible loss of information in the final products.

5. Conclusions

The topographic image in Figure 4 exhibited smoothing in the azimuth direction. This smoothing blurred the patch boundaries, making them less distinct. However, it did not specifically improve the amount of height change introduced at the patch boundaries. It was originally thought that presumming just prior to azimuth compression would help the patch boundary errors by reducing aliased frequencies in the azimuth spectrum and providing a more slowly varying azimuth phase spectrum. This would then result in an accurate re-indexing of the pixels in the deskew operation. However, after implementing the presumming, it was determined that to affect the indexing of the pixels so that the patch boundaries line up properly, presumming should be carried out before motion compensation.

As this research is continued, extensive changes will be made to the JPLIP processor so that presumming may be implemented prior to motion compensation. The manner in which array sizes are declared and passed among the programs will be changed so that changing array dimensions through resampling will be easier. This will eliminate the need for zero insertion.

References

[1] J. Curlander and R. McDonough, *Synthetic Aperture Radar: Systems and Signal Processing*, John Wiley & Sons, ISBN 0-471-85770-X, 1991.

[2] R. Bamler, "A Comparison of Range-Doppler and Wavenumber Domain SAR Focusing Algorithms," *Trans. on Geosci. and Remote Sensing*, vol. 30, no. 4, pp. 706-713, July 1992.

[3] R. Lanari, "A New Method for the Compensation of the SAR Range Cell Migration Based on the Chirp Z-Transform," *Trans. on Geosci. and Remote Sensing*, vol. 33, no. 5, pp. 1296-1299, Sept. 1995.

[4] R. Raney, H. Runge, R. Bamler, I. Cumming, and F. Wong, "Precision SAR Processing Using Chirp Scaling," *Trans. on Geosci. and Remote Sensing*, vol. 32, no. 4, pp. 786-799, July 1994.

[5] R. Lanari and G. Fornaro, "A Short Discussion on the Exact Compensation of the SAR Range-Dependent Range Cell Migration Effect," *Trans. on Geosci. and Remote Sensing*, vol. 35, no. 6, pp. 1446-1452, Nov. 1997.

[6] G. Morris and L. Harkness (ed.), *Airborne Pulsed Doppler Radar*, Artech House, ISBN 0-89006-867-4, 1996.

[7] A. Chu, Y. Kim, J. van Zyl, W. Tung, E. O'Leary, and Y. Lou, "The NASA/JPL AIRSAR Integrated Processor," *Proc. Int. Geosci. and Remote Sensing Symp.*, June 1998, pp. 1908-1910.

[8] D. Massonnet, F. Adragna, and M. Rossi, "CNES General-Purpose SAR Correlator," *Trans. on Geosci. and Remote Sensing*, vol. 32, no. 3, pp. 636-643, May 1994.

[9] W. Carrara, R. Goodman, and R. Majewski, *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*, Artech House, ISBN 0-89006-728-7, 1995.

[10] P. Rosen, Memorandum 3343-92-90, Jet Propulsion Laboratory, 1992.

[11] A. Oppenheim and R. Schafer, *Discrete-Time Signal Processing*, Prentice Hall, ISBN 0-13-216292-X, 1989.

[12] D. Dudgeon and R. Mersereau, Multidimensional Digital Signal Processing, Prentice Hall, ISBN 0-13-227638-0, 1984.