IFSAR for the Rapid Terrain Visualization Demonstration

Bryan L. Burns, Paul H. Eichel, William H. Hensley, Theodore J. Kim
Sandia National Laboratories
Albuquerque, NM 87185

Abstract

The Rapid Terrain Visualization Advanced Concept Technology Demonstration (RTV-ACTD) is designed to "demonstrate the technologies and infrastructure to meet the Army requirement for rapid generation of digital topographic data to support emerging crisis or contingencies". The primary sensor for this mission is an interferometric synthetic aperture radar (IFSAR) designed at Sandia National Laboratories. This paper will outline the design of the system and its performance, and show some recent flight test results.

The RTV IFSAR will meet DTED level III and IV specifications by using a multiple-baseline design and high-accuracy differential and carrier-phase GPS navigation. It includes innovative near-real-time DEM production on-board the aircraft. The system is being flown on a deHavilland DHC-7 Army aircraft.

1. Introduction

The United States Army has identified a requirement to rapidly produce fine-resolution Digital Topographic Elevation Data (DTED) to support operations. The need for rapid data generation is particularly keen for the finest resolution data sets -- those providing data on 10-m and finer posts. The cause is two-fold. First, it is currently impractical to collect and store a very high-resolution database of the entire world. Second, surface changes on the scale of 10 meters or less occur frequently. Such changes, due to both natural and human activity, give high-resolution databases a very short shelf-life.

This paper describes the overall system design and signal processing for a system developed at Sandia National Laboratories (SNL) to provide extremely rapid generation of digital terrain data at fine-resolution DTED levels.

2. System specifications

The RTV system produces data at DTED level III or DTED level IV. At either level the system produces four products: an orthorectified image, a digital elevation model (DEM), a coherence image, and a quality image.

The orthorectified image is sampled with one-fourth the sample spacing of the other three products, which all have the same sample spacing. The system is designed so that measurements at each individual post of the DEM are independent. The quality image indicates the estimated height accuracy of the DEM product based on radar data. All of the output products are saved in GeoTIFF files for dissemination. Table 1 shows the requirements for the RTV IFSAR system at both DTED levels.

Table 1. Product specifications for the RTV IFSAR system.

	Post Spacing		Absolute		Relative	
Level	ortho m	DEM m	CE90 m	LE90 m	CE90 m	LE90 m
III	2.5	10	10	10	3	2
IV	0.75	3	10	5	2	0.8

The IFSAR system is designed to meet these requirements for terrain with radar reflectivity higher than −15 dB. Lower reflectivity terrain than −15 dB is still measured; however, with less height accuracy than shown in Table 1.

Although real-time processing was not a system requirement, the desire to rapidly produce the output products lead to a system that performs most of the processing in real time.

3. Overall system design

The RTV system contains four important innovations as a part of its design. 1) Traditional phase unwrapping is avoided using a multiple-baseline approach. 2) Phase modulation due to multipath off of the aircraft body is mitigated by a combination of a shroud around each antenna and a diffraction grating around the radome on the outside of the aircraft. 3) Image formation and IFSAR processing are done on data at real-time rates. 4) Mosaicking of patches of imagery collected over a period of time reduces unavoidable time varying navigation errors.

Typically IFSAR systems use two antennas separated in elevation to produce estimates of terrain elevation. Phase difference information at the two antennas is used to estimate an incoming signal's angle-of-arrival at the

antennas. Unfortunately, this phase information is ambiguous if the antennas are far enough apart in elevation to produce high-accuracy DEMs. Traditionally, 2-D phase unwrapping techniques have been used to estimate the number of ambiguities in the data to produce the final DEM. While these techniques have been shown to work in many areas, there are many examples of terrain that do not unwrap correctly. Traditional 2-D phase unwrapping is eliminated by a new technique developed for RTV.

In RTV an elevation monopulse antenna replaces one of the two IFSAR antennas. The elevation monopulse antenna is designed so that no angle-of-arrival ambiguities occur within the elevation beam of the antenna. While the monopulse antenna is unambiguous, the precision of such a solution alone is insufficient to meet DEM requirements from reasonable ranges. To achieve an unambiguous and accurate solution, angle-of-arrival information from the monopulse antenna can be used to resolve the elevation ambiguities present in the IFSAR system. In this case the IFSAR result is formed using data from the sum port of the monopulse antenna and the second antenna. For best phase matching between the antennas the second antenna also has a monopulse feed identical to the first; however, only the sum port is used while the difference port is terminated. The result is a three phase-center system, two for monopulse and one for the second antenna. We refer to these signals as $\Sigma A1$, $\Delta A1$, and $\Sigma A2$. This combination may also be thought of as having two baselines. The small baseline system has no ambiguities within the elevation beam and the larger baseline system has many. The ratio of the two baseline lengths must be properly constrained so that the monopulse angle-of-arrival answer will have sufficient precision to correctly estimate the number of cycles of ambiguity present in the IFSAR data. This approach requires three channels of image formation rather than two. The effectiveness of the monopulse solution is shown in the section on results below.

Many IFSAR systems have experienced difficulty controlling the phase response of the system as a function of elevation angle. The lack of control in the phase response of the system usually puts "waves" of height error into the DEM that can change as a function of flight conditions. We have noted this phenomenon in our previous systems and traced the cause to low-amplitude signals reflected from the body of the aircraft that add with the direct-path signal. These multi-path signals cause phase errors that produce an unpredictable system phase response as a function of the elevation angle. In RTV we placed a radar-absorbing shroud around each antenna and a diffraction grating on the body of the aircraft to significantly reduce the amplitude of the signals reflected from the body of the aircraft (Figure 1). After critical examination of a large quantity of data, we have been unable to observe any undesired phase errors caused by

the multi-path phenomenon. Phase difference calibration as a function of elevation angle closely matches preflight antenna range measurements.

Figure 1. Antenna with shroud (left) and diffraction grating (right).

Real-time signal processing in the RTV system is broken into three primary pieces that are shown in Figure 2. Real-time image formation implements Overlapped Subaperture (OSA) [1] image formation on all three channels of imagery simultaneously. The parameters of image formation are set to overlap adjacent patches of imagery to assist mosaicking the final products. Once all three images are formed, IFSAR processing combines the information from all 3 channels into a preliminary set of data products that include an ortho image, a DEM, a correlation map, and a quality map for each patch of data. IFSAR processing includes the calculations required to remove ambiguities in the large baseline answer based on monopulse measurements.
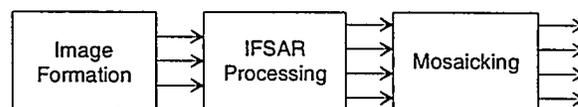

Figure 2. Top-level signal processing flow.

The system is designed to collect large areas by flying a set of stripmaps that are mosaicked together in the final stage of processing. Each stripmap is implemented as a collection of overlapping spotlight images. Adjacent patches are overlapped in azimuth by ten percent. Patches from adjacent lines are overlapped by ten percent in range to facilitate the mosaicking process.

A top-level functional block diagram is shown in Figure 3. While 3 antennas are shown, there are really only two antennas. Waveguide circuitry on the first antenna's monopulse feed horn adds and subtracts the microwave energy from the elements producing a sum and difference output. High-speed timing components and almost all of the RF hardware are mounted in VME 1. Motion measurement processors, a motion compensation and radar control processor, and processors that form two channels of real-time imagery are housed in this chassis as well. VME 2 houses processors for the third channel of real-time image formation and all of the IFSAR processing. Communication between VME 1 and VME 2 is done over a wide-bandwidth fiber channel interface. Image formation processing and IFSAR processing are implemented on a set of forty 300-MHz Power PC

DISCLAIMER

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

processors. Mosaicking processing is performed on the DEMS Workstation, which is a Sun Sparcstation 450 that has four 360-MHz CPUs.
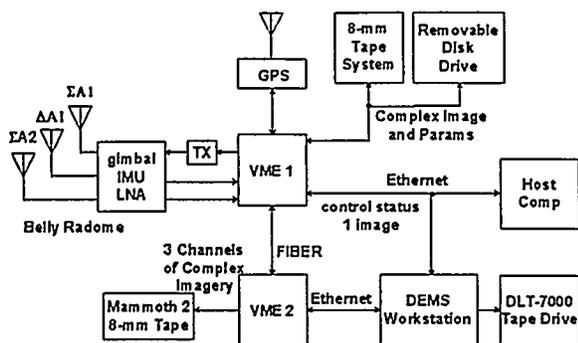

Figure 3. RTV functional block diagram.

Several parameters of the overall system are shown in Table 2. The elevation monopulse is done using amplitude comparison monopulse and the equivalent interferometric baseline can be computed as shown in Table 2. The ratio of the two baseline lengths must be constrained in order to avoid jump-cycle errors and is discussed further in section 7 below.

Table 2. General system parameters.

Parameter	Value	Units
Ground Speed, maximum	150-230	knots
Aircraft altitude	18-23	k ft
IFSAR Range, minimum	3.25	nm
IFSAR Range, maximum	6	nm
Radar center frequency	16.7	GHz
IPR 3 dB width specification	±10	%
IPR peak sidelobe amplitude	-30	dB
IFSAR Squint angle	90	deg.
σ_0 of clutter for height noise	-15	dB
required σ_n	-35	dB
Antenna beamwidth, azimuth	2.7	deg.
Antenna beamwidth, elevation	5.1	deg.
Antenna gain	32.7	dBi
Monopulse baseline length	0.038	m
Large baseline length	0.330	m
Peak transmitter power	700	W
Transmitter duty factor, max	35	%
Radar PRF	0.8-2	kHz
Presum, each channel	2	

The system parameters are adjusted depending on DTED levels as shown in Table 3.

Spatial accuracy requirements of the system are maintained by a combination of two components of the system: a GPS aided inertial measurement system and mosaicking computations.

Table 3. Radar system parameters for DTED III and IV.

Parameter	Level III	Level IV	Units
Nominal grazing angle	32.78	45	deg.
Range resolution	0.9144	0.3048	m
Azimuth resolution	1.067	0.4572	m
Azimuth pixel spacing	0.9063	0.381	m
DEM sample spacing	10	3	m
Ortho image spacing	2.5	0.75	m

4. Measuring the position and attitude of the antenna

The motion of the IFSAR antennas is measured using an inertial measurement unit (IMU) and two GPS receivers. The IMU consists of the instrument sensor assembly and inertial electronics from a production Honeywell H764G inertial navigation system that has been repackaged at SNL with stiffer isolators and called the High Accuracy Inertial Measurement System (HAIMS). The HAIMS is mounted behind the radar antennas (Figure 1) on the innermost axis of the three axis gimbal. The 1200Hz outputs of this IMU are used to produce a geographic navigation solution at 100Hz.

A P-code Interstate Electronics Corporation (IEC) GPS receiver outputs pseudorange and accumulated carrier phase measurements at 1Hz and satellite ephemeris data every minute. These measurements are used directly in an extended Kalman filter to update the navigation solution [2]. This filter has been modified to accept real time kinematic (RTK) position information from a NovAtel differential GPS receiver. The IEC and NovAtel data are not time synchronized so the Kalman filter runs at 2Hz when both sets of data are available.

The updated navigation solution is used to generate pilot's guidance and gimbal pointing commands. The position-based pointing commands are passed to the gimbal controller, which runs at 1200Hz and has its own estimate of attitude based on the high rate IMU data. The controller is based on near-time optimal control and generates motor drive commands based on pointing errors and its estimates of gimbal axis friction.

In addition to the 100Hz geographic navigation solution, the IMU data is also used to produce a 400Hz navigation solution in the Cartesian radar frame. This solution is carefully coupled to the updated geographic navigation solution to obtain the benefits of the GPS measurements without the discontinuities caused by the Kalman filter. The position and velocity of the antennas, the attitude of the IFSAR baseline, and the estimated pointing errors are passed to the radar control processor at the PRF rate.

5. Multi-channel SAR signal processing

In order to use only two channels of RF hardware for the three RF signals (ΣA1, ΔA1, and ΣA2), time-domain multiplexing of two receivers is done. Figure 4 shows one cycle of a timing diagram for the RF circuitry. The symbol ΣA1TX refers to the timing of the ΣA1 transmitted signal, ΣA1RX to the ΣA1 received signal, etc. To maximize the effectiveness of the large baseline we transmit on both antennas.
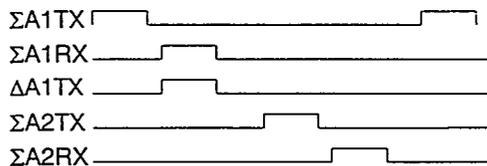

Figure 4. RTV RF timing diagram.

Two simultaneous receive paths are exercised on the first half of the cycle when both components of the monopulse antenna are received and quantized. Each triple-conversion receiver includes in-phase (I) and quadrature (Q) outputs at baseband and are quantized by two 10-bit A/D converters to maintain high dynamic range.

Careful design is inadequate to maintain the desired gain and phase alignment between the two receiver channels and the A/D converters. A full-bandwidth calibration path is used prior to each pass to measure the relative gain and phase of the receivers and the I/Q balance of the channels. In-flight testing of the system has shown that this approach maintains phase calibration between the three channels to much better than 1 degree of phase and the opposite sideband rejection of the I/Q calibration is better than -55 dB.

Three channels of phase history data are sent to different Mercury processors for real-time image formation processing. Image formation processing removes the phase term caused by time-domain multiplexing shown in Figure 4.

6. Interferometric signal processing

RTV represents a radical departure from previous interferometric SAR systems in many respects. In this section, we will summarize several key innovations that distinguish this system. In addition to the use of amplitude monopulse mentioned above, these include mapping of spotlight-mode data to absolute coordinates, and least-squares mosaicking of the spotlight-mode patches to the final output map.

It should be reemphasized that all of these processes are accomplished in real time on board the RTV aircraft. The signal processing proceeds on an heterogeneous, multi-processor architecture. The tasks of amplitude monopulse, ifsar phase, ambiguity resolution, scatterer location in 3-D, and gridding to a map projection all run on Mercury Power-PC nodes in VME 2 shown in Figure 2. The gridded patches are then transferred to a multi-processor SUN workstation for mosaicking into the final map. The processing proceeds with no human intervention of any sort, must keep up with the peak area coverage rate of the aircraft (~10 km^2 per minute at DTED III and ~3.5 km^2 per minute at DTED IV), and is finished when the aircraft lands.

7. Amplitude monopulse for ambiguity resolution

In many ways, the need for phase unwrapping the interferogram of principal values of phase has been the Achilles Heel of IFSAR. No single algorithm has emerged that is capable of produced the "right" unwrapped phase function in all situations. Indeed, where the phase has been spatially aliased, the ambiguity of cut-off regions of the interferogram relative to the rest is unknowable. As a practical matter, no matter what the system design, there are naturally occurring situations that defeat every 2-D phase unwrapping algorithm.

The RTV IFSAR system does not employ phase unwrapping. Instead, a third radar channel is used to produce unambiguous, albeit noisy, estimates of elevation angle for every scatterer. These estimates, in turn, resolve the cycle ambiguity of the interferogram on a point-by-point basis, thus avoiding the need for an unwrapping algorithm entirely [3]. Since every scatterer is treated independently, there can be no such thing as cut-off regions or artifacts resulting from spatially aliased data.

Like other IFSAR systems, the RTV platform has two physical antennas displaced orthogonal to the velocity vector. However, one of the antennas has two feeds and two independent radar processing channels. The two feeds are also arranged orthogonal to the velocity vector so that their boresights are separated in elevation angle. A given scatterer therefore is seen with a different amplitude by the two beams, depending on its elevation angle. That is, if the scatterer falls above the boresight of the antenna, it appears brighter in the channel associated with the lower feed and vice versa. By careful matching of gain and phase through both channels, it is therefore possible to estimate elevation angle unambiguously by taking the ratio of a scatterer's amplitude in the two channels and normalizing by their sum. To be sure, this is a noisy estimate of elevation angle, but is sufficient to resolve the cycle ambiguity of the scatterer's interferometric phase. The noise performance of the technique (known as amplitude monopulse) is enhanced by multilooking, much as is done for the interferometric phase.

The actual cycle ambiguity resolution of the interferogram is simple and fast. It is perfomed on a point-by-point basis. Let $\bar{b}$ represent the interferometer baseline vector, θ_{mono} the (unambiguous) elevation angle measured from the boresight of the antenna at the base of $\bar{b}$, ϕ the principal value of phase in the interferogram, and Φ the desired unambiguous phase. We have the standard IFSAR relationship:

$$\Phi = \frac{4\pi|\bar{b}|}{\lambda}\cos(\alpha) \qquad [1]$$

where α represents the angle between the line-of-sight vector to the scatterer and the baseline. If the interferometric baseline is orthogonal to the antenna boresight (a condition met by the active antenna steering of a spotlight-mode SAR), then α is merely the complement of θ. We can therefore *estimate* Φ from θ_{mono} as:

$$\hat{\Phi} = \frac{4\pi|\bar{b}|}{\lambda}\cos\left(\frac{\pi}{2} - \theta_{mono}\right) \qquad [2]$$

so a nearest-neighbor estimate of the cycle ambiguity number k is simply:

$$k = \left\lfloor \frac{\hat{\Phi} - \phi}{2\pi} + 0.5 \right\rfloor \qquad [3]$$

and the resolved ("unwrapped") phase is:

$$\Phi = \phi + 2k\pi \qquad [4]$$

This procedure remains error free as long as the *jump cycle error* criteria:

$$\left|\hat{\Phi} - \Phi\right| > \pi \qquad [5]$$

is not met. For small errors near boresight, this is equivalent to:

$$\left|\alpha - \left(\frac{\pi}{2} - \theta_{mono}\right)\right| > \frac{\lambda}{4|\bar{b}|} \qquad [6]$$

For RTV, various system parameters and performance specifications ensure that jump cycle errors have a probability of less than 1×10^{-6} at a clutter cross section of −22 dB. It is important to note, however, that even when jump cycle errors occur, they remain isolated and do not propagate, as do 2-D unwrapping errors.

8. Mapping to absolute coordinates

Following ambiguity resolution, the terrain height and backscatter amplitude imagery are mapped to an absolute coordinate system in a two step process. First, the raw radar measurements of range, doppler cone angle, and unambiguous interferometric phase are used to compute scatterer location in 3-D geocentric coordinates. This step is accomplished using the efficient direction cosine method, adapted for spotlight-mode SAR. See [4] for details of this method. It should be noted that the aperture phase center positions, baseline vector, and velocity vector all need to be known to a great deal of precision so as to keep absolute errors below specification in this process. See Section 5.

Second, these 3-space triples in geocentric coordinates are interpolated onto a desired map grid. For RTV, this grid is a UTM map projection with height expressed in meters above earth ellipsoid (WGS-84). Since the geocentric triples are irregularly spaced, this means that the UTM output posts are computed over a variable number of input triples. A distance-weighted interpolator is used for this process. Also, geocentric triples arising from IFSAR phase samples exhibiting low coherence or low radar backscatter are inhibited from the interpolation process. A void-filling algorithm follows the UTM projection to fill in small voids.

In addition to the DEM, the backscatter amplitude imagery is also mapped to a UTM grid. This orthorectified imagery has none of the layover, skew, or other spatial distortions normally associated with SAR imagery. RTV interpolates the gridded height data so as to produce orthorectified imagery at sixteen times the sampling rate of the DEM. The standard output map grids for RTV are shown in Table 1 above.

9. Mosaicking

Because a spotlight-mode SAR produces imagery in patches, the gridded DEM and Ortho products of the interferometric processing must be mosaicked together to construct large (15 x 15 minute Level 3 or 5 x 5 minute Level 4) final maps. The innovative least-squares technique developed for RTV not only accomplishes nearly flawless mosaics, it also plays a fundamental role in meeting the absolute error specifications for the system as a whole. The technical details of this method will be addressed in a future publication; here we will provide a brief outline.

Patches of data arrive at the mosaic processor already mapped to a common (UTM) coordinate system. Both orthorectified imagery and DEM data are available. We assume that the data in each patch is correct in absolute placement save for a 3-D displacement and a 2-D tilt. The displacement errors are largely attributed to GPS position errors and may be weakly or not at all correlated from patch to patch. The tilt errors are largely attributed to errors in attitude, principally knowledge of the correct local gravity vector. These tilt errors are assumed to be common to patches comprising a map. It is the job of the mosaicking process to determine the displacements

required to sum with each patch and the tilt required to sum with each map so as to produce a seamless final product.

While it is impossible to discern from the data themselves what the individual patch displacements are, we can determine the partial first differences of the patch displacements. By correlating each ortho image patch with its four neighbors in the areas where they overlap, we can directly measure the partial first differences in easting and northing. Having found those, we can compare the DEM patch with its four neighbors to determine the partials in height. Thus we determine the six partials of easting, northing, and height displacements with respect to easting and northing. These measurements are then inverted via a partial differential equation solver to obtain a least squares estimate of the 3-D displacements of each patch. Tilts manifest themselves as nonzero means in the first partials of height with respect to easting and northing. These are easily removed from the DEM data. The final maps are constructed by offsetting the data of each patch by its computed 3-D displacement and interpolating the data into a final map grid.

This procedure not only produces more pleasing final maps without noticeable artifacts and seams, it also improves the absolute positional accuracy of the data. The least squares estimation procedure reduces the variance of absolute displacement errors by the number of independent measurements. For example, if each patch represented an independent, zero mean sampling of GPS position, then the absolute error variance of the mosaic would be equal to that of any one patch divided by the number of patches. While the GPS errors are somewhat correlated from patch to patch, a typical 5 or 15 minute final map represents several hours of data collection. Over this time span, the final mosaic represents many hundreds of independent measurements with commensurate reduction of absolute error variance.

10. Results and conclusions

At this writing we have collected data over 5 different geographic areas totaling approximately 250 km^2. One of those areas is known as Twin Lakes, New Mexico because of a small town located there. Figure 5 is a DTED III orthorectified image of a small portion of the entire area with a chip from the upper left corner shown on the right. Figure 6 is a DTED IV orthorectified image of the same area with a chip of the upper left corner on the right. The chips best illustrate the significant difference between the two sample spacings (2.5 m and 0.75 m respectively). Figure 7 contains a gray scale, shaded relief of the DTED III DEM and Figure 8 contains a gray scale, shaded relief of the DTED IV DEM. In both of these figures the terrain slope has been exaggerated by a factor of 3 to enhance presentation of the terrain relief. In one piece of flat

terrain near the town of Twin Lakes, the total height variation (due to both terrain effects and errors) has a standard deviation of 0.2 m. This corresponds to an LE90 relative accuracy of 0.32 m, significantly better than the 0.8 m requirement. The Twin Lakes area does not have ground truth, so no absolute error comparison has yet been made for DTED Level IV.

A second area is a 7.5-minute tile including Kirtland Air Force Base in Albuquerque. NM. This area was collected at DTED Level III – it covers reference reflectors and ground truth maps produced from optical stereo. Comparison with 10,000 posts of the stereo-optical DEM shows a relative height error of 1.5-m LE90, and an absolute height error of 2.9 m. There are 18 surveyed reference reflectors in the scene also -- spread over most of the 7.5-minute tile. Comparison of the DEM with the reference reflectors gives a 3.7-m height offset, and a 0.66-m LE90 relative height accuracy. The mean absolute errors in easting and northing are 4.7 m and 6.7 m, respectively. These mean errors are well within the specification for using P-Code GPS alone, but they cause the absolute CE90 spec to be exceeded. The NovAtel differential GPS is not yet installed – it is expected to remove most of the absolute position error.

Accuracy results to date are summarized in Table 4. Items that have not yet been measured are marked as "n/m". It is important to note that the results are _not_ statistical – each one has been measured only once.

Table 4. Comparison of Ground Truth to DEMs

DTED Level	Type of Comparison	Absolute		Relative	
		CE90 m	LE90 m	CE90 m	LE90 m
III	Specification	10	10	3	2
	DEM–StereoDEM	n/m	2.9	n/m	1.5
	DEM–Reflectors	12.6	3.8	2.7	0.7
IV	Specification	10	5	2	0.8
	Height Variation	n/m	n/m	n/m	0.3

Integration of the differential GPS, as well as further calibration and statistical testing of the system are to be finished in 2001.

11. Acknowledgments

References

[1] B. L. Burns and J. T. Cordaro, "A SAR image-formation algorithm that compensates for the spatially-variant effects of antenna motion," Proc. SPIE, vol. 2230, pp. 14-24, 1994.

[2] Kim, T. J., et. al., "An Integrated Navigation System Using GPS Carrier Phase for Real-Time Airborne Synthetic Aperture Radar (SAR)", Institute of Navigation 55th Annual Meeting, Boston, MA, June 28-30, 1999.

[3] C. V. Jakowatz, Jr., D. E. Wahl, P. A. Thompson, "Ambiguity resolution in SAR interferometry by use of three phase centers," Proc. SPIE, vol. 2757, pp. 82-91, 1996.

[4] P. H. Eichel, "The direction cosine method of scatterer location extended to spotlight-mode IFSAR," 34th Asilomar Conf. On Signals, Systems, & Computers, Oct. 2000.
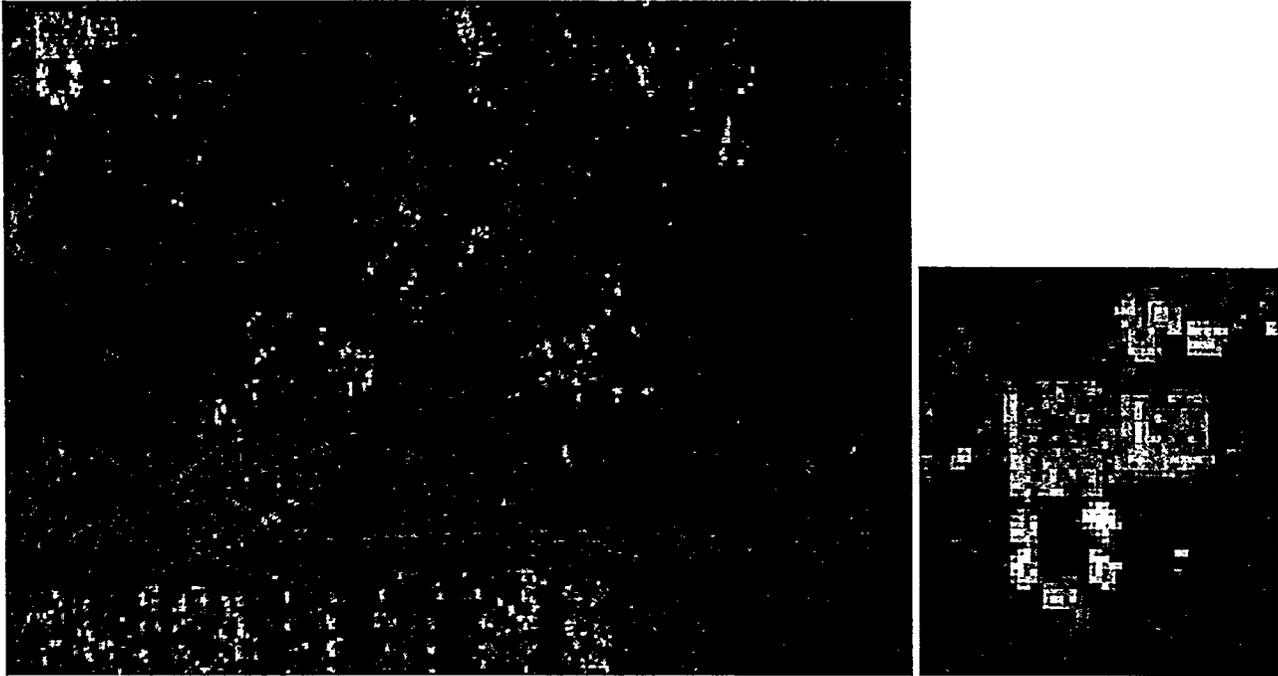
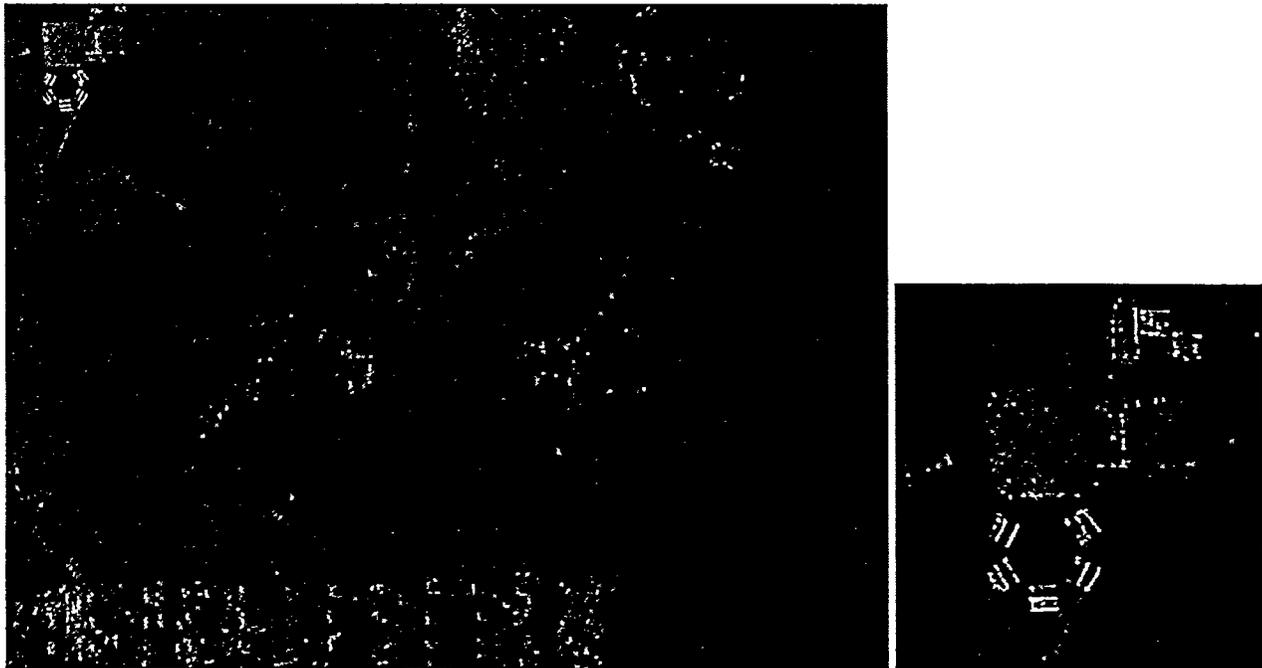**Figure 5. DTED III ortho image of Twin Lakes.**

Figure 6. DTED level IV ortho image of Twin Lakes.

Figure 7. DTED III Shaded relief of Twin Lakes.

Figure 8. DTED Level 4 DEM of Twin Lakes.