

Ground Processing of Cassini RADAR Imagery of Titan

Bryan W. Stiles, Yonggyu Gim, Gary Hamilton, Scott Hensley, William T. K. Johnson, Joanne Shimada, Richard D. West, and Phil Callahan

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA

Abstract--The Cassini RADAR instrument on-board the Cassini Orbiter is currently collecting SAR Imagery of the surface of Saturn's largest moon, Titan. This paper describes the ground processing of Cassini SAR data. We focus upon the unusual features of the data and how these features impact the processing. We exhibit a data dependent mechanism we have implemented for eliminating artifacts due to attitude and ephemeris knowledge error. Finally we describe how we trade-off SAR performance vs. area of coverage when we design our spacecraft pointing profiles.

I. INTRODUCTION

The Cassini RADAR instrument is a Ku-band RADAR transmitting at 13.78 GHz (2.17 cm wavelength). In addition to the SAR mode (935 kHz chirp bandwidth), it also operates in a scatterometer mode (BW=117 kHz), an altimeter mode (BW=4.68 MHz), and in a passive radiometer mode. Scatterometer and radiometer measurements are being collected for other icy Saturnian satellites, but SAR and altimetry data are only collected during Titan flybys. (Radiometer measurements of Saturn and Jupiter have also been obtained.) In contrast to all previous planetary and earth based SAR systems, the Cassini radar collects data on a hyperbolic trajectory as it flies by Titan while in orbit around Saturn. The closest approach of the Titan flybys on which the SAR is operated varies from 1000 to 2000 km. Typically SAR data is obtained while the spacecraft is less than 4000-5000 km from the surface. More distant SAR observations are currently under consideration for special cases. The Cassini radar is planned to generate SAR imagery for about 20% of the surface of Titan throughout the course of the Cassini mission. For more details on the Cassini Radar Instrument and its mission see [1].

Cassini SAR employs a burst-mode timing scheme. Many SAR systems operate in a continuous mode, which means that a train of chirped pulses is transmitted with regular time interval throughout the observation period. The interval between the pulses is such that returned echoes can be obtained between the transmitted pulses. The collection of pulses that contain reflected energy within the radar antenna's mainlobe, called the synthetic aperture, are processed in a coherent manner to obtain high along-track resolution in the resulting SAR image. The uncertainties in predicted Cassini spacecraft attitude and ephemeris are so large that they preclude the precise timing needed for continuous mode SAR. Even without such uncertainties our data rate and data volume constraints would make a continuous mode design unsuitable for the wide area coverage desired for the surface of Titan. For these reasons the Cassini radar utilizes a burst-mode SAR in which a train (burst) of 30-60 chirped pulses is transmitted followed by a long gap (about 400-800 pulses in length) in transmission until the return echo from the burst is received. After reception of the echo, the cycle repeats. The duration of the individual bursts are much shorter than the dwell time (the time over which a single point on the surface is within view of the antenna). Because the azimuth resolution of a SAR is inversely proportional to the dwell time, the burst-mode implementation provides lower resolution imagery than what would be obtainable from a continuous mode system. Although one could coherently process multiple bursts to overcome this limitation, Cassini SAR's large interburst gaps would result in unacceptably large grating lobes. Nonetheless, Cassini SAR offers more than an order of magnitude improvement over the real aperture resolution.

In order to increase coverage, Cassini employs five different antenna feeds (beams) that are separated in elevation. Four of these beams

are offset and elongated so that the beam width in elevation (cross-track direction) is roughly four times larger than minimum beam width obtainable with the antenna. Bursts are repeatedly transmitted sequentially from each of the five feeds. This configuration allows Cassini SAR to image a swath width that is 16 times wider than what could be achieved with the central narrow antenna beam alone. The extra coverage comes at the expense of five times fewer looks (spatial averaging of pixels is needed to reduce speckle noise) for each resolution element. The loss is most severe at closest approach to Titan where typically only three looks are obtained because of the rapid flyby.

The viewing geometry of Cassini RADAR is different from other imaging radars because the Cassini Orbiter flies by Titan during its orbit of Saturn rather than orbiting Titan itself. As a result, the SAR imaging geometry varies greatly resulting in both range and azimuth resolution that vary from 300 m to 2 km along the observed swath. Titan SAR coverage throughout the mission is limited by the need to share Titan flybys with sensors with incompatible imaging geometries. Thus to maximize SAR coverage of Titan's surface, SAR ambiguity levels and range/Doppler orthogonality are traded against area coverage. The range of viewing geometries requires a robust SAR processing algorithm. Additionally, simulation tools are used to insure the instrument commanding and pointing designs are consistent with adequate SAR performance.

This paper describes the Cassini SAR processor, providing a general overview of the algorithms and the special challenges required in the processing of two Titan flybys (TA, October 2004) and (T8, October 2005). TA was an early Titan flyby in which the first SAR observations were made. Being one of the first passes the Titan-relative spacecraft ephemeris and attitude knowledge errors were relatively large. Spacecraft pointing and location knowledge errors can severely impact the SAR imagery due to rapid variation in Doppler and range centroids and the necessity of imaging points far down on the antenna patterns in order to maximize swath area. Applying the predicted spacecraft ephemeris and attitude estimates obtained from the navigation team into the processing generally leads to radiometric artifacts throughout the image especially between the 5 separate antenna beams. Even though the navigation team has subsequently generated improved estimates of the ephemeris and attitude data for TA (and other Titan passes), the radiometric artifacts were

reduced but not sufficiently eliminated when processed with these data. Correcting the ephemeris and attitude information from the SAR data is difficult because of the number of unknowns including not only ephemeris and attitude errors but also knowledge of the shape and local topography of Titan. A data dependent method of adjusting the range and Doppler centroids used in azimuth and range compression was employed in the processor resulting in a significant improvement in image quality.

In the most recent flyby, T8, we intentionally designed the instrument pointing to maximize area of coverage at the expense of reduced SNR, reduced ambiguity levels, and non-orthogonal range and Doppler dimensions. The strategy employed to optimize the trade-off between SAR performance and area of coverage is described in Section IV. The effectiveness of this strategy is examined by evaluating the actual T8 imagery that was obtained.

II. SAR PROCESSOR OVERVIEW

This section provides a necessarily brief description of the Cassini SAR processing algorithm. So far no detailed description of the Cassini SAR processor has been published. Cassini SAR processing was based largely on previous work done for the Magellan SAR observations of Venus. A description of the Magellan SAR processor is available in [4]. For a detailed discussion of general SAR theory and the full range of issues involved in processing SAR data see [5].

A. *Synopsis and inputs to processor*

The Cassini SAR processing is performed at JPL subsequent to down linking of the raw echo data to Earth. No on-board SAR processing is performed. The inputs to the ground processor are 8 to 2 bit BAQ compressed real offset video sampled voltages, a copy of the parameters used to command the radar, and associated spacecraft clock time tags. The relevant time calibration, spacecraft attitude and ephemeris, and planetary ephemeris information necessary to process the data is provided to us by the Cassini Navigation and Attitude Control System teams. Prior to SAR processing, the data is passed through a preprocessor which converts encoded radar commanding parameters to engineering units, performs some initial geometrical computations, and decompresses the BAQ compressed echo

data. The SAR processor computes a single-look, complex-valued, radiometrically calibrated SAR image from each burst. The single burst images are interpolated using 2-D sinc interpolation onto an oversampled grid on the surface of Titan. The contribution from each pixel is “detected” by computing the square of the magnitude of the interpolated complex quantity. For each pixel in the grid, the detected backscatter values from all bursts that cover the pixel are averaged together to produce the final multi-looked image. Normalization of the incidence angle variation in the backscatter via a backscatter law and noise subtraction are performed as post-processing steps.

B. Pulse Segmentation and Range Compression

As previously described the Cassini radar transmits a series of pulses and records the returned echoes from the entire pulse train in a contiguous block of ADC samples. The first step in the SAR processing is to segment the echo data into separate pulses (necessary for the coherent processing of the synthetic aperture) and perform range compression on each pulse. The high duty cycle of our transmitted pulse train required a novel pulse segmentation algorithm. Using our knowledge of the location and attitude of the spacecraft relative to Titan we estimate the starting and ending time of the return echo for each burst. After omitting any pulses that do not fall completely within the range gate, there are N remaining pulses. We divide the burst train in N overlapping regions of equal length. The regions are necessarily overlapping because of the high 70% pulse duty cycle. The high duty cycle was chosen in order to maximize SNR, but it leads to some of the energy from consecutive pulses returning simultaneously. This energy comes from different ranges on the ground and is thus separated in the range compression process. This pulse segmentation therefore guarantees maximal SNR and simplifies calibration by recovering all energy from the returned pulses.

Range compression is performed by convolving each segment with a matched filter formed from the transmitted waveform. Cassini is an offset-video system, so the echo data consists of real-values rather than complex (IQ) samples. The matched filter is a digitally sampled estimate of the Doppler-shifted return echo after downshifting to baseband. The filter was modified slightly in order to zero out the DC term in the echo samples. The range compressed analytic signal is obtained using fast Fourier transform techniques in the frequency domain to perform the convolution after base banding of the signal. The final number of complex-valued

range compressed samples for each pulse is half that of the original real-valued echo samples.

C. Azimuth Compression

Azimuth compression refers to signal processing of the along track pulse data to obtain fine along-track resolution. Here pulse data within a burst at a fixed range (referred to as a range bin) forms the input signal for the azimuth compression. We refer to this signal as a range line. Each range line is multiplied by a matched filter in the form of a linear FM chirp (representing the approximate phase history of a point on the ground) and transformed into the frequency domain. The final result is a complex value for each image pixel from which a magnitude and a phase for each range and Doppler frequency bin can be derived. The center frequency of the matched filter is chosen so that for each range, r , the Doppler bin with the maximal antenna response is centered within the azimuth bandwidth (equal to the azimuth sampling frequency, i.e. Pulse Repetition Frequency (PRF) of the transmitted signal). The center frequency of the matched filter is the Doppler centroid $f_c(r)$ of the returned echo. The chirp rate of the matched filter is the rate at which $f_c(r)$ changes with time. For Cassini SAR the rate of change in $f_c(r)$ is small enough (because of the short burst length) that it can be neglected. However, the variation of f_c with range is large enough over most of the swath to require its inclusion in azimuth compression.

D. Calibration

After range and azimuth compression a 2-D (range,Doppler) complex-valued SAR image for a burst is obtained. The size of the single-burst images varies along the swath. At closest approach, the usable region of the single-burst image corresponds to 40 km by 7 km on the surface. Near the end of the swath the usable region is 140 km by 30 km on the surface. The long (range) dimensions provided are for the 4 elongated beams. The narrow central beam is a factor of 4 smaller in range. The image needs to be calibrated to remove the systematic variations in the amplitude due to image geometry and antenna effects as indicated by the radar equation. From the radar equation the correction factor, X , is given by:

$$X = \sqrt{\frac{P_t \lambda^2 G_a^2 G_r G_c G_r A}{64 \pi^3 r^4}} \quad (1)$$

Here P_t is the transmit power, λ is the wavelength, G_a^2 is the two-way antenna gain, A is the area of the pixel on the ground G_l is the commanded attenuator gain, G_c is the gain through the range and azimuth compression stages, and G_r is the gain of the receiver in ADC sample number per Watts. The square root is used because the SAR pixels are corrected in the complex domain (amplitude and not power).

The antenna pattern was estimated as a function of azimuth and elevation using passive radiometer measurements of raster scans across the sun. These sun scans were performed while Cassini was passing by Jupiter so that the sun was much smaller than the antenna beam width. The receiver gain was estimated using receive only noise measurements. Using a system noise temperature T_s , derived from scatterometer and radiometer measurements, we compute receiver gain directly from the mean of the square of the sampled receiver noise data. This computation subsumes any biases in the transmit power and allows us to calibrate the Cassini SAR data in an absolute sense. For more details on Cassini Radar calibration see [2].

E. Usable Pixel Computation

The SNR and ambiguity levels vary within a burst image. Because all the data for each burst is processed, portions of the single burst SAR image can have poor ambiguity performance or calibration. These portions of the image need to be eliminated in order to avoid corrupting the final multi-looked image product. Competing with the desire to maximize image quality for each pixel is the desire to keep as much of the image as possible to maximize looks and avoid introducing gaps in the coverage. Two criteria are used to determine those pixels in the burst image that will be retained to generate the final multi-looked image. First, the pixel must have a sufficiently high signal-to-ambiguity ratio. Secondly, it must not be so far down on the antenna pattern that it cannot be accurately compensated for antenna gain (the further down on the antenna pattern the steeper the gain curve and the greater sensitivity to pointing errors). For nominal SAR flybys of Titan, it is sufficient to include all pixels that fall within the 10-dB two-way gain contour of the antenna pattern. There are exceptional flybys such as T8 mentioned below where points within the nominal gain contour have unacceptably poor signal-to-ambiguity ratio. For the exceptional flybys, we compute the signal-to-ambiguity ratio for each

pixel and exclude those below some threshold. Prior to obtaining data we chose a value of 14 dB. After evaluating the T8 data it was decided to use 5 dB. This omitted the more egregious ambiguous regions, but kept imagery on the edges of the swath that were deemed to be useful by the science team.

F. Multi-looking

After each burst has been processed into a single look SAR image in range and Doppler space, it is interpolated onto a surface grid, detected, and accumulated. This procedure is similar to what was done for the Magellan SAR to Venus (also a burst mode radar) [4] and is also similar to the way some back projection SAR processors operate. Although it is possible to accumulate bursts in a standard range/Doppler coordinate system, (e.g. SRTM burst mode processing) interpolating to a surface grid is more efficient and directly converts the data to the desired coordinate frame for science analysis. The rapidly changing Doppler centroid further complicates accumulating bursts in range/Doppler coordinates. It varies significantly from burst to burst because: 1) The antenna boresight moves along the surface at speeds of 1-3 km/s; 2) the required burst duty cycle of 7% results in large times between bursts; and 3) Cassini radar's five antenna beam cycle causes the interburst gaps to be multiplied by a factor of 5.

The surface grid employed is an oblique cylindrical map projection in which the nadir track is the equator of the projection and its prime meridian passes through the point of closest approach. Titan is assumed to be spherical with a radius of 2575 km. This map projection was chosen because it yields approximately equal area latitude and longitude pixels for typical Cassini SAR swaths. The resolution of the grid (256 degrees/pixel) was chosen to insure that the SAR data was oversampled by at least a factor of 2. For each burst, a rectangular portion of the grid enclosing the usable Doppler and range pixels is chosen. Each pixel in the surface grid is checked to see if it falls within the usable portion of the SAR image. Pixel values in the ground image are obtained using a 2-D sinc interpolation for those pixels that have sufficient sample support around the corresponding point in the range/Doppler image. The detected values for all pixels from all bursts that contribute to a pixel on the surface grid are averaged to obtain the final image value. As the number of pixels (looks) contributing to a given output pixel varies spatially, an additional

layer maintaining the number of looks is maintained. The final pixel value is a somewhat biased estimate of the backscatter (normalized radar cross section of the area covered by the pixel) due to the presence of thermal noise. For most of the swath the bias is small because the SNR is large, but for low backscatter areas (dark areas in most image representations) and for the extreme ends of the swath where the SNR is lowest, the bias due to thermal noise can be significant. For this reason, a noise subtraction technique has been developed to eliminate this bias as described in the next section.

G. Noise subtraction

Biases in pixels due to thermal noise are removed by creating a noise only image that can be subtracted from the original SAR image. The noise only image is computed by replacing each echo in the SAR data with zero mean Gaussian noise with a given variance σ_g^2 and processing the resultant noise-only bursts in an identical manner to that which was done for the real SAR data. This method insures that the noise value subtracted from each pixel has undergone identical processing as the backscatter value for that pixel.(including the somewhat complex and rapidly varying radiometric correction factor) To insure that the subtracted noise is itself not noisy a Monte Carlo approach where 10 noise images are averaged is used to reduce the random errors in the noise bias estimate. The sample variance σ_g^2 used to generate the noise only images can be estimated in two different ways: (1) theoretically for a known system noise temperature similar to what was done in estimating the receiver gain or (2) empirically from the portion of the echo data in which we know no signal is present. We preferred the first method because noise-only portions of the SAR are scarce by design.

III. DOPPLER AND RANGE CENTROID TRACKING

In order to avoid artifacts due to poor pointing or ephemeris knowledge a data-driven method for refining the range and Doppler centroids used in range and azimuth compression was employed. The algorithm employed is similar to other data driven algorithms that have used in previous SAR systems including Magellan, SIR-C and SRTM. The improvement in the radiometric image quality obtained using this method is illustrated

particularly well by the first Titan SAR flyby, TA. The TA flyby occurred on October 25, 2004.

For our first images the *a priori* spacecraft ephemeris and pointing predicts used to process the data were suboptimal. (Updated and improved ephemeris data were generated several weeks after the flyby.) As shown in Figure 1, the initial SAR images were particularly poor. Especially prominent were large alternating dark and bright streaks along the seams between the different antenna beams.




Figure 1. Portion of TA SAR swath processed using predicted ephemeris and attitude data. No range or Doppler centroid tracking performed.




Figure 2. The same SAR data as in Figure 1 but with range and Doppler Tracking applied.

Initial corrections (images generated within the first 24 hours of the flyby) employed a manual correction that removed a large majority of the artifacts. Since it was apparent the problem was likely to persist on subsequent Titan passes, an automated range and Doppler centroid tracking mechanism was developed. The technique involves performing SAR processing twice. First, the SAR processor is run with the nominal ephemeris and pointing using only the central antenna beam. The central beam is used because it is 4 times narrower in range than the other four beams. For each burst the 2-D radiometric correction factor (X from equation 1) image for the burst is correlated with the 2-D detected and uncalibrated (that is the radiometric correction is not applied during SAR processing) azimuth and range compressed burst data. Doppler and range centroid offsets are chosen for each burst by determining the offsets that yield the maximal correlation. This technique relies on the fact that for a uniform radar scene, if the calibration is done correctly the uncalibrated, detected SAR image would be proportional to the radiometric correction factor. Because the real radar scene is not uniform, the burst-by-burst Doppler and range centroid estimates are noisy. We regress two 12th order polynomials in time from the burst-by-burst estimates. These polynomials are then used to compute the Doppler and range centroid offsets to apply during range and azimuth compression during generation of the final SAR image. It is assumed that the same offsets can be applied to all five beams. This approximation was sufficient for removing the most noticeable artifacts in all the 4 Titan flybys to which it has been applied. Figure 2 illustrates the improvement of the SAR image in Figure 1 obtained using the automated technique. The algorithm breaks down if the offsets become large compared to the SAR processing window. The only case where this has happened so far was a region in the TA swath near periapsis processed with the predicted ephemeris and attitude. In this case, a 20 km range centroid offset was required.

As we have accumulated more Titan passes, the a priori estimates of the ephemeris and pointing have improved resulting in improvement in the initial SAR images. However, even when the reconstructed ephemeris and pointing are employed there are still noticeable artifacts. Compare Figures 3 and 4. Figure 3 is an image of the portion of the TA swath near closest approach processed using the best available ephemeris and attitude data. Figure 4 is the same image processed using the Doppler and range centroid

tracking algorithm. The tracking algorithm removes almost all of the residual artifacts including the dark lines parallel to the swath (range centroid errors) and the scalloping perpendicular to the swath (Doppler centroid errors). The algorithm only corrects calibration. Geo-location errors remain. The location errors cannot be removed until the exact cause of the range and Doppler centroid errors is determined.



Figure 3. Portion of TA SAR swath near closest approach processed with the best available ephemeris and attitude data. No range and Doppler tracking applied.



Figure 4. Same portion of TA SAR swath as in Figure 3 but with range and Doppler Tracking applied.

The magnitude of the range and Doppler centroid offsets can be used to estimate the scale of the geo-location problem. Typically, it is necessary to apply corrections on the order of 2 km in range and 500 Hz in Doppler. The corrections can vary by as much as 4 km in range and 1000 Hz in Doppler over the course of a SAR Titan pass. The root cause that requires these corrections is still under investigation. A variety of possible contributors have been postulated including ephemeris, pointing, and timing knowledge errors, errors in our knowledge of Titan's shape, and large scale topography. Similar behavior has been observed in Cassini radar altimetry data as well.

Errors in ephemeris would have to be on the order of 1 km or attitude errors around 1 mrad to produce the range centroid offsets we are observing. The Cassini Navigation and Attitude Control Teams tell us these numbers are at least an order of magnitude too large to be reasonable. We are currently focusing on a possible 0.2 to 0.3 second timing knowledge error. Such an error would explain the general shape of the offsets we are observing. Even if the putative timing error turns out to be real, actual topographic variation on Titan may be enough to require Doppler and range tracking. However, if we can eliminate all but topographic error, then the range and Doppler centroid offsets themselves may become scientifically interesting. With other sources of error eliminated they can be used to estimate a low resolution surface height profile along the SAR swath.

IV. SAR PERFORMANCE AND COVERAGE TRADES

Currently, all the planned SAR passes for the entire Cassini mission are expected to image about 20% of the surface of Titan. For this reason the Cassini Radar Instrument Team is interested in increasing the area of coverage of each SAR pass as much as possible. An examination of trades that result in slight reduction in SAR performance while improving coverage has been conducted. A good trade is defined as one in which the region of Titan we would have covered without modifying our performance requirements is largely unaffected and the modified region of coverage includes a large amount of additional coverage of useful data albeit of poorer quality. Such a trade-off was performed for the Titan flyby T8, one of the latest flybys.

The nominal pointing design for a SAR swath is obtained by first selecting an incidence angle and PRF profile that: 1) insures no gaps within the SAR swath, 2) meets a minimal SNR, 3) meets a minimal range and azimuth ambiguity isolation criteria, and 4) maximizes the cross track width of the swath so long as the other three criteria are met. Taken together these four criteria serve to maximize the contiguous usable cross-track swath width. Once incidence angle and PRF profiles are determined, SAR pointing is chosen to minimize the amount of variation in the Doppler within a range bin (referred to as iso-doppler pointing). The primary reason for iso-doppler pointing is to obtain rectangular pixels and thus improve our spatial resolution. As the acute angle of the pixel parallelogram decreases the resolution along the long diagonal becomes markedly worse. In the extreme case in which range bins and iso-doppler lines are parallel, the resolution in the direction perpendicular to range devolves to the real aperture azimuth resolution, and the utility of azimuth compression is entirely lost. Cassini flies by Titan on a hyperbolic trajectory. For most of the swath it is impossible to obtain Doppler and range orthogonality because the plane perpendicular to the spacecraft velocity vector does not intersect Titan. For the remainder of the swath where orthogonality is possible, the constraints imposed to maximize contiguous usable cross swath width impose non-orthogonality everywhere except closest approach. Iso-doppler pointing chooses the attitude that maximizes the acute angle of the pixel parallelogram given these other constraints.

In T8, we employed a different spacecraft pointing design named *pushbroom* to increase surface area covered by the SAR. Due to Cassini's flyby geometry, for iso-doppler pointing the speed at which the antenna boresight moves along the ground (along-track speed) varies significantly with time from periapsis. The along-track speed is typically 3.3 km/s at closest approach, but slows down to about 1.7 km/s 10 minutes later. As a result, SAR images away from closest approach have a higher number of looks (> 10) and therefore less speckle noise and less variance due to thermal noise than those taken at or near periapsis. One way to image more surface area is, of course, to increase the along-track speed while reasonably maintaining the requirements of contiguous usable cross-track. The increase in along-track speed is achieved by gradually rotating the spacecraft about the axis approximately perpendicular to the direction of spacecraft velocity. Various rotation rates were examined and the surface coverage and radar performance between 10 and 20 min

from periapsis were assessed. It was found that a rotation rate of 0.6 deg/min after 10 min from periapsis could increase coverage by more than 12% for T8 without significant reduction in contiguous, usable cross track swath width for the nominal portion of the swath. The difference in pixel acute angles between the iso-doppler pointing and the pushbroom pointing is slight. The simulated signal to ambiguity ratio is greater than 10 dB for the vast majority of the pushbroom swath. The largest impact is SNR. The noise equivalent backscatter value is as high as -5 to -7 dB near the end of the nominal swath.

Figure 5 depicts the outbound quarter of the T8 SAR swath. The white line is the boundary between the nominal coverage and the extended coverage obtained through push-broom. The region above the line has high contrast and does not appear to be overly noisy compared to other parts of the image. Furthermore the fine grooves in the top right portion of the image demonstrate good resolution capability. The area below the line (12.8% of the entire T8 swath area) contains complex spatially varying backscatter information that is of interest to the Cassini Radar Science Team. The qualitative examination of the image tends to validate our decision to use pushbroom on this flyby. The reduction in SNR does not lead a significant decrease in the image quality due to the large amount of contrast in the image.

V. SUMMARY

The Cassini SAR processor has four unusual features needed to deal with novel aspects of the acquired Cassini data. First it utilizes a overlapping pulse-segmentation approach in order to separate echo energy that returns simultaneously from consecutive pulses. Second, it employs a signal to ambiguity ratio estimation algorithm in order to eliminate ambiguous portions of the compressed burst images. This technique removes egregious ambiguous regions, but has an advantage over simpler but more conservative techniques in that it preserves more of the unambiguous data. Third it employs a SAR noise subtraction technique to remove biases in the backscatter estimate for radar-dark and low SNR portions of the SAR swaths. Fourth it employs a Doppler and range centroid tracker for removing known artifacts from the SAR images. The results of employing the Doppler Tracker have been exhibited. It significantly improves image quality for cases in which ephemeris and attitude knowledge are known to be suboptimal and also offers improvement even when the best obtainable ephemeris and attitude data is used.

The reasons for the residual range and Doppler centroid errors are currently under investigation.

Finally we have employed a strategy for trading off SAR performance in order to improve SAR coverage. The imagery we have obtained from Titan pass T8 vindicates our decision to employ this trade-off.



Figure 5. Latest acquired portion of T8 swath. (10-18 minutes after periapsis) White line indicates where the iso-doppler SAR swath would have ended.

REFERENCES

- [1] C Elachi, M. D. Allison, *et al*, "RADAR: The Cassini Titan Radar Mapper," *Space Science Reviews*, vol. 115 pp 71-110, 2004.
- [2] C. Elachi, R. D. West, *et al*, "Cassini Radio Detection and Ranging (RADAR): Earth and Venus Observations," *J. Geophys. Res.*, vol. 106 no. A12 pp 30271-30279, Dec. 1, 2001.
- [3] C. Elachi, S. Wall, *et al*, "Cassini Radar Views the Surface of Titan," *Science*, vol 308, no. 5724, pp 970-974, May 2005.
- [4] R. S. Saunders, G. H. Pettengill, R. E. Arvidson, W. L. Sjogren, W. T. K. Johnson, and L. Pieri, "The Magellan Venus Radar Mapping Mission," *J. Geophys. Res.*, vol 95, p 8339, 1990.
- [5] I. G. Cumming and F. H. Wong, *Digital Processing of Synthetic Aperture Radar Data: Algorithms and Implementation*, Artech House Remote Sensing Library, 2005.