

A New Hybrid-Beam Data Acquisition Strategy to Support ScanSAR Radiometric Calibration

I. G. Cumming, *Member, IEEE* and D. C. Bast, *Associate Member, IEEE*

Dept. of Electrical and Computer Engineering
The University of British Columbia, Vancouver, BC, Canada V6T 1Z4.
tel: 604-822-4623 fax: 604-822-5949 ianc@ece.ubc.ca dannob@telus.net

April 17, 2003

ABSTRACT

Wide-swath SAR coverage is provided by RADARSAT using a multiple-beam scanning strategy called ScanSAR. Each beam covers a different range, and is allocated a fixed period of time in which to transmit and receive radar pulses. During SAR processing, the data from each beam must be “stitched” together to form a complete image of the scanned area. This data must be radiometrically calibrated to compensate for antenna beam patterns. However, incorrect measurements of the satellite roll angle cause errors in radiometric calibration, and can lead to visible artifacts in the image (e.g. banding).

A new ScanSAR data acquisition technique is proposed that improves roll angle estimation through the use of radar pulses, transmitted by one beam and received by another. The new data are called “hybrid beam data” and can be utilized with modified versions of existing roll estimation algorithms. This paper shows how the hybrid beam data are collected, accommodating PRF, range gate delay and other timing changes as beams are switched.

Key words: ScanSAR, banding, radiometric calibration, roll angle estimation.

1 Introduction

RADARSAT and ENVISAT provide wide-swath coverage by using a scanning strategy called ScanSAR that employs multiple beams covering different ranges. The antenna dwells on each beam for a fixed amount of time, transmitting and receiving a sequence of radar pulses known as a burst. At the end of the burst, the radar antenna switches to the next beam by adjusting its elevation gain pattern, and again dwells there for a fixed length of time. This process is repeated until all the beams have been scanned, at which point the antenna returns to the first beam and the scanning cycle repeats.

During the ScanSAR processing, these bursts must be “stitched” together in range and azimuth to form a contiguous image. However, the received data have a significant variation of energy *versus* range, due to the range spreading factor, incidence angle effects and the elevation beam patterns of the antenna. The variations are most noticeable for the imagery regions where adjacent beams overlap. The profile of gain *versus* range must be corrected using information on the beam patterns, known as Range Dependent Gain Corrections (RDGCs), in RADARSAT terminology. Each beam has a unique RDGC that is defined in terms of antenna gain *versus* elevation angle, and is provided to the user via the satellite payload file. However, the radar data are collected as a function of *range* rather than elevation angle. Consequently, errors in the measurement of the satellite roll angle lead to a range offset in the application of the RDGC, which causes a radiometric error in the corrected image. Other correction errors come from incorrect values in the RDGC table and dynamic changes in the beam pattern itself, but we concentrate on what is usually the dominant error – the roll angle effect.

In current processors, the satellite roll angle is estimated by fitting the received radar energy to stored beam patterns. Estimation errors are usually less than 0.1° , but errors as large as 0.3° have been observed in operational data. Roll errors of this magnitude result in significant gain variations in the processed image when the nominal radiometric corrections are applied (a roll error of 0.1° may lead to a radiometry error as large as 0.5 dB, and radiometric discontinuities as low as 0.3 dB can easily be seen [1]). Abrupt gain changes at the ScanSAR beam edges cause the most noticeable error, as they create a banding effect. This effect is illustrated in Figure 1, where the beam stitching points can be clearly seen. A similar effect occurs in azimuth, if the Doppler centroid is not estimated accurately, and the azimuth beam pattern is not well known [2].

In this paper, we present a new data acquisition strategy that provides more data to

the statistical roll angle estimator. The method requires an overlap in range when the ScanSAR beams are switched. The method is designed for the RADARSAT ScanSAR beam arrangement where the beams are switched in order. In principle, the method could also be applied to ENVISAT, but as the ENVISAT beams are switched in the pattern 1, 3, 5, 2, 4, there is a certain amount of azimuth decorrelation in the beam overlap area, and the method may not work as well. For this reason, we focus on RADARSAT parameters in this paper.

The paper is structured as follows. After the introduction in Section 1, Section 2 begins with a description of the current technique of acquiring ScanSAR data, including how the timing parameters are adjusted when the antenna beams are switched. It ends with a short summary of current roll angle estimation algorithms. A new beam switching method is then proposed in Section 3, which utilizes the invalid data area for each burst to acquire “hybrid beam” data. The feasibility of hybrid beam data collection is illustrated in Section 4 for cases where the pulse repetition frequency increases or decreases between bursts. Simulation experiments are described in Section 5, showing the roll angle estimation improvement to be gained from the hybrid data. Finally, the issues of implementation and the pros and cons of the proposed method are discussed in Section 6.

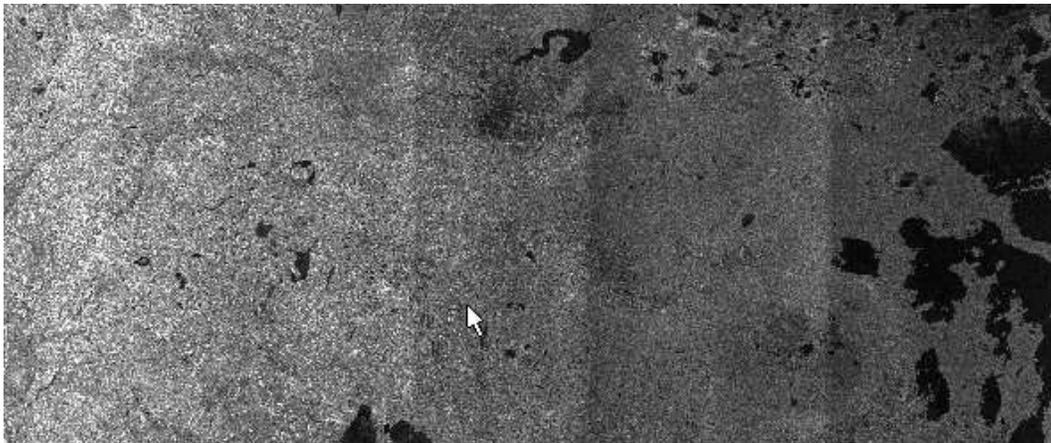


Figure 1: An example of the radiometric banding that sometimes occurs in RADARSAT-1 ScanSAR images. The image is made up of four beams with range increasing from left to right. Image ©CSA.

2 Current ScanSAR Operation

In this section, the current pattern of ScanSAR data acquisition and the present methods of roll angle estimation are briefly outlined.

2.1 Current ScanSAR Data Acquisition

In SAR data collection, each range line arises from the reception of an echo from a single radar pulse that is scattered off the Earth’s surface. The round-trip distance between the satellite and the Earth is approximately 2000 km, and the radar echo is received approximately 7 ms after transmission of each pulse. In RADARSAT, the pulses are transmitted at rates of 1100 to 1450 Hz, called the pulse repetition frequency (PRF), corresponding to an inter-pulse period or pulse repetition interval, $PRI = 1/PRF$, of approximately 0.7 to 0.9 ms (satellites with shorter antennas have a correspondingly higher PRF and lower PRI). Since the travel time of a specific pulse is much longer than the inter-pulse period, the corresponding echo is not received until a number of intervening pulses have been transmitted. The number of extra pulses “in the air,” m_i , between any transmit and receive event of the i_{th} beam is found from

$$m_i = \text{int}\left(\frac{2R}{c} \text{PRF}_i\right) \quad (1)$$

where R is the slant range to the target and c is the speed of light.

Figure 2 shows the patterns of two overlapping beams.¹ Beam 1 is given a starting range of 1000 km, and a PRF of 1230 Hz is chosen to place the echo in the center of the inter-pulse period for $m_1 = 8$. The sampling rate is 20 MHz. Both beams are drawn with range or time referenced to the start of Beam 1 echo reception. The beam shapes are arbitrarily selected so that they can be distinguished in the figures, however the echo durations are approximately the same as those used in practice.

Figure 3 shows the time lines of the transmitted pulses and the corresponding, received echoes. The pulses are numbered under each one, and the corresponding echoes are numbered above each echo. The vertical dashed line indicates the time that the beams are switched from Beam 1 to Beam 2, called the beam switchover point. The transmitter is turned off m_1 pulses before this point, so that the complete set of echoes for the current

¹ These patterns are arbitrarily chosen to provide visual distinction between Beams 1 and 2 in the timing examples of Figures 3, 4, 7, 8, 9 and 10. The actual RADARSAT-1 beam patterns are shown in Figure 6, and are used in the simulation experiments of Section 6.

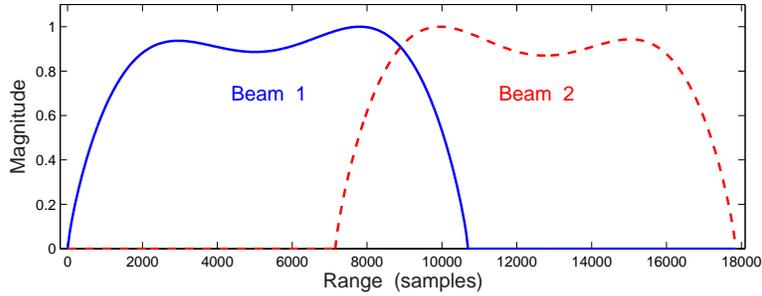


Figure 2: Gain patterns assumed for Beams 1 and 2 in the timing diagrams. The origin of the x-axis is referenced in range or time to the start of the Beam 1 echo reception.

burst are received before the switchover. The transmitter is turned on again right after the switchover, in this case with pulse number 12. No echoes are received until m_2 pulses later. Only part of a complete ScanSAR burst cycle is shown, including three echoes before the switchover, and three or four echoes after the Beam 2 echoes start coming in.

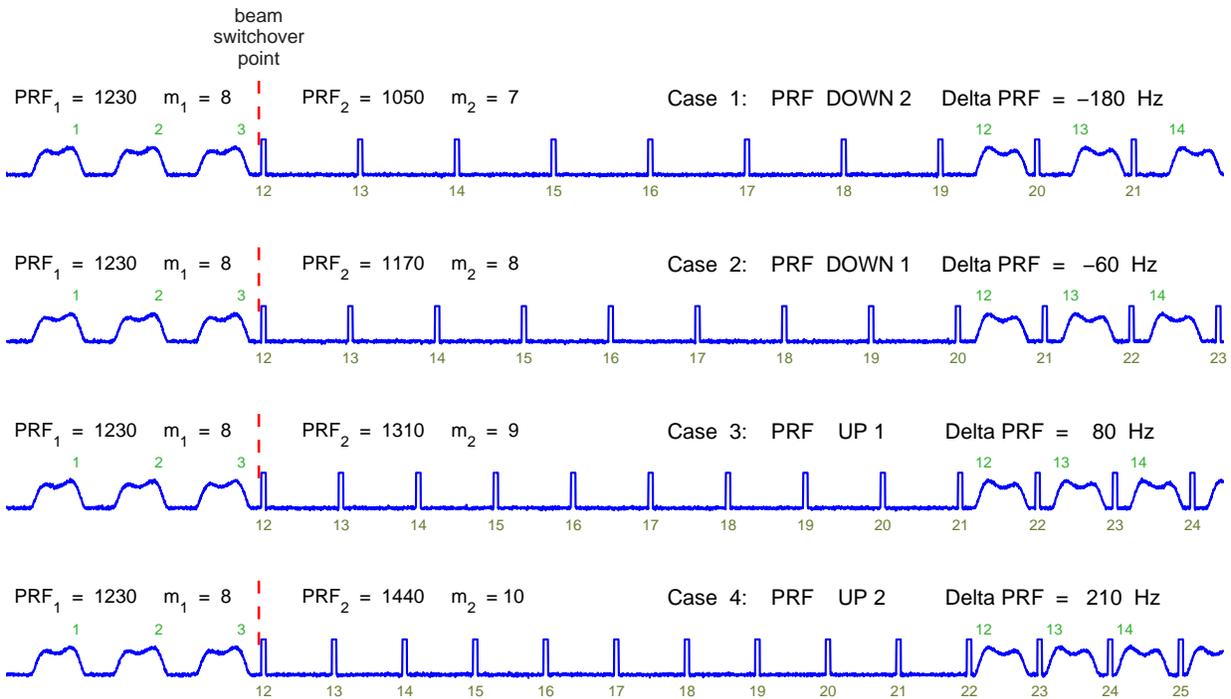


Figure 3: Illustrating the time line of the current RADARSAT-1 ScanSAR data reception.

When the beams are switched, the PRF must be changed so that the echoes of the new beam are received within the interior of the new inter-pulse period. In Figure 3 and subsequent timing diagrams, four possible PRFs are shown for Beam 2, 1050, 1170, 1310 and 1440 Hz. These four PRFs are selected as being representative of values that make $m_2 = 7, 8, 9$ and 10 , and allow the Beam 2 echoes to be received within a centrally-located

receive window for the given operating range. A choice of PRFs is needed, because the range has changed, and because the PRF should be selected to avoid the strong nadir return from below the satellite. However, the PRF cannot be changed too much because of ambiguity constraints [3] – Cases 1 & 4 may be outside the feasible range of PRFs, but are included here to illustrate specific timing features of the hybrid beams.

In Figure 3, it can be seen that as the PRF changes to different values, the number of “missing” echoes changes according to (1). Also, it is noted that Pulse 12 occurs at the same time in each of the four cases, as the PRI change has not had an effect yet and Echo 12 is received at the same time in each case, as its time delay is solely a function of the new range, not of the new PRI.

In summary, here are the key points illustrated in Figure 3:

1. The transmitter is turned off for the last m_1 pulses before the beam switchover.
2. All the remaining m_1 echoes are received prior to the switchover point.
3. The antenna elevation pattern is switched from Beam 1 to Beam 2 at the beam switchover point.
4. The transmitter begins transmitting pulses immediately after the switchover, with the timing parameters changed to the Beam 2 values.
5. After the switchover, the first m_2 receive periods do not contain valid data.

It is important to emphasize that, at the beam switchover point, all of the governing timing parameters (PRF, range gate delay T_{RGD} , and receive window length) change to accommodate the different range of the new beam.

Figure 4 shows the same data, but in a different format – the data in Figure 3 are cut into segments at the beginning of each pulse and arranged vertically. The data are arranged as they would be in a two-dimensional computer signal memory, with each new line beginning at the time that the transmit pulse is triggered.

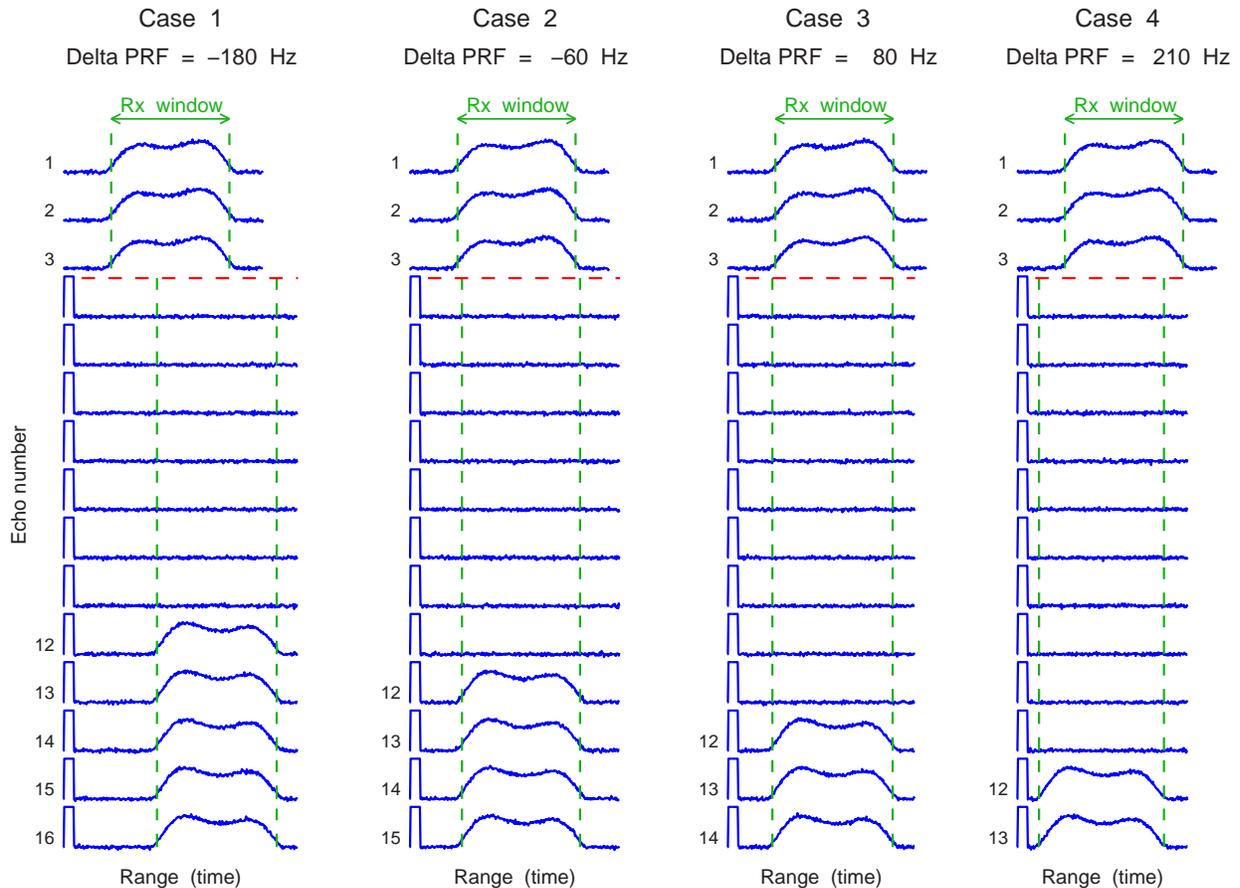


Figure 4: Illustrating the format of normal data reception in a hypothetical computer signal memory. The vertical dimension is the echo number, associated with the “azimuth” position of the satellite.

In the memory of Figure 4, each new range line is drawn below the previous one, to form a two-dimensional array of data. Each range line begins with the trigger of the transmit pulse, followed by a range gate delay, a receive window, and another small delay before the line ends. Time increases from left to right within each range line, and each range line begins PRI seconds after the preceding range line.

Each data segment shown in Figure 4 is referred to as a “range line” in SAR terminology, but only the data within the vertical dashed lines are downlinked and recorded, in practice. Note that the start time and the duration of the receive window usually change at the beam

switchover.

2.2 Current Roll Angle Estimation Algorithms

Several roll angle estimation algorithms exist based on Luscombe's idea [4] of using the overlap region between beams. In the overlap region, the same area on the ground is imaged separately by each beam. The collected data are then radiometrically corrected, assuming a certain roll angle and beam pattern. Radiometric differences among the corrected data from each beam are used to estimate the actual roll error and improve the correction. Bamler [5] refines this method by using a linear fit to the corrected data. Goulding [6] modifies this idea by computing the log ratio of the corrected data, and adjusting the roll angle until the ratio of the beam energies becomes equal to one. Jin [7] uses a different approach, involving the convolution of a kernel derived from the antenna pattern derivatives with the data. Dragosevic [8] expands on Jin's algorithm by working on all beam pairs together and adding parameters such as gain offsets for each beam.

While these algorithms provide adequate estimates much of the time, there are some scenes in which the roll angle estimate has sufficient error to cause an observable banding in RADARSAT-1 ScanSAR images. These algorithms rely on the use of the outer edges of each beam (the skirts), giving the algorithms two main disadvantages. First, the signal to noise ratio (SNR) of the received data gets quite low when the data comes from the edges of the beam, hence the estimates tend to be based on the noisiest of the received data. Second, the uncertainty in the beam gain is greatest at the beam edges, in both the actual (dynamic) beam pattern and the stored RDGC tables. Each of these effects leads to increased errors in the roll angle estimate, and substandard radiometric correction.

3 New Data Acquisition Method

The pattern of ScanSAR data collection can be further examined by drawing the regions in range/azimuth “space” from where the received data come. Figure 5 shows these regions for the current data collection pattern. It is clear that the first m_2 range lines in each burst contain only noise, since these receive events have no corresponding transmit event. This area is marked as “invalid data” in the figure. In the RADARSAT case, there are from 7 to 12 range lines in the invalid data area.

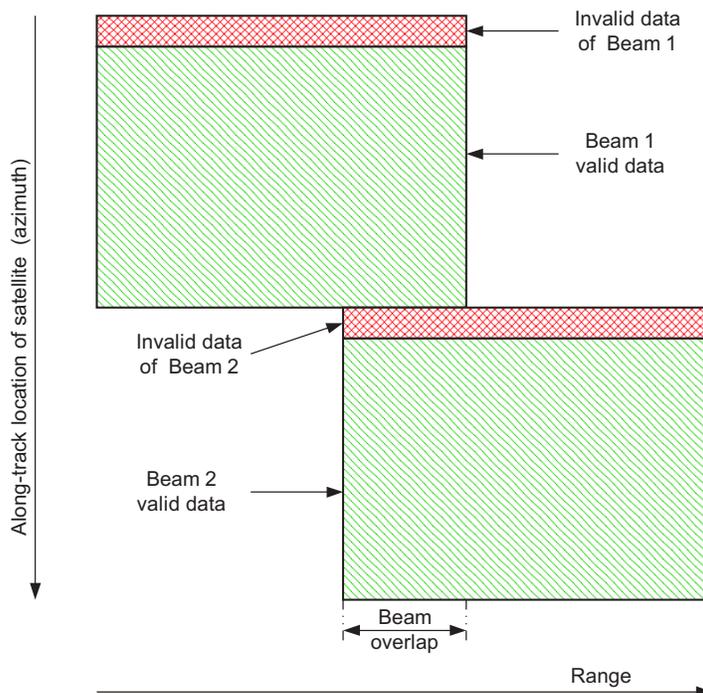


Figure 5: The range/azimuth data collection pattern of the received ScanSAR signal data.

We propose to use part of this unused area to provide extra information to the roll angle estimator. This can be achieved by maintaining pulse transmission in Beam 1 right up to the beam switchover point, giving m_1 extra pulses. The resulting m_1 echoes are received in the first m_2 receive periods of Beam 2. This new data is called “hybrid beam” data because it results from two different beam patterns: Beam 1 on transmit and Beam 2 on receive.

In RADARSAT, there are a number of different beams available that offer various elevation angles and elevation beam widths. These beams are designed to provide contiguous coverage in range when used in groups of 2, 3 or 4 beams for ScanSAR data acquisition. The beams used for ScanSAR include some of the STANDARD beams, S1 to S7, with a

100 km range swath width, and each of the WIDE beams, W1 to W3, with a 150 km swath. RADARSAT-1 offers 4 different ScanSAR modes. Two of these modes provide NARROW ScanSAR coverage (300 km range swath) and the other two provide WIDE ScanSAR coverage (500 km range swath). But in all of these modes, only five different *overlapping* beam pairs are used, and these are shown in Figure 6. The actual two-way beam patterns are plotted against elevation angle, and are shown to scale.

Hybrid beams formed by transmitting on one beam and receiving with the adjacent beam are also shown in Figure 6, such as the W1W2 beam in the top plot. Hybrid beams are available whenever the switch between beams creates an overlap in the range coverage. They are narrower than the regular beams because of the non-linear rolloff of the skirts (their elevation beamwidth is annotated in the figure).

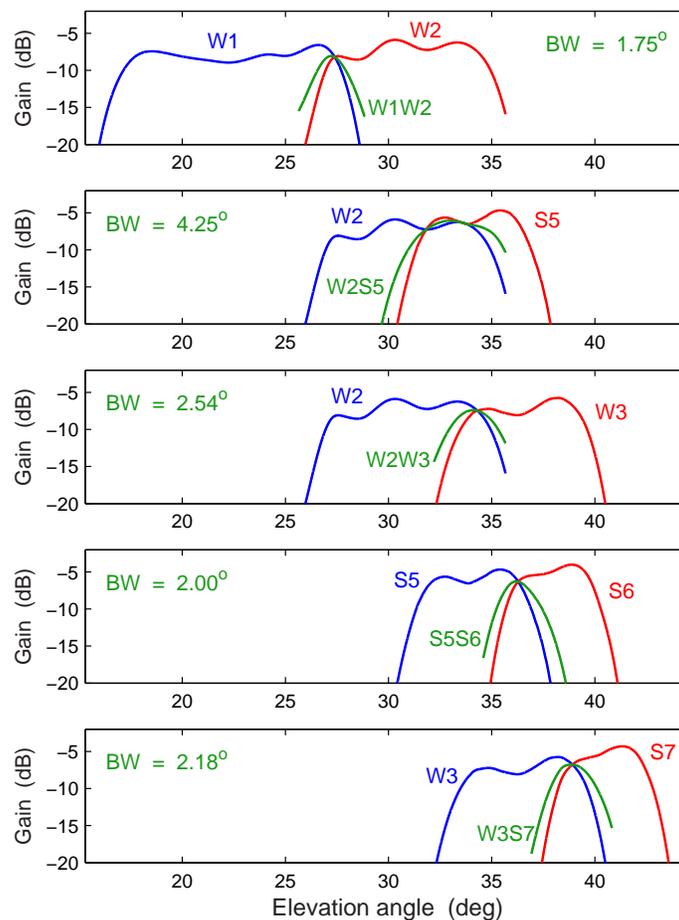


Figure 6: The five overlapping beam pair combinations used in RADARSAT-1 ScanSAR, and the corresponding hybrid beam patterns.

3.1 Benefits of Hybrid Beam Data

There are five advantages to incorporating hybrid beam data into the roll estimation algorithm:

1. The hybrid beams have narrower widths than the regular beams, improving the beam edge detection sensitivity.
2. The hybrid beams have less beam pattern uncertainty in the skirts of the beam.
3. The hybrid beam data have higher SNR in the low gain skirts than the regular beam data in the same range cell.
4. The hybrid data has less speckle decorrelation with the regular beam data, as they are closer in azimuth than the two regular beam data sets.
5. Hybrid data provides *extra* data for the statistical estimator.

The third point is a result of the combination of the stronger, better known portion of the near beam pattern, represented by its one-way gain G_x , with a weaker, less certain outer portion of the far beam pattern, G_y , up to the point of equal gain for the two beams (the interface point). After this interface point, G_y has a higher SNR and is better known than G_x . The hybrid pattern $G_x G_y$ is a combination of the two regular beams. In terms of both SNR and beam pattern certainty, it can be seen that $G_x^2 > G_x G_y > G_y^2$ on the near side of the interface and $G_y^2 > G_y G_x > G_x^2$ on the far side. In other words, the hybrid beam data has a higher SNR and greater beam pattern certainty than data acquired from the outer edge (skirt) of either the near or far beam.

While the addition of the extra data is the most obvious of the benefits, the other four points mean that the benefits obtained by the hybrid data are greater than just being proportional to the extra lines added.

3.2 How Hybrid Beam Data are Received

The hybrid beam data are generated by simply maintaining radar pulse transmission for the last m_1 pulses before the switchover, rather than stopping transmission. Since it takes a time of approximately $(m_1 + 0.5) \text{PRI}_1$ for the signal to travel to the Earth and back, the echoes of these Beam 1 pulses are received on Beam 2.

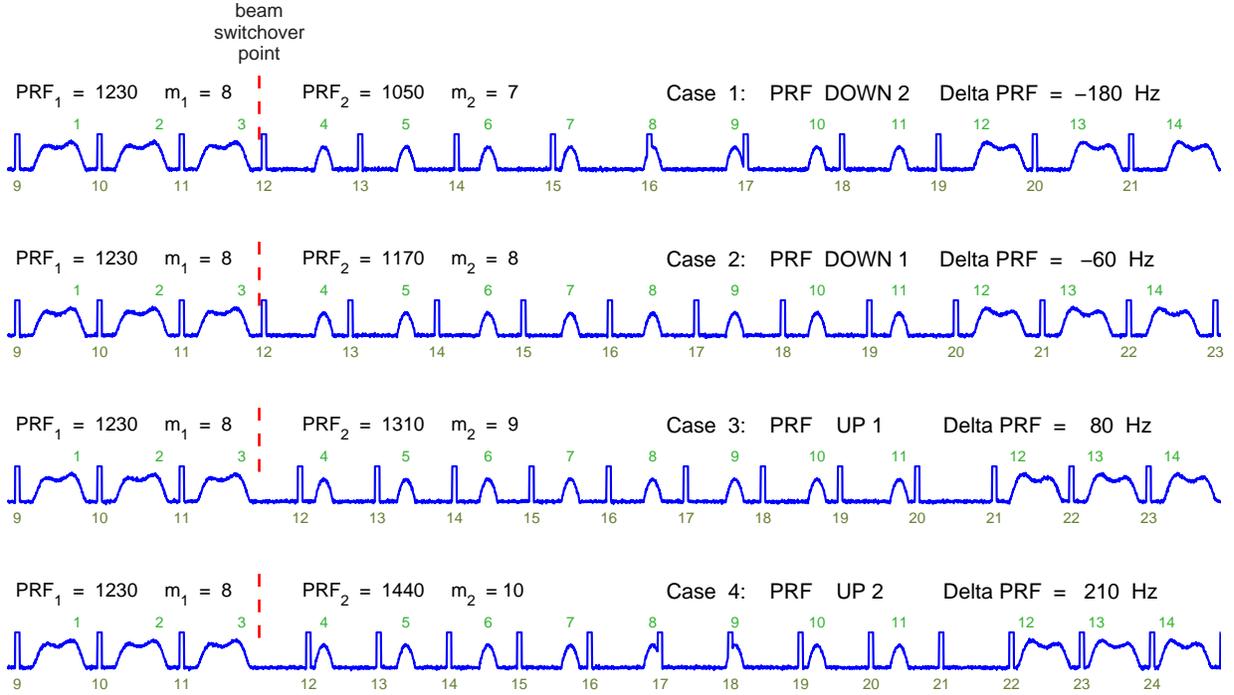


Figure 7: Illustrating the time line of the received hybrid beam echoes for four possible changes in PRF. Beam 2 illuminates a farther range than Beam 1.

The timing diagram of the proposed scheme is shown in Figure 7, for the same four PRF₂ cases as in Figures 3 and 4. Note that the m_1 echoes, 4 through 11, are the hybrid beam echoes, derived from the hybrid beam patterns of Figure 6.

Figure 7 illustrates five main features of the new pulse/echo timing:

1. The Beam 1 transmitter remains on until switchover, creating m_1 new pulses.
2. All previous Beam 1 echoes are received with Beam 1 prior to the switchover, as before.
3. The transmitter begins transmitting with Beam 2 timing parameters immediately after the switchover, as before.
4. The first m_1 echoes after the switchover point are received with Beam 2, but result from the hybrid beam patterns. These echoes are weaker and narrower than the main lobe of regular-beam echoes.
5. The trigger of the first pulse after the switchover (pulse 12) has been delayed by a small amount in Cases 3 and 4 to optimize the reception of the hybrid echoes.

It should be noted that the absolute receive time of the hybrid echoes is independent of the Beam 2 timing, which is why the trigger delay is sometimes needed. This is explained in Section 4.2.

4 Adjusting the Timing for the Hybrid Echoes

In this section, we discuss how timing parameters can be optimized for hybrid beam data reception. Figure 8 further illustrates the timing of receiving the m_1 extra echoes by redrawing Figure 7 in the format of the computer signal memory, as done in Figure 4. The m_1 echoes received after the switchover result from pulses transmitted using the timing parameters of Beam 1 (e.g. PRF_1), but are received using the timing parameters of Beam 2 (e.g. PRF_2). However, the Beam 2 timing is *not ideal* for reception of the hybrid echoes, as can be seen from Figure 8.

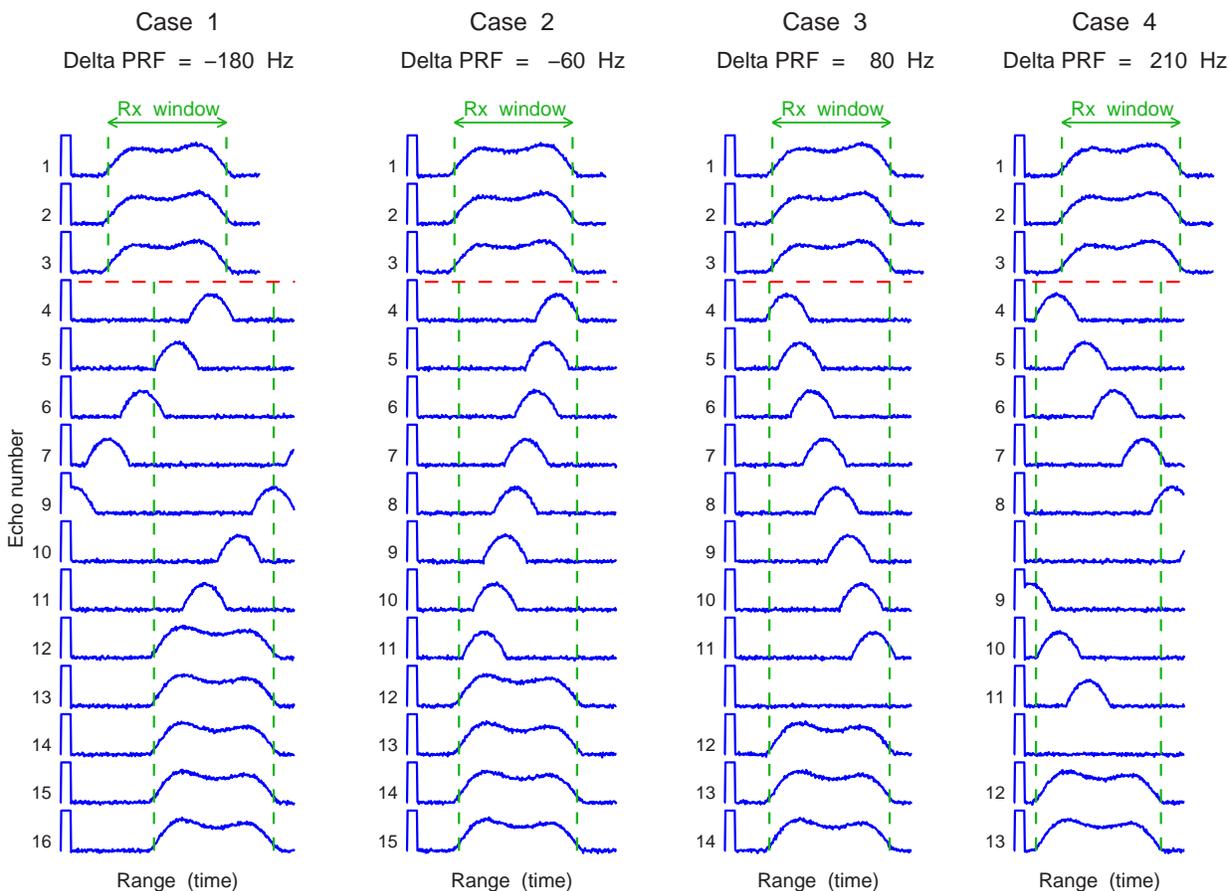


Figure 8: Illustrating the format of the hybrid beam data of Figure 7 when divided into range lines and stored in computer signal memory.

Since PRI_2 does not equal PRI_1 , the hybrid beam echoes *walk through* the receive window. The echoes walk left when the PRF decreases (Cases 1 & 2), and walk right when the PRF increases (Cases 3 & 4). Also, there may be a *shift* in the receive window after the beam switchover, as the range gate delay T_{RGD} may change. Therefore, it may be

advantageous to adjust the transmit timing in order to receive the maximum hybrid beam data (as outlined in Section 4.2).

This *change* in the start time of each hybrid beam echo in signal memory is given by

$$\Delta T_{\text{hyb}}(n) = \Delta T_{\text{RGD}} + (n - 1) \Delta \text{PRI} \quad (2)$$

where ΔT_{RGD} is the change in range gate delay and ΔPRI is the difference in the inter-pulse period between Beams 1 and 2. The “time” variable n is the range line number, counting from one after the switchover.

It is the *change* in PRI, and possibly in T_{RGD} , at every beam switchover that causes the shift and walk observed in Figure 8. The *absolute* arrival time of the hybrid echoes is fixed by the transmit timings of Beam 1 and *cannot be changed* by adjusting the regular Beam 2 timing parameters. Moreover, the PRI and receive window location of Beam 2 must be set according to the timing requirements of the valid Beam 2 data, rather than the requirements of the hybrid data.

Despite these constraints, most of the hybrid beam data can be received within the Beam 2 window by simply making a change to the trigger time of the first pulse transmitted after the switchover (see eqn. (3) below). This delay effectively moves the Beam 2 receive window such that it is positioned to open when the hybrid echoes return, but does not affect the relative timing of Beam 2 transmissions and receptions. The delay is visible in Cases 3 and 4 in Figure 7, but is not evident in Figure 8, as the origin of the horizontal axis in the latter figure is keyed to the pulse itself.

After this one-time delay, the pulse trigger and receive window timings remain stable at their normal parameters throughout the burst. In Figure 8, it can be seen that *all* hybrid echoes can be received in Cases 2 and 3 (the most common cases), while *most* of them are received in Cases 1 and 4. The nature of the timing change depends on whether the PRF change is positive or negative, as discussed in more detail in the next two sub-sections. The final sub-section discusses the 2-beam RADARSAT-1 ScanSAR case, where range *decreases* from Beam 2 to Beam 1.

4.1 Timing Adjustments when the PRF Decreases

It is instructive to first examine the timing of Case 2. It is the simplest case, and no pulse trigger delay (PTD) is needed to receive all the hybrid echoes.

In Case 2, examine the timing of the hybrid Echoes 4 through 11 in Figure 8, and compare them with the timing of Echoes 3 and 12. Echo 3 comes from Beam 1 and Echo 12 comes from Beam 2, while Echoes 4 through 11 come from the hybrid beam. In Case 2, the pulse trigger delay is zero. The hybrid Echo 4 arrives at approximately the same time as the *end* of Echo 3, assuming ΔT_{RGD} is small, because the hybrid beam has the same range as the right hand side (far range) of Beam 1, and the new PRI has not yet affected the timing. Echoes 5 to 11 then walk left because of the positive PRI change. By the time Echo 11 is received, the “walking” has almost finished, and hybrid Echo 11 has almost the same relative receive time as the beginning of Echo 12, the first echo received from a Beam 2 transmission. This is because the range of the hybrid beam corresponds to the left (near range) part of Beam 2. The walking actually ends with Echo 12, as Pulse 12 is the last transmit event governed by the Beam 1 timing. Thus, with no further adjustments, all of the hybrid echoes of Case 2 fall nicely within the Beam 2 receive window, as long as ΔT_{RGD} is small.

Case 1 is a little more complicated, because the larger PRF change causes the hybrid echoes to walk through the receive window faster, and about half of them move out of the receive window. There are still $m_1 = 8$ hybrid echoes, but they are compressed into $m_2 = 7$ PRI₂ intervals (Echoes 8 and 9 arrive in the same interval). Including a pulse trigger delay would not help in this situation. There is no way that they can all be received without a disruptive change in the receive window at each pulse. Such a change is not recommended because of its complexity, and because such a large change in PRF does not happen frequently.

4.2 Timing Adjustments when the PRF Increases

A positive PRF change is covered by Cases 3 and 4 of Figures 7 and 8. This PRF change *decreases* the inter-pulse period of Beam 2, and results in a *right walk* of the hybrid beam data. To modify the strategy recommended above for the “PRF decrease” case, the first hybrid echo (#4) is placed at the *left edge* of the receive window, in order to accommodate as many “walking” echoes as possible.

In order to place the first hybrid echo at the beginning of the Beam 2 receive window, the first pulse trigger of Beam 2 must be *delayed* by an appropriate amount. This once-off delay shifts Echoes 4 to 11 left in the receive window. The pulse trigger delay required to

achieve the appropriate shift is

$$T_{\text{PTD}} = \Delta T_{\text{RGD}} + T_{\text{Rx}_1} - T_{\text{Rx}_h} \quad (3)$$

where T_{Rx_1} is the receive window duration for Beam 1 and T_{Rx_h} is the duration of the hybrid echo. Since T_{Rx_1} is much larger than T_{Rx_h} , and ΔT_{RGD} is relatively small, then T_{PTD} is always positive, but less than PRI_2 . This means that the transmit pulse is always delayed (not advanced) and that this delay is less than the time required for a whole range line, PRI_2 . The one-time pulse trigger delay can be seen in Figure 7 by comparing the transmission time of Pulse 12 between Cases 1 and 2 (where there is no delay) and Cases 3 and 4, where there is a delay of about half of PRI_2 .

From Figure 8, it can be seen that hybrid Echo 4 has been placed correctly at the beginning of the receive window by the delay in Cases 3 and 4. Subsequent echoes then walk to the right. In the common Case 3, the walk is moderate, and all m_1 hybrid echoes are received within the window. However, in the less common Case 4, where the PRF increase is much larger, the walk is too fast for all the hybrid pulses to be received.² In fact, using a PTD in this case does not help much, and can be omitted.

² Note that the PRI directly after Echo 11 does not contain an echo in Cases 3 and 4, owing to the increase of PRF and the use of the PTD. Also, there is no echo in the PRI following Echo 8 in Case 4, because of the fast walk.

4.3 Timing Adjustments when the Range Decreases

The previous two sub-sections consider the most common RADARSAT-1 ScanSAR case where the range *increases* from Beam 1 to Beam 2. However, there is a RADARSAT ScanSAR mode consisting of just two beams, where hybrid beam data collection can also be achieved during the switch from the far range beam to the near range beam. This scenario is illustrated in Figure 9.

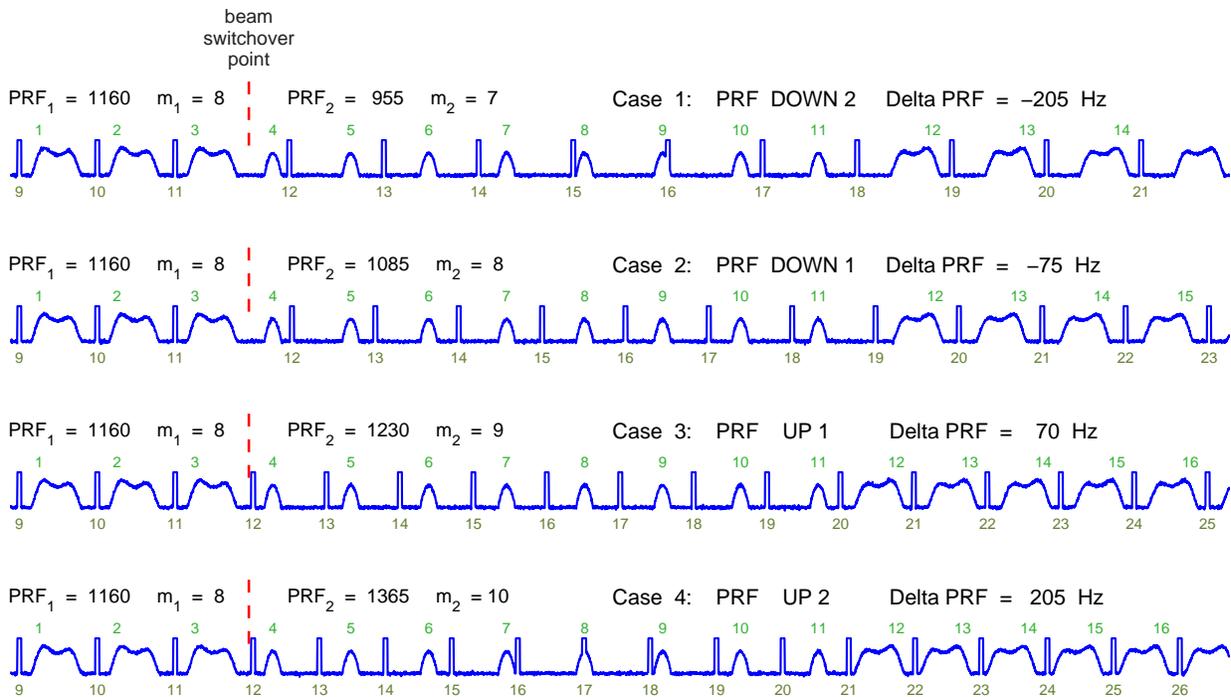


Figure 9: The time line of the received hybrid beam echoes when the range decreases.

The main difference in hybrid reception when the range decreases is that the first hybrid echo arrives near the beginning of the PRI, rather than the end (see Cases 3 and 4 in Figure 9). This timing is favorable for hybrid echo reception when the PRF increases (Cases 3 and 4 where the echoes walk right), and a PTD is generally not needed in these cases. When the PRF reduces at the switchover (Cases 1 and 2), a non-zero PTD is needed to move the first hybrid echo to the right side of the new receive window:

$$T_{PTD} = \Delta T_{RGD} - T_{Rx_2} + T_{Rx_1} + PRI_2 \quad (4)$$

As $T_{Rx_2} \gg T_{Rx_1}$, one inter-pulse period, PRI_2 , is added in (4) to get the desired positive pulse trigger delay. The hybrid echoes shift to the right side of the receive window by T_{PTD} , and then walk left. However, the first hybrid echo is lost, since the positive T_{PTD} shifts

the first receive window beyond its arrival time. Then there are only $m_1 - 1$ hybrid echoes available for reception.

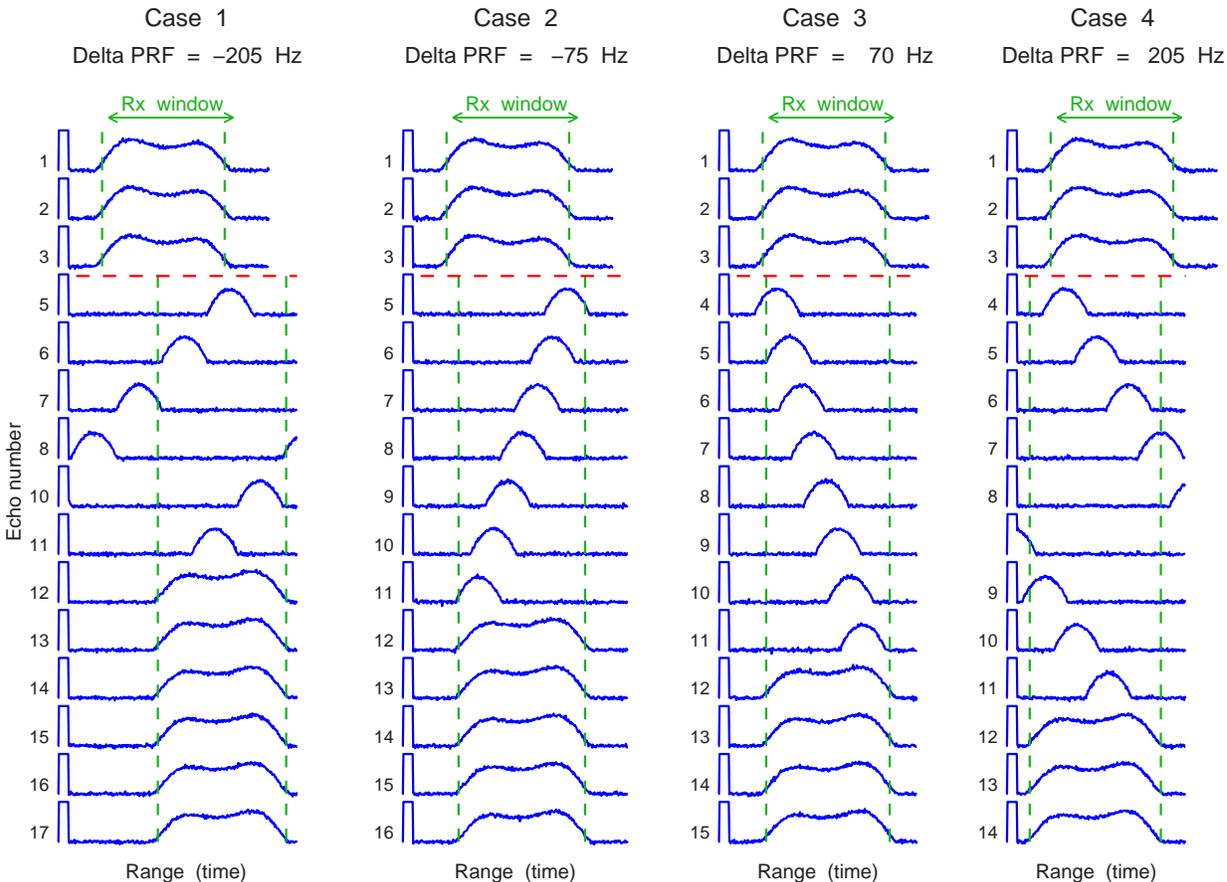


Figure 10: Illustrating the signal memory format of the hybrid beam data of Figure 9.

Figure 10 illustrates how the hybrid echoes fit within the receive window for the four PRF cases. In Case 3, with a 70 Hz PRF increase, all $m_1 = 8$ hybrid echoes can be received. In Case 2, with a 75 Hz PRF decrease, $m_1 - 1 = 7$ hybrid echoes can be received with the aid of the pulse trigger delay (4). In less common cases with a larger PRF change, about half of the hybrid echoes can be received.

5 Simulation Results

5.1 Simulating Hybrid Beam Data

Simulations are performed with current RADARSAT-1 ScanSAR parameters:

- to understand the operation of the current roll angle estimation algorithms;
- to modify a selected estimation algorithm to incorporate hybrid beam data;
- to quantify the improvement in roll angle estimation with the addition of hybrid beam data.

Raw SAR signal data are simulated, including the additional lines from hybrid beams, and placed in CEOS format data files. A variety of real SAR images are used as radiometric models, to ensure realism in the roll estimation algorithms. The simulated data incorporates the real elevation beam patterns of Figure 6 and an azimuth beam pattern based on a sinc function. Gaussian speckle noise is independently applied to both the in-phase and quadrature components of the raw data, and various levels of receiver noise is added to obtain specific values of image SNR. Random roll angle deviations between -0.2° and $+0.2^\circ$ are introduced to create an offset between the stored beam elevation patterns and the received data.

Beam pattern uncertainty is simulated by increasing or decreasing the gains on the edges using a linear function (zero to 100% of the applied error) from the inside to the outside of the overlap regions. The signal data are range compressed and applied to the roll angle estimation algorithms. More details are given in [9] and [10].

5.2 Roll Angle Estimation Algorithms

The various roll angle estimation algorithms available in the literature [4–8] are examined, and studied through simulation. The Goulding algorithm is chosen for use in this study for the following reasons:

- its implementation on regular data is simple;
- it is easy to modify to incorporate hybrid data;
- operational experience with real data is available.

The Goulding method works by finding the roll angle that equalizes the corrected energy between the data from the near beam and the far beam. First, the received energy in each

beam is corrected using the RDGC tables and an assumed roll angle. Then, the ratio of near beam to far beam energy is calculated at each range cell in the beam overlap region, using data from contiguous bursts, and averaged over range. Finally, the roll angle estimate is adjusted in the direction that drives the energy ratio closer to unity. After this iteration, the received energy is then radiometrically corrected using the new roll angle, and the process is repeated until the estimate converges. Utilizing an energy ratio specific to each range means variations in image intensity with range are no longer a factor. This algorithm is used on the regular data simulated in this paper.

Two estimation algorithms are developed to make use of the new hybrid data. In the first algorithm, we use only the hybrid data to get a feel for its properties. The energy of the hybrid data is plotted *versus* range, and a peak finding method is used to find the range at which the energy is highest. This algorithm works reasonably well, as the curve has a sharp, well-defined peak. This method is found to be fairly independent of SNR and beam uncertainty. However, its results are not as good as the Goulding algorithm used with conventional data, because it only has access to 10% of the received data, and suffers from bias due to changing radiometry with range.

The second estimation algorithm is designed to utilize both the hybrid and the regular data, to obtain the maximum benefit from *all* the available received data. The new method modifies Goulding’s algorithm to use *three* energy ratios: the near/far beam ratio, the near/hybrid ratio, and the far/hybrid ratio. Energies at the same range are again used to compute the ratios, thereby eliminating radiometric variations as factors in the algorithm. The difference between the three ratios and unity is minimized in a manner similar to Goulding’s algorithm, by adjusting the roll estimate, re-applying the RDGCs, and re-computing the energy ratios in an iterative procedure. The algorithm is found to have robust convergence properties.

5.3 Simulation Results

First, the requirements for roll angle accuracy are established by simulating the radiometric consequences of incorrect roll estimates. Various images, including pure noise, are radiometrically corrected using roll angles with different levels of error. First, each beam is corrected separately, resulting in a radiometric discontinuity where the beams are joined in range. Second, a linear blend is introduced over the beam overlap region to mask the discontinuity. The step discontinuity that arises in the first form of radiometric correction is the

most distracting form of error, and it is found that a roll error as small as 0.04° causes a radiometric discontinuity big enough to be just observable on a high-quality monitor. In the second case, it is found that a roll error up to 0.075° can be tolerated before producing an observable artifact. As blending is in common use, 0.075° is taken as the roll angle accuracy criterion for our study.

It is found that the main parameters affecting roll estimation for both the Goulding and modified Goulding method are first, the SNR of the received data, and second, the beam pattern uncertainty. A summary of the SNR results is given in Figure 11 as a function of average scene brightness, using the correct beam patterns. Image SNR is examined as a data parameter, to test if the algorithm works well over dark areas, such as water. Three bursts of received data are simulated using 66 different SAR scenes. A burst typically consists of 90 regular beam lines plus 10 hybrid beam lines. The receiver noise level is adjusted in each simulation run to model a certain scene brightness, as indicated along the horizontal axis of Figure 11. This scale assumes a noise equivalent sigma nought of -23 dB, corresponding to the lower end of the RADARSAT specifications.

Each symbol in Figure 11 represents the roll angle estimation error of the five beam pairs, averaged over the 66 different simulated scenes. The blue crosses give the estimation results when the hybrid data are not used, and the red circles indicate when the hybrid data are used in combination with the regular data. The solid lines are 3rd order polynomial fits through the data, to show the trends of the results.

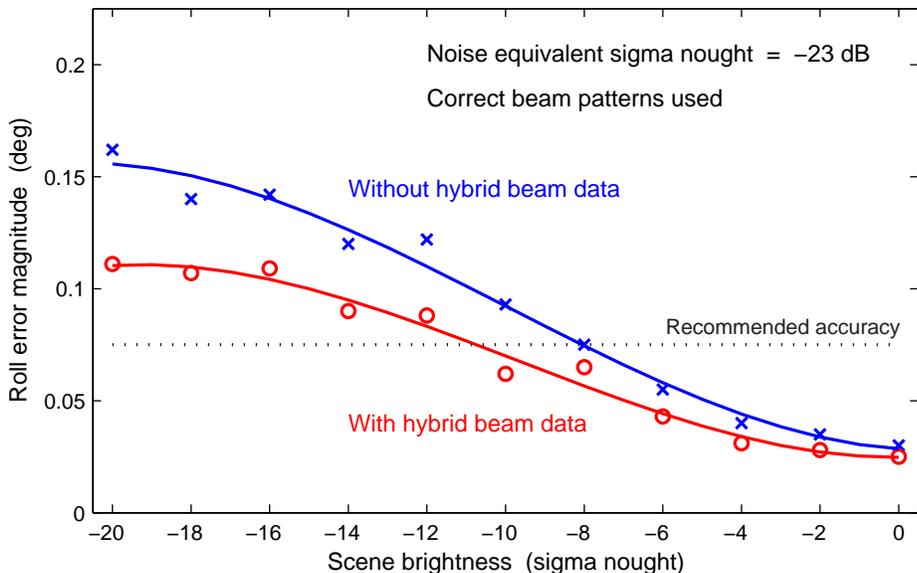


Figure 11: Simulation results with varying σ° but no beam pattern uncertainty.

Figure 11 shows that roll accuracy significantly improves when the hybrid data are utilized. For typical mean σ° values between -15 and -8 dB, the roll accuracy improves by approximately 0.03° when the hybrid data are used. However, for scene brightness below -10 dB, the roll angle error is above the recommended 0.075° limit. Further experiments are conducted using the worst scene, and it is found that approximately 12 bursts must be used in the estimator to lower the error to 0.075° for scene brightness down to -20 dB if the beam patterns are well known. As 12 to 23 bursts are received per second in RADARSAT-1, a one second estimation cycle is suitable.

When 0.2 dB uncertainty in the beam patterns is added to the simulation, both the conventional and hybrid beam estimation results degrade. However, the hybrid beam results are now 0.04° more accurate than the conventional beam results in the -15 to -8 dB scene brightness region, showing that the use of hybrid beam data reduces sensitivity to beam pattern errors. The increase in accuracy obtained from the hybrid beam data often amounts to the difference between noticeable or non-noticeable radiometric errors in the final imagery.

6 DISCUSSION

6.1 Implementation of Hybrid Beams on RADARSAT

Acquisition of hybrid beam ScanSAR data can be easily implemented within the current system design of RADARSAT-1. No fundamental hardware or software changes are required; only simple changes to the beam command table are needed [11]. For example, if a delay is required at the beginning of the burst, then a one-pulse “beam” can be introduced to the command table with a PRI equal to the required delay, and the transmit pulse suppressed. In a case where the trigger delay equals zero, only the number of transmitted pulses per beam changes from the present operation. Receiving, recording and down-linking the hybrid beam data occurs in the currently unused portion of the data burst.

In current RADARSAT-1 operation, it is possible for PRF changes of up to 200 Hz to occur, although typical changes are on the order of 50 Hz. In order to verify the feasibility of implementing the acquisition of the hybrid data, both positive and negative PRF changes of up to 200 Hz are examined on all five ScanSAR beam combinations. It is confirmed that the size of the PRF changes, as well as the amount of overlap between beams, are the two parameters determining the number of hybrid echoes that can be received within the adjusted receive window.

Figure 12 shows the number of valid hybrid echoes that can be received per burst for a given PRF change, in the case of $m_1 = 9$. The circles correspond to the narrow W1W2 hybrid beam; the crosses correspond to the wide W2S5 beam. The beam pairs not shown lie between the two cases illustrated, but are closer to the W1W2 case. It can be seen that a small PRF change allows all nine hybrid echoes to be received, while a large PRF change reduces this number. Most of the beam combinations can tolerate a PRF change of at least 50 Hz and still receive all nine hybrid echoes. The exception is the W2S5 hybrid beam, where the large beam overlap reduces the number of complete hybrid echoes that fall within the receive window. This reduction can be alleviated by widening the receive window just for the m_1 hybrid echoes, or by allowing a narrow window to track the hybrid echo walk, at the expense of increased control complexity. A longer term option is to design beams with less overlap.

Note that when the received data need to be coherently processed over bursts, the pulse trigger delay must be taken into account, as well as the PRF change. They must also be taken into account in the image registration equations.

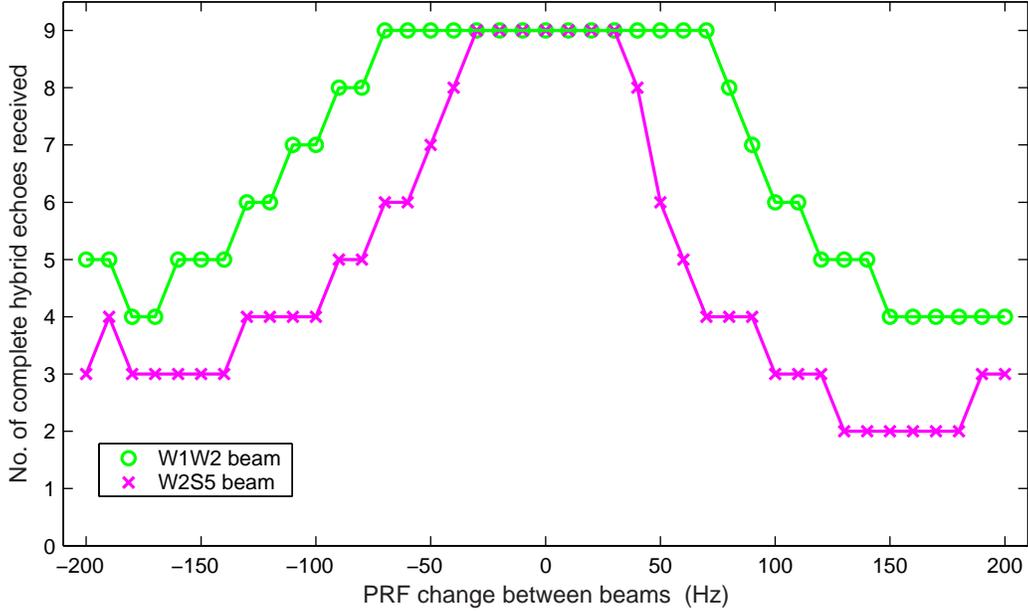


Figure 12: The number of hybrid echoes that can be received in RADARSAT as a function of the size of the PRF change.

6.2 Estimation Procedures

We have not considered estimator implementation in any detail, but can offer a few comments. An estimate of the roll angle is one of the main parameters needed for radiometric correction, but not the only one. Another parameter is the beam pattern itself. Updating the beam patterns is usually done off-line, but the interesting possibility remains of combining the estimation of the roll angle with the estimation of the beam gains, as proposed by Dragosevic [8]. Combining the estimators is especially useful if the beam gains vary from day to day, as beam pattern errors cause roll angle estimation errors as well as radiometric errors, so the estimation problems are inter-related. Modern SAR satellites tend to have active array antennas, increasing the likelihood that their gains change with time. Also, future satellites may have a transmit elevation pattern different from the receive pattern, complicating beam pattern measurement.

However, if the beam gains vary slowly, for example, over a few months, it is likely better to update the gains periodically using received data from a radiometrically-stable region such as the Amazon, rather than try to update the gains when the radiometry is highly variable. Despite the possibility of a combined estimator, we did not include the updating of the beam gains in our simulations, as our purpose is to show the increased accuracy of using the hybrid beam data, rather than to optimize the estimation algorithm.

If the beam patterns are to be updated within the roll estimation procedure, the hybrid data offer the advantage of being differently affected by offsets in the RDGC table, and thereby enhance the robustness of the estimator.

The roll angle estimates can be made during the SAR processing, or as a prior or post-processing operation. Radiometric corrections are made on every range line, but estimates are made less frequently, about one per second. More accuracy can be obtained by filtering the roll estimates using a model of the satellite attitude dynamics.

6.3 Advantages and Disadvantages of Using Hybrid Data

The invalid data area currently occupies about 10% of every burst. The ability to utilize this otherwise wasted down-link capacity to improve roll angle estimation is a clear advantage for radiometric calibration. The simulation results show that radiometric banding can be reduced from current levels, even below the visible threshold if the scene SNR is strong enough and the beam patterns are known with sufficient accuracy.

The main disadvantages of using hybrid beam data seem to be that:

- extra transmit power is required, and
- extra downlink capacity is required

These effects can be lessened by reducing the duty cycle of the hybrid beams to the minimum required for adequate estimation, and by shortening the receive window for the hybrid data.

In addition, a number of system and programming changes may have to be made to the data handling and estimation programs:

- extra commands have to be added to the beam command tables
- the downlink format specs may need to be changed
- the data handling in the ground receiver may have to be changed
- the roll estimation programs need to be changed

7 CONCLUSIONS

A new data acquisition method that utilizes the invalid area at the beginning of each ScanSAR burst is proposed. This method involves data from radar pulses transmitted by one beam and received by a subsequent, overlapping beam. These data are useful for improving the accuracy of ScanSAR roll angle estimation, in order to improve overall radiometric calibration and reduce visible banding artifacts.

The hybrid beam data are received by maintaining radar pulse transmission as antenna beams switch. A simple delay at the start of each burst aligns the data with the available receive window to maximize hybrid data reception. The proposed data acquisition is simulated and results confirm that an improved roll angle estimate can be calculated using hybrid beam data. The improvement is approximately 0.03° using 30 hybrid lines, and it is estimated that the use of 100 hybrid lines are sufficient to ensure adequate estimation accuracy and eliminate radiometric banding in almost all cases.

Acknowledgments

The authors would to thank Bob Hawkins of CCRS for reviewing the initial research and providing suggestions for improvement, and Tony Luscombe of MacDonald Dettwiler for reviewing the manuscript and providing feedback on the technical feasibility of the proposed method. Thanks also go to the Canadian Space Agency and MacDonald Dettwiler for providing image data and technical documents, and to MacDonald Dettwiler and NSERC for providing funding for the project. The authors would also like to thank the anonymous reviewers, whose suggestions led to improvements in the manuscript.

References

- [1] R. K. Hawkins, J. Wolfe, K. Murnaghan, and W. C. Jefferies, “Exploring the elevation beam overlap region in RADARSAT-1 ScanSAR,” in *Proc. CEOS WGCV-SAR Workshop*, (Tokyo), NASDA/ESA, April 2-5, 2001.
- [2] R. K. Hawkins and P. Vachon, “Modelling SAR scalloping in burst mode products from RADARSAT-1 and ENVISAT,” in *Proc. CEOS Workshop on SAR*, (London), ESA Publication SP-520, Sept. 2002.
- [3] R. W. Bayma and P. A. McInnes, “Aperture size and ambiguity constraints for a synthetic aperture radar,” in *Synthetic Aperture Radar* (J. J. Kovaly, ed.), pp. 118–123, Norwood, MA: Artech House, 1976.
- [4] A. P. Luscombe, “Using the overlap regions to improve ScanSAR calibration,” in *CEOS SAR Calibration Workshop*, vol. WPP-048, (Noordwijk, The Netherlands), pp. 341–346, ESA/ESTEC, Sept. 1993.
- [5] R. Bamler, “Roll angle estimation in SIR-C ScanSAR processing.” JPL Interoffice Memo 3340-94-030, 1994.
- [6] M. Goulding, “Roll angle estimation for ScanSAR processing in the CDPF.” MacDonald Dettwiler Technical Report, 1997.
- [7] M. Jin, “Optimal range and Doppler centroid estimation for a ScanSAR system,” *IEEE Trans. on Geoscience and Remote Sensing*, vol. 34, pp. 479–488, Mar. 1996.
- [8] M. V. Dragosevic, “Roll angle measurement and compensation strategy for RADARSAT ScanSAR,” in *CEOS SAR Workshop*, (Toulouse), ESA-CNES, October 26–29, 1999.
- [9] D. C. Bast and I. G. Cumming, “RADARSAT ScanSAR roll angle estimation,” in *Proc. Int. Geoscience and Remote Sensing Symp., IGARSS’02*, vol. 1, (Toronto), pp. 152–154, IEEE/CRSS, June 24–28, 2002.
- [10] D. C. Bast, “ScanSAR radiometric calibration based on roll angle estimation,” Master’s thesis, Dept. of Electrical and Computer Eng., The University of British Columbia, 2356 Main Mall, Vancouver, BC, Canada V6T 1Z4, May 2002.
- [11] Tony Luscombe, MacDonald Dettwiler, Private conversation, September 2002.