

Ultra-High-Resolution Near-Coastal Wind Retrieval for QuikSCAT

Michael P. Owen, Keith M. Stuart, and David G. Long
Brigham Young University, MERS Laboratory
459 CB, Provo, UT 84602

ABSTRACT

The SeaWinds on QuikSCAT scatterometer measures near surface ocean winds using radar backscatter values and a geophysical model function. QuikSCAT data is limited in coastal regions due to land contamination of the backscatter measurements. However, wind retrieval in near coastal areas can be successfully accomplished by estimating the amount of land contamination in the backscatter measurements and eliminating measurements which exceed a specified threshold. In order to accurately assess the amount of land contamination in a given measurement, a detailed knowledge of the antenna spatial response is required. The land contribution ratio is used as the contamination metric and is calculated using the spatial response of each QuikSCAT measurement. The land contamination threshold changes during wind retrieval as a function of wind conditions in the local area which allows retrieval closer to the coast as wind speeds increase. To ameliorate processing time the QuikSCAT spatial response is calculated and tabulated prior to wind retrieval. Subjective comparisons to estimated coastal winds show that determining a land contamination threshold using the percent land contribution metric provides accurate wind retrieval up to 25km closer to the coast than current methods. As a result wind speeds can be accurately retrieved as close as 2.5km from the coast.

Keywords: SeaWinds, QuikSCAT, wind retrieval, near-coastal

1. INTRODUCTION

The SeaWinds on QuikSCAT scatterometer offers valuable analysis of wind speeds over the ocean around the world. Operating in all weather conditions, QuikSCAT data is limited by contamination of measurements near land. Land contamination of QuikSCAT measurements is particularly troublesome within 30km of the coast because it makes accurate wind retrieval impossible. When land contamination is present, the retrieved wind speed can be up to 25m/s higher than the true wind speed. In order to overcome the limitations caused by land contamination, we calculate the spatial response for each QuikSCAT measurement and use this knowledge to determine the amount of contamination using a metric called the land contribution ratio (LCR).

Using the antenna spatial response in a tabulated form the LCR can be calculated for each measurement. Discarding measurements with LCR values above a pre-determined threshold allows uncontaminated wind retrieval using the remaining measurements. The LCR, when used as a land contamination metric, allows for the elimination of significantly land-contaminated measurements and enables retrieval of uncontaminated winds up to 25km closer than previously possible.

To show the effectiveness of the LCR as a land contamination metric, this paper briefly summarizes the QuikSCAT mission in Section 2. Section 3 explains the calculation, tabulation, and usage of the antenna spatial response function. Section 4 derives and explains the LCR, and Section 5 discusses results of high-resolution wind retrieval using the LCR.

2. QUIKSCAT OVERVIEW

Launched in June of 1999 aboard the dedicated spacecraft QuikBird, the QuikSCAT scatterometer is a dual pencil-beam rotating antenna that measures the normalized radar backscatter (σ^0) of the ocean and thereby infers wind speed and direction. QuikSCAT takes measurements at two incidence angles each at a different polarization: V-pol at 54° and H-pol at 46° . Each antenna beam has full azimuthal diversity [1].

Copyright 2007 Society of Photo-Optical Instrumentation Engineers
This paper was (will be) published in Proceedings of SPIE Vol. 6680,
26-30 August 2007, and is made available as an electronic reprint
(preprint) with permission of SPIE. One print or electronic copy may
be made for personal use only. Systematic or multiple reproduction,
distribution to multiple locations via electronic or other means,
duplication of any material in this paper for a fee or for commercial
purposes, or modification of the content of the paper are prohibited.

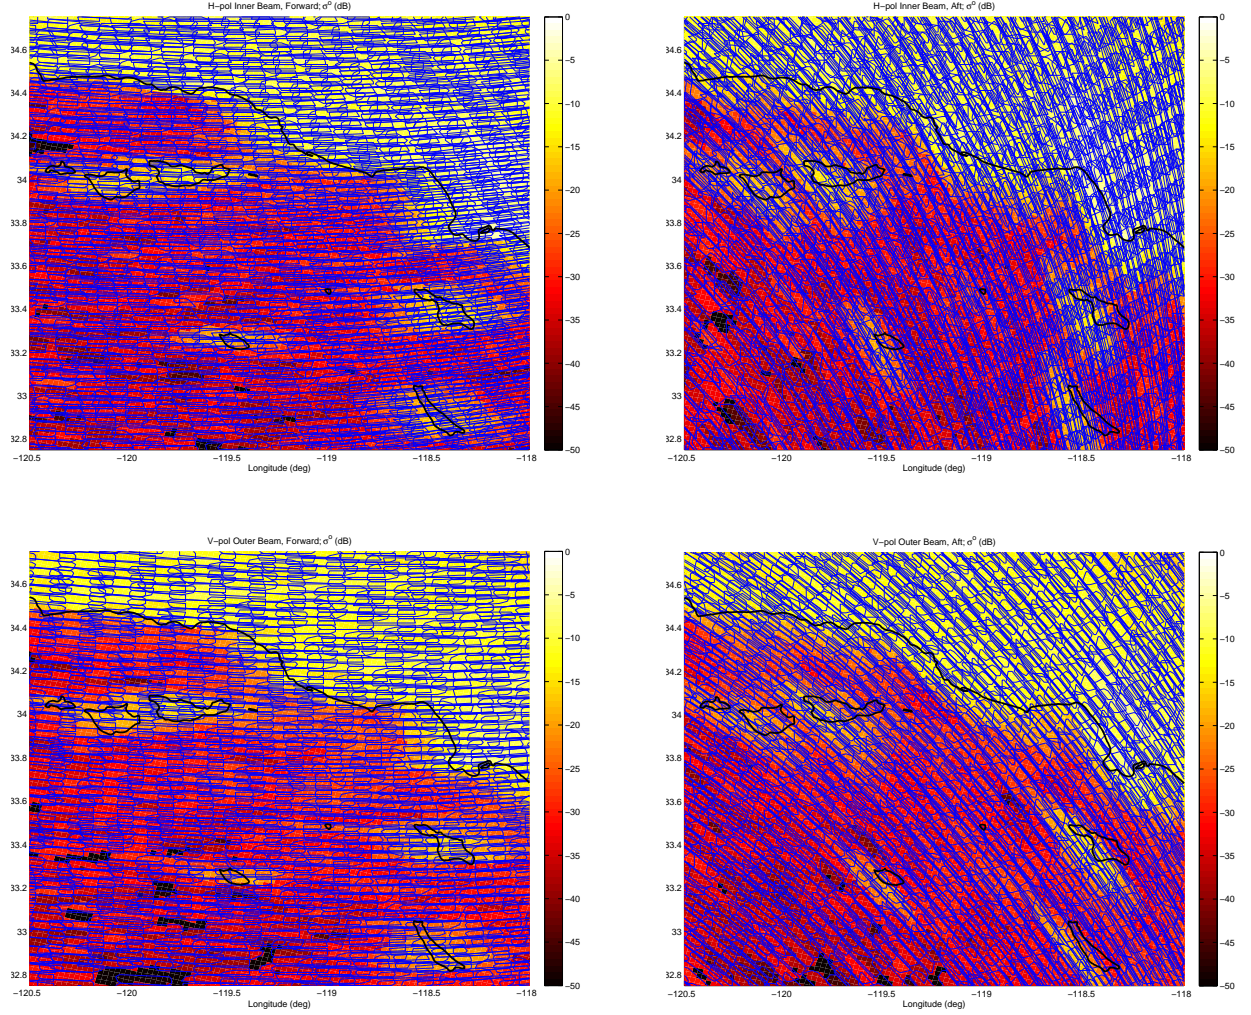


Figure 1. σ^o fields from a single QuikSCAT pass for forward and aft looking horizontally and vertically polarized slices together with 3dB contours and the coastline. Land contamination is apparent as σ^o values greater than -20dB. Note that contamination spreads outward from land principally parallel to the longer side of the slices.

Each pulse from the antenna is processed using an onboard range and Doppler processor and divided into 8 to 12 different regions each with separate σ^o values. These subdivided regions, called “slices”, are typically represented using the 3dB contour of their spatial response which has approximate dimensions of 6km by 25km [2].

Due to a rapid roll-off of the slice side-lobes, previous work has implemented a simple binary mask which only characterizes the 3dB contour. Slices can be used to produce ultra-high resolution wind and σ^o fields which have a 2.5km by 2.5km pixel resolution [3, 4]. Although lower resolution products exist, this paper addresses only high resolution products and consequently only the slice spatial response is discussed.

The true backscatter value (σ_{True}^o) for any measurement is the integral of the surface σ^o over the spatial response of the antenna

$$\sigma_{True}^o = \iint_{A_{slice}} \sigma^o(x, y) h(x, y) dx dy \quad (1)$$

where $h(x, y)$ is the spatial response of a slice, $\sigma^o(x, y)$ is the surface backscatter value and the bounds of integration are the bounds of the spatial response function.

σ_{True}^o can be written as the sum of the land and ocean backscatter values,

$$\sigma_{True}^o = \sigma_{Land\ Contribution}^o + \sigma_{Ocean\ Contribution}^o \quad (2)$$

or in integral form

$$\sigma_{True}^o = \iint_{A_{Land}} \sigma^o(x, y) h(x, y) dx dy + \iint_{A_{Ocean}} \sigma^o(x, y) h(x, y) dx dy \quad (3)$$

where A_{land} and A_{ocean} are the regions of the footprint consisting of land and ocean respectively.

Land contamination occurs for slices where the spatial response lies partially over land. As a result, $\sigma_{Ocean\ Contribution}^o$ which is the desired input for wind retrieval, is distorted or obscured by the larger value of $\sigma_{Land\ Contribution}^o$. While small levels of land contamination are tolerable in wind retrieval, many slices within 30km of the coast are intolerably biased. To illustrate how land contamination of σ^o is effected by proximity to land, the high resolution σ^o fields from one orbit are plotted on the coastline with the 3dB contours for each slice in Fig. 1. Land contamination is expressed as σ^o values greater than -20dB. Note that land contamination typically spreads outward from land parallel to the longer edge of the slices which is a consequence of the spatial response pattern for slices.

Each σ^o measurement is related to wind speed using a tabulated function known as the geophysical model function (GMF). The process of transforming overlapping σ^o measurements from both antenna polarizations to wind speeds through the GMF is known as wind retrieval. When land contaminated measurements are used in wind retrieval, the resulting wind speeds in coastal areas can be biased 5 to 25m/s higher than the true wind speed. Wind speeds so biased are too inaccurate to be used in a practical way and must therefore be discarded or ignored. Previous wind retrieval algorithms avoided land contamination by discarding all measurements within a fixed distance from the coast making wind retrieval close to the coast impossible.

After wind retrieval there exist multiple possible wind vectors, called “ambiguities”, for each wind vector cell (WVC). Ambiguity selection is the process by which one “correct” ambiguity is selected in each WVC. The ambiguity selection algorithm used in this paper is initialized by the 25km resolution wind product from the Jet Propulsion Laboratory (JPL) known as L2B. To select the high resolution ambiguities the high resolution ambiguity closest to the L2B wind vector is first chosen. The wind field is then passed through a median filter [5] to remove directional errors and to mitigate noise.

Prior to wind retrieval, the land contamination for each slice must be evaluated to determine which slices are uncontaminated and will therefore provide accurate wind retrieval results. To do so requires a detailed knowledge of the spatial response for each slice which must first be calculated. Due to the the complexity of the spatial response, the 3dB contour has insufficient information to determine the level of land contamination in any given slice. Hence an improved three-dimensional response function approach is developed.

3. QUIKSCAT SPATIAL RESPONSE

Many important factors such as orbit geometry, sampling resolution, and flight perturbations increase the complexity of computing the spatial response of SeaWinds. Due to the complexity involved, real-time processing of the response function for each slice is far too computationally expensive; therefore, the response function is pre-computed, tabulated, and stored prior to use. As a result, desired antenna response patterns can be accurately produced from the tables using common interpolation techniques. The trade-offs regarding this reconstruction process involve computation time, effective resolution, and the inherent process noise of the reconstruction process.

3.1 Calculation of the Spatial Response

Each received backscatter measurement σ_{True}^o can be interpreted as the integral of a ground target’s backscatter σ^o with the weighting function $h(x, y)$, as in Eq. 1. The weighting function can be factored into $h = h'/h_o$ where h' is the actual antenna spatial response with

$$h_o = \iint \frac{G^2(x, y) G_F(x, y)}{R^4(x, y)} dx dy \quad (4)$$




Figure 2. Contour plot of the first vertically polarized response function in the table for slice 4. The latitude and longitude are in degrees from the center of the egg. Contours are spaced 3dB apart.

which is the normalizing coefficient. Therefore, h is the normalized spatial response function, G is the antenna gain, G_F is the egg- or slice-dependent gain of the digital processor onboard SeaWinds, and R is the slant range [2]. Figure 2 illustrates h for a particular slice measurement.

Calculating h is computationally intensive due the numerous variables involved which include beam type; resolution mode; orbit parameter perturbations; spacecraft roll, pitch, and yaw; elevation grid spacing and resolution enhancement; azimuth grid spacing; and desired orbit time and azimuth angle [6]. To calculate the spatial response integral, each variable adds an extra dimension of complexity which results in an exponential growth in computation.

Due to changes in footprint size and orientation, computation of the spatial response h requires detailed knowledge of the antenna pattern G , the onboard processor G_F , the onboard Doppler pre-compensation algorithm, the observation geometry which determines the slant range R , as well as the x, y coordinate system used to define the spatial response.

For example, Fig. 3 illustrates nominal SeaWinds slice geometry from both fore and aft looking views at both polarizations. Note the changes in the 3dB footprint as a function of rotation angle: the slices behind the flight path are elongated and the slice boundaries change orientation as a function of azimuth angle.

3.2 Tabularization of the Spatial Response

Because computation of the spatial response involves multiple complex calculations, a lookup table-based approach becomes an effective method to satisfy the demands for both low computation time and high aperture accuracy.

The tables provide the spatial response $h(x, y)$ where x, y are the spatial coordinates corresponding to a differential field comprised of longitude u and latitude v perturbations. These differentials are added to the nominal longitude and latitude reported by JPL in the standard SeaWinds L1B datasets. The sum of the measurement center provided by JPL and the u, v perturbations derived from the tables result in the geographic location of the aperture footprint.

The tables consist of each set of the above calculations at 36 equidistant azimuth angles for 36 orbital locations. The orbit positions are unequally spaced to effectively characterize the orbit pattern of SeaWinds. Note that this table is organized for speed of computation and must be interpolated for use at a desired location. To further decrease computer requirements, a separate table is produced for each antenna beam.

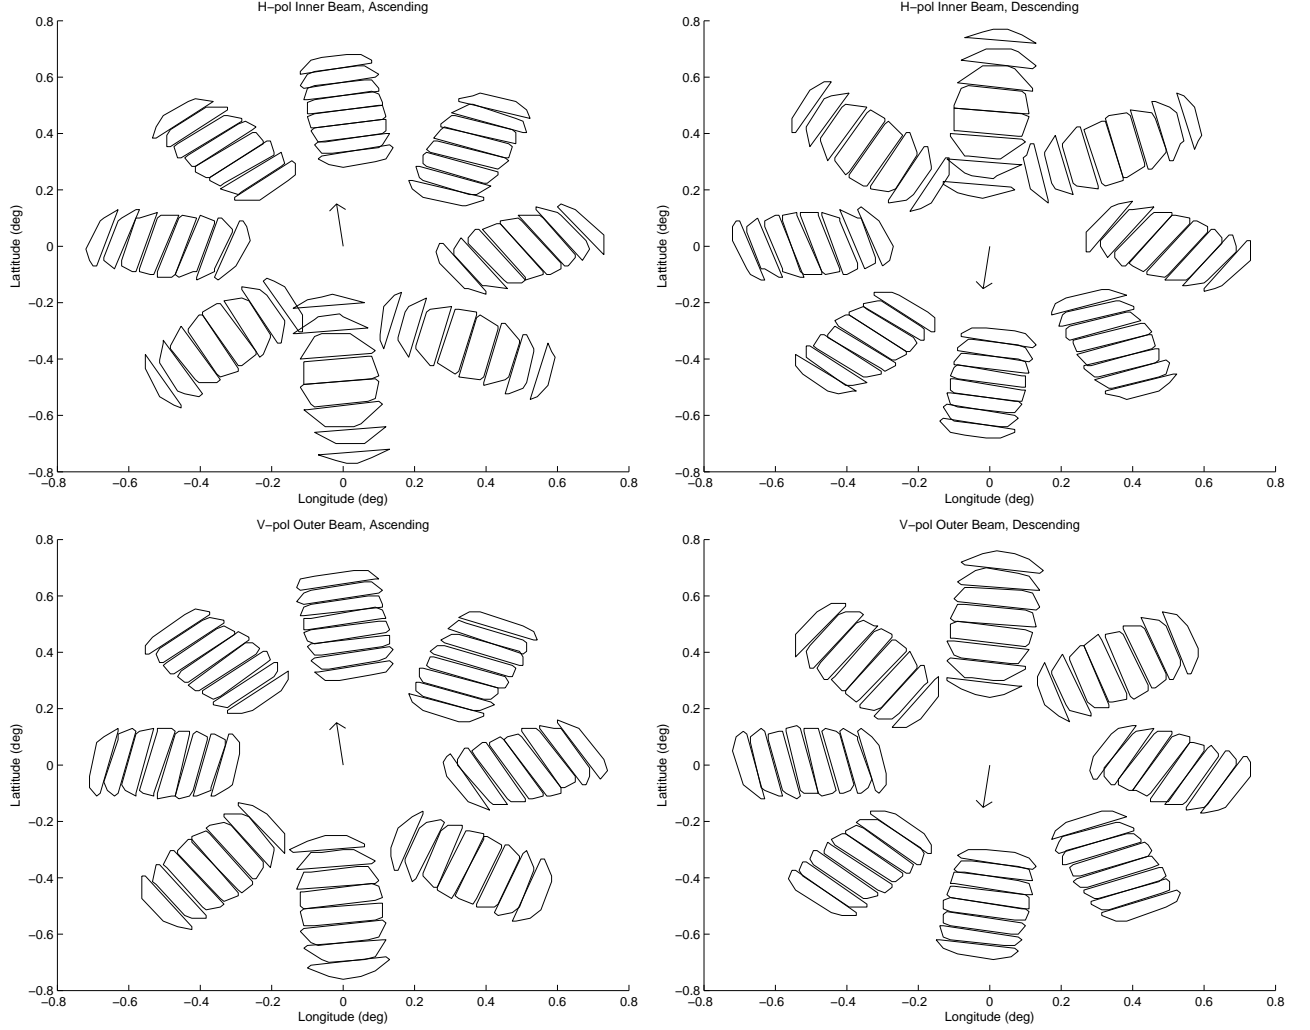


Figure 3. Illustration of the variations of 3dB slice geometries for the SeaWinds scatterometer. Both beam patterns at both ascending and descending equatorial crossings are depicted. Each set of slices correspond to different azimuth measurements. Arrows depict the line of flight [2].

In order to conserve file size, the tables are stored as raw 2 byte integers. The order of information is first m and n representing the size of the spatial response. Next is the latitude perturbation array with $m * n$ elements and similarly the longitude perturbation array. Following is the spatial response array for each of the 8 slices, i.e. $m * n * 8$ elements. Next is a set of 3 significant points spatially describing the orientation of the slice apertures due to range and Doppler filtering. Last are 8 centroid locations corresponding to each slice. The tabularized spatial response functions are available online at <http://scp.byu.edu> under “Software” then “SeaWinds Xshape.”

3.3 Usage of Tabulated Response

Usage of the tabulated spatial response function is complicated by the changes in slice orientation due to differences in range and Doppler processing for each azimuth angle. When the antenna is pointing in the along track direction this filtering process causes slices to be nearly perpendicular to the antenna azimuth direction. However, as the antenna turns, changes in the iso-range and iso-Doppler lines cause a skew in the slice direction relative to the antenna azimuth angle affecting the side-lobes of the spatial response. The side-lobes of the spatial response function have a large impact on land contamination and are thus vital in the accurate calculation of

land contamination for a slice; therefore, compensating for the skew in slice direction is a fundamental part of the interpolation process.

In order to accurately interpolate the spatial response between look-up table entries and compensate for the slice skew, two reference vectors per slice set are stored in the look-up table. One vector describes the main axis of slice orientation while the other vector is perpendicular and characterizes the main axis of each measurement orientation (the antenna azimuth direction). The two vectors are used to generate a transformation matrix which leads to smooth interpolation between the spatial response entries in the table.

To obtain the desired response function from the table, the entries are interpolated in both orbit time and in azimuth angle. This interpolation involves several steps.

The reference vectors for the two nearest azimuth angles from the two nearest orbit positions are extracted. Each set of vectors is mapped from latitude and longitude to a rectilinear grid using a tangent plane approximation described in [7] where

$$R_\phi = R_E \cos(\Delta\phi + \phi_0) \quad (5)$$

$$x = R_\phi \sin(\Delta\theta) \quad (6)$$

$$y = R_E \sin(\Delta\phi) + R_\phi (1 - \cos(\Delta\theta)) \sin(\phi_0). \quad (7)$$

x and y are kilometers, R_ϕ is the radius of the local latitude line, $\Delta\phi$ and $\Delta\theta$ are respective latitude and longitude perturbations, ϕ_0 and θ_0 are the latitude and longitude of the tangent location of the plane, and R_E is the radius of the earth.

With both sets of reference vectors on a rectilinear grid, the vectors for the upper azimuth angles are rotated counter-clockwise by the azimuth spacing used in the tables (10°) using the rotation matrix in order to be on the same axis as the lower azimuth vectors. The reference vectors are then linearly interpolated between the two corresponding azimuth angle entries to obtain vectors for the desired orbit time. The resulting vectors are then linearly interpolated again to obtain the reference vectors for the desired azimuth angle of the desired orbit time.

With the interpolated reference vector for the desired azimuth and orbit time, a transform matrix is calculated to compensate for the slice skew. The matrix

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \quad (8)$$

with coefficients a, b, c, d is calculated from the reference vectors for the nearest entry in the table in both orbit time and azimuth angle to the interpolated reference vectors. x, y are the components of the starting vector and x_t, y_t are the components of the interpolated vector.

Once the transform matrix coefficients are calculated, the look-up table's latitude and longitude perturbation fields are transformed to the desired azimuth angle and orbit time in the following 4 steps.

First the latitude and longitude arrays are mapped to a rectilinear grid using the tangent plane approximation in Eq. 5 - 7. Once the latitude and longitude perturbations are on a rectilinear grid, the x, y grid is transformed using Eq. 8 and then rotated around the egg center point to an x_r, y_r grid using

$$\begin{bmatrix} x_r \\ y_r \end{bmatrix} = \begin{bmatrix} \cos(\beta) & \sin(\beta) \\ -\sin(\beta) & \cos(\beta) \end{bmatrix} \begin{bmatrix} x_t \\ y_t \end{bmatrix} \quad (9)$$

where a, b, c, d are those calculated using the reference vectors and β is the difference in azimuth between the nearest table entry and the desired position.

After rotation, the x_r, y_r grid is mapped back to a latitude and longitude perturbation grid, $\Delta\phi_r$ and $\Delta\theta_r$ respectively, using the inverse tangent plane approximation where

$$\Delta\theta_r = \sin^{-1} \left(\frac{x_r}{R_\phi} \right) \quad (10)$$

$$\Delta\phi_r = \sin^{-1} \left(\frac{y_r - (1 - \cos(\Delta\theta_r)) \sin(\phi_0) R_\phi}{R_E} \right) \quad (11)$$




Figure 4. σ^o versus LCR for a pass over Catalina Island and the California coast for both horizontal and vertical polarizations. When the LCR is high, above -20dB, it indicates slices are partially over land. Slices completely over land, at 0dB LCR, correspond to the σ^o of land in the region, -11dB.

and variables are those used for Eq. 5 - 7.

This sequence of rotating the latitude and longitude perturbations saves computation time because the response function is not interpolated or resized. All necessary rescaling happens during the conversion of latitude and longitude to the rectilinear grid and during the rotation and transformation steps. Although the sampling of the response function is not identically spaced in latitude and longitude, the transformation is one to one in the local area. Each $\Delta\theta, \Delta\phi$ pair correspond to one response function value for each slice in the table; therefore, each $\Delta\phi_r, \Delta\theta_r$ pair correspond to the same response function values as the unrotated perturbation fields.

This efficient tabulation/interpolation scheme enables use of the full spatial response which was once too computationally intense. The usage of the spatial response in land contamination detection and removal using the LCR is one application of the full spatial response.

4. LAND CONTRIBUTION RATIO

An understanding of the spatial response for slices allows land contamination calculations to be made precisely. Although using the 3dB contour as a land contamination indicator provides insights into the nature of the land contamination problem (see Fig. 1), it does not provide a sufficient metric for contamination detection. The 3dB contour of a slice is the shape used successfully in resolution enhancement algorithms; however, it is insufficient for land contamination. Rather, using the complete slice spatial response, we determine the amount of land contamination for any slice and quantify the amount using a metric called the land contribution ratio (LCR). The LCR is derived from Eq. 3 as follows:

To calculate the LCR we assume that σ^o for land is constant. The approximate ratio of $\sigma_{Land\ Contribution}^o$ normalized by the σ^o of a land only measurement (σ_{Land}^o) can be written as

$$LCR = \frac{\sigma_{Land\ Contribution}^o}{\sigma_{Land}^o} = \iint_{A_{land}} h(x,y) dx dy \quad (12)$$

which we define as the land contribution ratio (LCR).

Rather than using a full continuous response function and integration, we simplify computation by using an approximation summation at a 1 km resolution. The simplified land contribution ratio is then estimated using

$$LCR = \sum_{x,y} L(x,y)h(x,y) \quad (13)$$

where x and y are kilometers away from the slice center, $L(x,y)$ is a land indicator function consisting of a 1 for land and a 0 for ocean. The bounds of summation over x and y from the center of the slice can vary depending on the desired accuracy.

To use the LCR for land contamination removal, the LCR is calculated for every slice measurement while creating the high-resolution σ^o fields using the AVE reconstruction algorithm [3]. Slices with a LCR greater than a pre-determined threshold are deemed “land contaminated” and are discarded.

To illustrate the accuracy with which the LCR can identify land contaminated σ^o values, Fig. 4 shows the LCR and σ^o for each slice in Fig. 1. LCR values near 0dB represent slices over land and LCR values of -80dB are slices entirely over the ocean. σ_{Land}^o in this region is around -11dB and σ_{Ocean}^o for this rev. is around -35dB. As shown in the figure, σ^o values below LCR values of around -32dB are not land-contaminated. This indicates a threshold of -32dB above which slices can be discarded.

The susceptibility of wind retrieval to land contamination varies as a function of the cross-track swath location, wind speed, and land reflectivity in the proximity of the measurement. Wind speed and direction errors are larger in the nadir region and at the edge of swath regions. Land contamination has effects on wind retrieval that vary accordingly. As wind speed increases the backscatter from the ocean increases making it less susceptible to land contamination. Thus, an increase in the reflectivity of land in the local region increases wind susceptibility to land contamination by increasing $\sigma_{Land\ Contribution}^o$.

Because the impact of land contamination on wind retrieval varies with the current wind conditions, LCR thresholds must be set individually for each slice based on local wind conditions. Compass simulations are performed for various wind vectors with varying land contamination to generate appropriate thresholds. LCR thresholds, determined by compass simulation, are tabularized in a look-up-table to be used during real-time processing.

Choosing the threshold during LCR processing requires a wind speed estimate, a land brightness estimate and the cross-track swath location of the measurement. The cross-track location of every slice is reported in the data, wind speed estimates are obtained from the JPL L2B wind product and land brightness estimates are obtained by evaluating slices which lie over land for each polarization.

The appropriate threshold is set according to the tabularized compass simulation results for each slice after which the LCR is calculated. If the LCR for any slice is greater than the threshold the slice is discarded. Wind retrieval is performed once land-contaminated measurements have been discarded.

5. HIGH RESOLUTION RESULTS

Prior to discarding land-contaminated measurements, land contamination in wind retrieval is apparent in wind speed fields as high wind speeds next to the coast. These wind speeds range from 5m/s to 25m/s higher than the true wind speed in the WVC. Figure 5 shows the high resolution wind speed for the Catalina Island region without LCR processing. The Los Angeles area (34°N, -118.5°E) is particularly contaminated because reflections from man-made objects are typically brighter than both land and ocean backscatter.

Previously, land contamination was avoided by creating a land mask extending 30km out from land. All measurements in the extended land mask were arbitrarily deemed contaminated and not used in wind retrieval. Figure 6 shows high-resolution wind speeds when all slices within 30km of the coast have been removed. Note that despite the land mask, which makes wind retrieval next to the coast impossible, some land contamination is still apparent as pixels with excessive wind speed along the edge of the mask.

When land contamination is identified and removed using the LCR, high-resolution wind speeds show several promising results (see Fig. 7) when compared to previous results. Most apparent is the region next to land where




Figure 5. High-resolution wind speed in m/s and wind direction vectors plotted without LCR processing.




Figure 6. High-resolution wind speed in m/s and wind direction vectors using a 30km land mask.




Figure 7. High-resolution wind speed in m/s and wind direction vectors after LCR processing.

no wind is retrieved which is much smaller than the land masked region of Fig. 6. This is a result of removing only land-contaminated slices during LCR processing rather than masking out all slices.

As a consequence of the missing slices, the distance to the nearest WVC with a wind estimate changes as a function of orbit geometry and wind conditions due to the local LCR thresholds. Wind speed, land brightness, measurement proximity to land, cross-track location and slice orientation all affect whether wind can be retrieved in a given WVC, and consequently, they also determine the nearest WVC to the coast with a wind estimate. The proximity to the coast of the nearest retrieved WVC can thus vary from as little as 2.5km (one WVC) to as much as 25km from the nearest coastline.

Comparing Fig. 7 with Figs. 5 and 6 distinctly shows the advantages of LCR processing. Further, wind vectors retrieved after LCR processing are validated by the compass simulation results and maintain approximately the same error levels found in the rest of QuikSCAT data.

6. CONCLUSION

In summary, the QuikSCAT slice spatial response function allows the LCR to be an accurate and computationally feasible measure of the land contribution to each slice measurement. Wind retrieval using the LCR as a metric for detection and removal of land contamination thus allows more accurate wind vector estimates closer to the coast

than previously possible. Although LCR processing requires significantly more computation, the advantages it provides to near coastal wind retrieval outweigh the additional computational cost.

REFERENCES

1. C. W. M. W. Spencer and D. G. Long, "Improved resolution backscatter measurements with the SeaWinds pencil-beam scatterometer," *IEEE Transactions on Geoscience and Remote Sensing* **38**, pp. 89–104, January 2000.
2. I. Ashcraft and D. Long, "The spatial response function of SeaWinds backscatter measurements," in *Proceedings of SPIE Vol. 5151 Earth Observing Systems VIII*, W. L. Barnes, ed., pp. 609–618, SPIE, Bellingham, WA, 2003.
3. D. S. Early and D. G. Long, "Image reconstruction and enhanced resolution imaging from irregular samples," *IEEE Transactions on Geoscience and Remote Sensing* **39**(2), pp. 291–302, 2001.
4. D. G. Long, J. B. Luke, and W. Plant, "Ultra high resolution wind retrieval from SeaWinds," *IGARSS*, pp. 1264–1266, 2003.
5. S. J. Shaffer, R. S. Dunbar, S. V. Hsiao, and D. G. Long, "A median-filter-based ambiguity removal algorithm for NSCAT," *IEEE Transactions on Geoscience and Remote Sensing* **29**, January 1991.
6. I. Ashcraft, "How to make and use the x-factor table," Tech. Rep. MERS 99-03, MERS Lab, Brigham Young University, 1999. <http://mers.byu.edu/docs/reports/MERS9903.pdf>.
7. I. Ashcraft, "Projecting an arbitrary latitude and longitude onto a tangent plane," Tech. Rep. MERS 99-04, MERS Lab, Brigham Young University, 1999. <http://mers.byu.edu/docs/reports/MERS9904.pdf>.