Direction Interval Retrieval With Thresholded Nudging: A Method for Improving the Accuracy of QuikSCAT Winds

Bryan W. Stiles, Brian D. Pollard, and R. Scott Dunbar

Abstract—The SeaWinds scatterometer was developed by NASA JPL, Pasadena, CA, to measure the speed and direction of ocean surface winds. It was then launched onboard the QuikSCAT spacecraft. The accuracy of the majority of the swath and the size of the swath are such that the SeaWinds on QuikSCAT Mission (QSCAT) meets its science requirements despite shortcomings at certain cross-track positions. Nonetheless, it is desirable to modify the baseline processing in order to improve the quality of the less accurate portions of the swath, in particular near the far swath and nadir. Two disparate problems have been identified for these regions. At far swath, ambiguity removal skill is degraded due to the absence of inner beam measurements, limited azimuth diversity and boundary effects. Near nadir, due to nonoptimal measurement geometry, (measurement azimuths approximately 180° apart) there is a marked decrease in directional accuracy even when ambiguity removal works correctly. Two algorithms have been developed: direction interval retrieval (DIR) to address the nadir performance issue and thresholded nudging (TN) to improve ambiguity removal at far swath. We illustrate the impact of the two techniques by exhibiting prelaunch simulation results and postlaunch statistical performance metrics with respect to ECMWF wind fields and buoy data.

Index Terms—Algorithms, backscatter, ocean winds, scatterometry.

I. INTRODUCTION

SEAWINDS on QuikSCAT (QSCAT) was designed by NASA JPL to measure ocean surface wind fields. End-to-end simulations performed to estimate the performance of QSCAT prior to its launch indicated that the directional accuracy of the wind vectors varies across the swath. Postlaunch comparisons between scatterometer data and analytical wind fields support this conclusion, as does visual inspection of the scatterometer wind fields. At far swath, ambiguity removal skill is degraded due to the absence of inner beam measurements, limited azimuth diversity and boundary effects. Near nadir,[1] due to nonoptimal measurement geometry (measurement azimuths approximately 180° apart), there is a marked decrease in directional accuracy even when ambiguity removal works correctly. Two algorithms have been developed to improve wind retrieval

[1]In this context, nadir is taken to mean along the ground track of the satellite. The antenna is never actually pointed perpendicular to the ground.

in these regions: direction interval retrieval (DIR) to address the nadir performance issue and thresholded nudging (TN) to improve ambiguity removal at far swath. These two algorithms are now performed operationally by the JPL Ground Processing System. The combined algorithm is referred to as direction interval retrieval with TN (DIRTH). The DIRTH wind vectors (as well as wind vectors produced without DIR) are included in the QuikSCAT wind data product. In this paper, we discuss the underlying theory behind the new algorithms (Section II) and empirically compare their results (Section III) to those from the baseline wind retrieval method. Wind vector accuracy relative to analytical wind fields (ECMWF) and buoys are presented for both the baseline and DIRTH cases.

A. Review

Before discussing the new algorithms in detail, we first review some of the general theory of wind scatterometers as well as some features peculiar to the QSCAT instrument. A scatterometer is a microwave radar that measures the normalized backscatter cross section σ_0. Geophysical model functions (GMF) have been developed empirically to map ocean wind speed and direction to σ_0 [2]–[6]. The theoretical basis of this relationship is the action of wind on small-scale (capillary) ocean surface waves, which in turn affect the ocean surface backscatter [7] (for a detailed overview of scatterometry, see [1]). Fig. 1 illustrates the QSCAT-1 model function, the most recent GMF used for QSCAT [5]. The y-axis in Fig. 1 is σ_0 expressed in dB. The x-axis is the angle between the radar beam and the wind direction. This figure represents the model function for V-polarized σ_0 with a 54° incidence angle. Cases with different incidence angles and polarization are numerically different but similar in form.

From Fig. 1, one can conclude that a single σ_0 measurement does not contain enough information to uniquely determine a wind vector. Multiple measurements from different look geometries are required. In the past, the need for multiple looks has been met by using multiple fan-beam antennas. This configuration has been used for all previous spaceborne scatterometers including the NASA Scatterometer (NSCAT) [1] and the European Research Satellite scatterometers (ERS-1 and ERS-2) [17]. How QSCAT's look geometry is designed to obtain these differing measurements is illustrated by Fig. 2.

Unlike previous scatterometers, QSCAT employs two conically scanning antenna beams. This configuration allows QSCAT to obtain a wider measurement swath and thus more frequent global coverage than fan-beam scatterometers. The

Fig. 1.   QSCAT1 geophysical model function. The x-axis represents the azimuth component of the angle between the wind direction and the radar beam. The y-axis is the normalized backscatter cross section in dB. The multiple curves represent different wind speeds. All the plots are done for V polarization and 54° incidence angle which correspond to the QSCAT outer beam.


Fig. 2.   QSCAT viewing geometry. (a) and (b) depict the sweet-spot and nadir-viewing geometries, respectively. Black squares depict 25 km × 25 km wind vector cells. The four spacecraft locations from which the cell is viewed are depicted by dots. Arrows from the dots to the squares indicate the viewing azimuths. Circles depict rotations of the inner and outer beams. The nadir track is represented by a dashed line and the outer edges of the swath by solid lines.

two beams differ in incidence angle (46°, inner beam; 54°, outer beam) and polarization (H pol, inner beam; V pol, outer beam). For most of the swath, every 25 km × 25 km cell on the ground is measured using four different look geometry configurations. Fore and aft measurements are obtained for each beam. The viewing geometry differs across the swath. For the outer portions of the swath, the viewing geometry is suboptimal: no inner beam measurements are available and as the extreme edge of the swath is approached, the azimuth diversity of the measurements approaches zero. At nadir, both beams are available, but the antenna azimuths are nearly 180°

apart between fore and aft looks. As we shall see in Section II, this is also a suboptimal viewing geometry. For a more detailed discussion of the QSCAT instrument, see [8].

B. Baseline Wind Retrieval

Before we discuss changes to the wind retrieval algorithm, we first describe the baseline technique. Wind retrieval was implemented for the NASA Scatterometer (NSCAT) and, initially, for QSCAT as a two step procedure: 1) a pointwise maximum likelihood estimator to calculate a set of likely wind vectors and 2) a median filter to select the best vector from the set. In the

Fig. 3. Wind vector probability maps for (a) a sweet spot wind vector cell (cross-track index 20) and (b) a nadir wind vector cell (cross-track index 38.) The x-axis is wind direction clockwise from north. The y-axis is wind speed in m/s. Darker regions are more probable. Regions of 80% probability are enclosed by thick solid lines. Thin solid lines depict the best speed ridges for the inner and outer forelooking measurements. Dashed lines depict the best speed ridges for the two sets of aft-looking measurements. Each peak of the objective function is denoted by an X.

first step, a search is performed in the space of all possible wind directions and speeds (a speed range of 0–50 m/s is presumed) to maximize a function that represents the likelihood that a certain trial wind direction and speed is the actual ocean surface wind speed and direction existent when the σ_0 measurements were obtained. The likelihood function is

$$f(u, \phi) = \sum_i \left(\frac{\sigma_{0i} - \sigma'_{0i}(u, \phi)}{\sigma_i} \right)^2 \qquad (1)$$

The value of the ith σ_0 measurement is σ_{0i}. The corresponding calculated σ_0 measurement $\sigma'_{0i}(u, \phi)$ is obtained from the trial wind speed u, direction ϕ and measurement geometry via the GMF. The expected standard deviation, σ_i, of the ith σ_0 measurement is computed from known characteristics of the instrument noise. The trial wind speeds and directions corresponding to maxima of the likelihood function are found by a heuristic search technique. The maximum likelihood estimator alone has been shown to be insufficient to choose a unique wind vector [10]. For a small set of measurement azimuth angles, multiple wind vectors may yield the same set of σ_0 values. Even if there are enough measurements from enough different azimuth angles to preclude this possibility, the addition of noise can still lead to multiple solutions of significant likelihood. For this reason, a discrete set of feasible solutions are obtained rather than a single solution. The solution set thus obtained is the set of local maxima of the likelihood function. For NSCAT, this solution set resulted in acceptable directional accuracy. The likelihood function dropped off quickly in the neighborhood of the local maxima, so that the chance of the true wind vector being far away from every vector in the solution set was small. As we shall see, the previous statement is not true for scanning pencil-beam scatterometers like QSCAT.

After the maximum likelihood step is performed, median filtering is used (with a 7×7 window size) to select among the available solutions (ambiguities) [13], [14]. This process is initialized by selecting the closest of the two most likely ambiguities to an analytical wind field. The initialization step is referred to as *nudging*. The *two* most likely ambiguities are used for historical reasons. For NSCAT, by far the most common set of ambiguities was two likely ambiguities approximately $180°$ apart and two other ambiguities with much lower likelihood values.

II. ALGORITHM

Two algorithms have been developed, direction interval retrieval (DIR) to address the nadir performance issue and TN to improve ambiguity removal at far swath.

DIR is a set theoretical estimation technique [9]. It is similar to the conventional (NSCAT) wind retrieval technique in that first a set of wind vectors are determined which are consistent with the data (solution set), then median filtering is used (spatial information incorporated) to select a solution vector from this set. DIR differs from the conventional method in that the solution set is not a finite set of vectors, but rather a set of disjoint one-dimensional (1-D) curves in the two-dimensional (2-D) space of wind speed and direction. The range of wind direction spanned by each of these curves is determined by a probabilistic analysis of the noise on the measurements and its effect on the directional discrimination information available (see Section II-A)

TN is a technique for optimizing the manner in which ambiguity removal is initialized. In the baseline wind retrieval algorithm, the closest of the two most likely ambiguities to a co-located numerical weather product (NWP) wind vector is used to initialize the median filter. With TN, the number of ambiguities available for initialization is not limited to two; instead it is determined by thresholding the likelihood values associated with the ambiguities. In this manner, fewer ambiguities are considered in regions of high instrument skill and the impact of the NWP field is lessened. On the other hand, in regions of lower instrument skill, more ambiguities are considered and the impact of the NWP field is heightened (see Section II-B).

A. Direction Interval Retrieval

For QSCAT, the rate at which the likelihood value drops off from the maxima varies with cross-track distance. For wind vector cells near nadir, there are large ranges of direction over which the likelihood value is relatively similar and it is inaccurate to represent the set of likely wind vectors by the likelihood maxima alone. The DIR method addresses this problem by calculating a solution set for each wind vector cell which includes a range of wind directions around each likelihood maxima. The extent of the ranges is determined independently for each wind vector cell according to the specific shape of the likelihood function for that cell. Fig. 3 depicts the difference in the shape of the likelihood function for a nadir QSCAT wind cell and a wind cell in the optimal viewing geometry (sweet spot) portion of the swath. The likelihood values are depicted by gray scale pixels. Darker pixels indicate regions of greater likelihood. The bold lines in the plots enclose regions of 80% probability. Given the
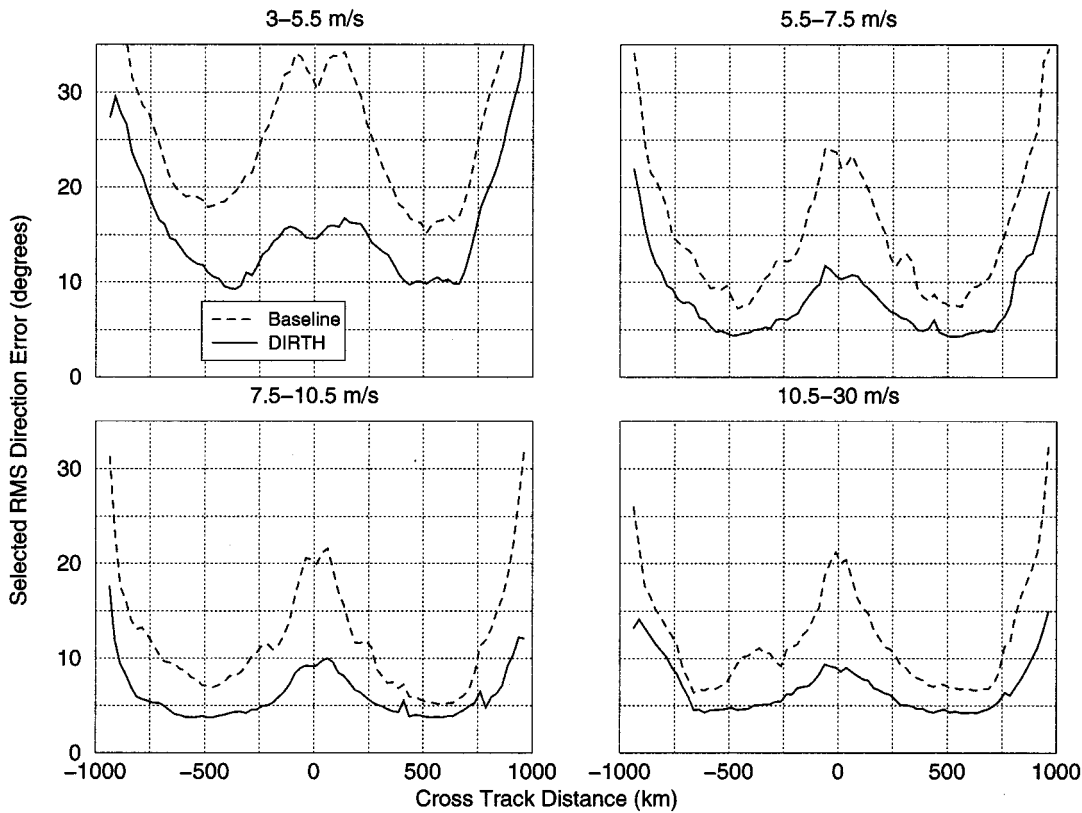
Fig. 4. Simulated RMS direction error versus cross-track distance. Baseline and DIRTH methods are compared for four ranges of wind speeds using data from 25 simulated orbits. Cross track distance is defined to be the cross-track component of the wind vector cell position in km, with the nadir track set to zero.

QSCAT noise model, the true wind vector falls within this region with 80% probability. Notice that the 80% regions in the NSCAT-like cell are much smaller than in the nadir wind cell, so that choosing the likelihood maxima alone as a solution set is a much better choice for the former than the latter. The thin lines represent the best speed ridges for the four sets of measurements (inner/fore, outer/fore, inner/aft, outer/aft). The best speed ridge is the curve $u = b(\phi)$ in the 2-D space of wind direction and wind speed determined by choosing the wind speed, u, which maximizes the likelihood function, $f(u, \phi)$ for each wind direction, ϕ. For the nadir cell case, the four best speed ridges calculated from the four sets of measurements all nearly intersect for a large range of directions, leading to decreased directional discrimination.

The DIR technique is a set theoretical estimation technique [9] that incorporates information from the σ_0 measurements and a model of the noise on those measurements in order to construct the solution set. Allowing the technique to consider all possible sets of wind vectors would be time prohibitive, so a simplifying assumption must be made regarding the types of sets to be considered. In the baseline technique, solution sets are four or fewer points on the best speed ridge corresponding to local likelihood maxima. In DIR, solution sets are generalized to four or fewer segments of the best speed ridge, with each segment including a local maxima. This choice of solution set is justified by the empirical observations that likelihood drops off sharply for speeds away from the best speed ridge and that whenever the wind direction is determined accurately, the wind speed is as well.

The endpoints of the segments are determined by estimating error bounds in a manner similar to techniques described in [11] and [12]. These techniques estimate probability distributions (and confidence intervals) for each measurement and then combine information by intersecting solution sets derived from confidence intervals on each measurement. The DIR technique instead estimates a joint probability distribution for all the measurements and then directly computes the solution set, yielding a more accurate (and more time-consuming) result. Since most of the information needed for the calculation is already available from the maximum likelihood [ML] estimator and the search space is limited to one dimension (by the best speed ridge assumption) computational efficiency is not a problem.

We assume the various noisy σ_0 measurements come from mutually independent Gaussian distributions. The means and variances of these distributions can be calculated in a manner consistent with the ML estimator algorithm. For a given wind vector and known measurement geometry, the geophysical model function is used to compute the conditional mean for each σ_0 measurement. The variances are computed using known characteristics of the instrument noise. The means and variances are then used to determine $P(\{\sigma_{0i}\}|u, \phi)$, the conditional probability density of obtaining the σ_0 measurements given the wind vector. The conditional probability is related to the likelihood function $f(u, \phi)$ by

$$P(\{\sigma_{0i}\}|u, \phi) = k \exp\left(\frac{f(u, \phi)}{2}\right) \qquad (2)$$

Fig. 5. Simulated RMS speed error versus cross-track distance. Baseline and DIRTH methods are compared for four ranges of wind speeds using data from 25 simulated orbits.

for some constant k. However, since the purpose of wind retrieval is to find the most likely wind vector for a given set of σ_0 values rather than vice versa, a more relevant probability density function is $P(u, \phi | \{\sigma_{0i}\})$, the probability density of wind vectors given an observed set of σ_0 values. This function when integrated over any region in wind vector space yields the probability that a wind vector within that region has occurred given the observed backscatter measurements. The two probability density functions are related by Bayes' Theorem,

$$P(u, \phi | \{\sigma_{0i}\}) = \frac{P(\{\sigma_{0i}\} | u, \phi) \, P(u, \phi)}{P(\{\sigma_{0i}\})} \qquad (3)$$

where $P(u, \phi)$ is the a priori probability density of wind vectors and $P(\{\sigma_{0i}\})$ is the a priori probability density of σ_0 observations. For a given set of measurements, $P(\{\sigma_{0i}\})$ is a constant.

In order to restrict the solution space to the best speed ridge we let $P(u, \phi) = 1/2\pi$ for (u, ϕ) on the best speed ridge and 0 everywhere else. The ML estimator as used in the baseline technique omits $P(u, \phi)$. This is mathematically equivalent to assuming that in the absence of measurements, all wind speeds (within the range of 0–50 m/s) and directions are equally likely. We make the same assumption with the additional constraint of limiting the nonzero probabilities to the best speed ridge. The extra constraint is a heuristic simplification which improves the computational efficiency of the algorithm. It is justified by the observation that probability density drops off rapidly with increasing distance from the best speed ridge.

By combining (2) and (3) and limiting consideration to wind vectors on the best speed ridge we get

$$P(\phi | \{\sigma_{0i}\}) = c \exp\left(\frac{f(u(\phi), \phi)}{2} \right)$$

for which the constant c is chosen to satisfy the probabilistic identity

$$\int_0^{2\pi} P(\phi | \{\sigma_{0i}\}) \, d\phi = 1.$$

Now that the estimation of the probability density function (pdf) has been obtained, the solution set segments are determined by thresholding the probability. Given a threshold T, a set of directional intervals around each of the local maxima is selected such that the sum of the widths of the intervals is minimized, and the integral of the pdf over the intervals is T.

The choice of the threshold T is an important consideration. A value that is too low, i.e., 0.1, results in a solution set that is too small to sufficiently represent the uncertainty in the measurements. In such a case the DIR technique will not go far enough in reducing the near nadir directional error (the baseline technique is identical to DIR with $T = 0$). A value which is too high, i.e., 0.95, overestimates the uncertainty in the measurements allowing the ambiguity removal step to oversmooth the data. In simulation, $T = 0.8$, the value used in producing the QSCAT wind data product, is found to be reasonable. Performance is found to be insensitive to small changes in T. The choice of T deserves further study because simulation studies and analysis wind field comparisons are insufficient to deter-

Fig. 6. RMS direction difference from ECMWF. Baseline and DIRTH methods are compared for four ranges of wind speeds. Difference is plotted versus cross track distance, using 4500 orbits of QSCAT data. Data flagged for rain is omitted from the statistics.


Fig. 7. RMS speed difference from ECMWF. Baseline and DIRTH methods are compared for four ranges of wind speeds. Difference is plotted versus cross track distance using 4500 orbits of QSCAT data. Data flagged for rain is omitted from the statistics.

mine its impact on mesoscale phenomena. Depending on how well mesoscale phenomena are preserved in the current product T may be decreased to reduce smoothing or increased to improve noise removal.

Once the solution set has been calculated for each wind vector. Ambiguity removal is performed to select a unique solution vector from each solution set. A two-step procedure is employed. First, one of the disjoint segments that composes each solution set is selected by performing ambiguity removal in the usual manner.[2] Ambiguity removal is performed on the local likelihood maximas and the segment which encloses the selected maxima is chosen. Next, a unique vector within the chosen segment is selected by iteratively choosing the vector which is closest in direction to the median vector of the surrounding 7×7 window.[3] Each wind vector cell is initialized by the maxima within the selected segment. Wind vectors are not updated until after each median filtering pass is complete. Passes continue until no wind vectors change by more than a threshold amount ($5°$) or a maximum number of passes (100) is exceeded. In practise, the maximum number of passes is seldom reached and typically the vast majority of the wind vectors are determined by the fourth pass.

B. Thresholded Nudging

The baseline nudging algorithm chooses an ambiguity to initialize the median filter. The algorithm only allows one of the two most likely ambiguities to be chosen. The rationale for that limit is based on NSCAT experience, where we assume that the scatterometer can choose the correct streamline and want the nudging field to select the proper ambiguity from that line. The other reason for limiting to two the number of ambiguities from which the nudging field can choose is to limit the influence of the nudging field and to rely on the scatterometer information as much as possible. If all ambiguities are allowed to be selected by the nudging field, the retrieved wind field would be very close to the nudging field, defeating the point of making the measurement.

The QSCAT situation is somewhat different from the NSCAT situation. In the outer swath, the scatterometer cannot always select the correct streamline. A significant percentage of the time (10–15% in simulation), the ambiguity closest to the truth is the third or fourth ranked ambiguity. Given that situation, one method that suggests itself is to use more ambiguities for nudging in the outer swath.

The likelihood function f can be converted into an estimate of probability (see previous section). We calculate *relative likelihood* $r(i)$, a quantity proportional to $P(\{\sigma_{0i}\}|u_i, \phi_i)$ for each ambiguity (u_i, ϕ_i).

$$r(i) = c \exp\left(\frac{f(u_i, \phi_i)}{2}\right) \tag{4}$$

with some constant c. The value of c is chosen for each wind vector cell so that the relative likelihood of the first

[2]with the exception that the median filter is initialized using TN. See the next section for more detail.

[3]The window size was chosen to correspond to the size used by the baseline median filtering algorithm. Additional window sizes deserve further study both for DIR and the standard algorithm.
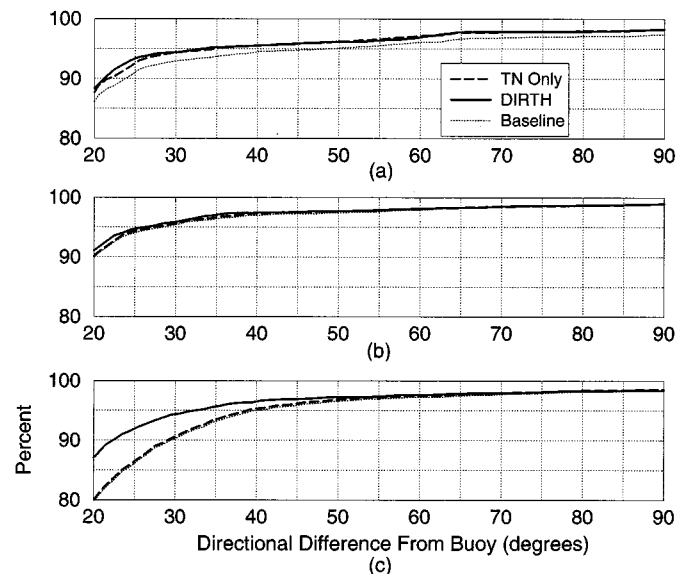

Fig. 8. Cumulative probability of directional difference from buoys. Three QSCAT data sets were compared: DIRTH, TN only and baseline data. Results were further divided into three cross track locations: (a) far swath (750–900 km), (b) sweet spot (400–700 km) and (c) nadir (0–300 km). The x-axis is the directional difference from the buoy wind vector. The y-axis is the cumulative probability. A point on the curve (x, y) indicates that $y\%$ of the buoy hits had directional differences less than $x°$. Buoy hits were restricted to a set of 24 moored NDBC buoys more than 100 km from land. A buoy was co-located with a QSCAT wind vector if it was within 30 min in time and 25 km in distance. Rain-flagged cells and cells with buoy wind speeds less than 7 m/s were omitted.

ranked (highest likelihood) ambiguity is one ($r(1) = 1$). We set the maximum rank for nudging by choosing the number of ambiguities above a certain threshold M in relative likelihood. For example, if there are four ambiguities and $r(1) > r(2) > r(3) > M > r(4)$, then the closest of the first, second, and third ambiguities to the nudging field is used to initialize the median filter. The optimal threshold is a function of the quality of the nudging field. The value used in the QSCAT data product, $M = 0.2$, was chosen to optimize ambiguity removal skill in simulation. As we shall see in Section III-C, comparisons of QSCAT data with buoys indicate that TN with $M = 0.2$ results in improvement in ambiguity selection over the baseline technique. These results imply that $M = 0.2$ is reasonable for real data but not necessarily optimal.

III. STATISTICAL ANALYSIS

In this section, we discuss empirical comparisons between the baseline wind retrieval technique and the DIRTH method. Three types of results are presented: comparisons to analytical wind fields, comparisons to buoys, and simulation results. Simulation studies performed prior to launch were used in the development of the DIRTH algorithm. Although the simulation results are superseded by results obtained later from real data, they have the advantage that they can be compared with "truth." Real QSCAT data can only be compared with analytical wind fields which have their own sources of error, or buoy measurements which are sparse and not entirely consistent with scatterometer measurements. The simulations were performed using an extensive high fidelity simulation package which has been developed over a number of years using lessons learned during the NSCAT mission (see [15] for an overview of the simulation strategy

Fig. 9.   Example-selected ambiguity wind field. An example of the selected wind vectors in the QSCAT wind data product as computed with TN but without direction interval retrieval. Note the noisy nature of the vectors along the nadir track.

employed and [16] for a detailed description of the QuikSCAT End-To-End Simulator). For these reasons, we have included a discussion of our simulation results along with the other comparison metrics.

A. Simulation Study

Twenty five orbits were simulated using $1° \times 1°$ ECMWF fields as truth. Fig. 4 illustrates the selected (after ambiguity removal) root mean square (RMS) direction error for the baseline and DIRTH techniques. Results are depicted as a function of cross track distance for four ranges of true wind speeds. Wind speed error results are depicted similarly in Fig. 5.

Applying DIRTH significantly improves the simulated directional RMS error across the swath for all four ranges of wind speeds. The most noticeable improvement is at nadir where RMS direction error is improved by more than $10°$ for all wind speed ranges. For most wind speed ranges and cross track positions the RMS speed error is also improved although not as significantly. One notable exception is a slight increase in RMS speed error (< 0.02 m/s) for wind speeds between 3 and 5.5 m/s and cross track distance between 400 and 700 km on both sides of the swath. This degradation in performance may be due to quantization errors in the interpolation scheme

used to calculate the DIRTH speed once the DIRTH direction has been determined. In any case the problem is small.

Because these results are simulated, there are of course a number of caveats worthy of examination. Perhaps, the most unrealistic aspect of the simulation is the use of $1° \times 1°$ true wind fields thereby eliminating small scale changes in the wind field from consideration. Clearly, the DIR method, because it makes greater use of spatial information in determining the wind vector, is more prone to losing high resolution information. Since DIR is applied preferentially to cases in which the quality of the directional information is subject to doubt, the loss of useful information is theoretically minimal. Clearly, the simulation results alone do not address this issue satisfactorily. We justify our choice of true field, by pleading the lack of realistic fields with small scale variation. All available small scale fields of which we were aware are obtained from previous scatterometers and thus have edge effects due to time difference between passes or ambiguity removal problems. Such edge effects might drastically impact median filtering thereby producing confusing results. Another major caveat of the simulation is that the geophysical model function was treated as the truth. Although random measurement errors were modeled extensively, the GMF was assumed to be free from systematic error.
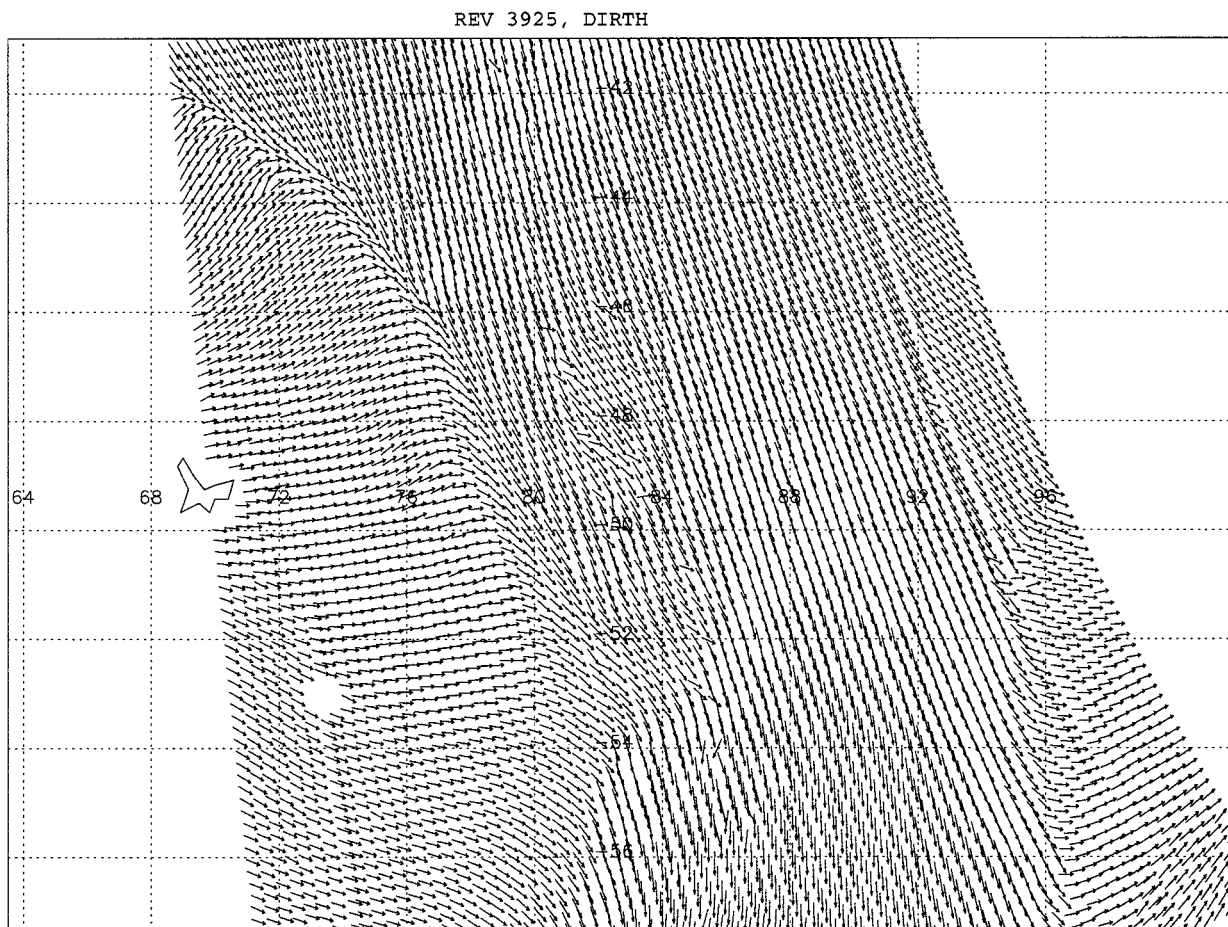
REV 3925, DIRTH

Fig. 10.   Example DIRTH wind field. The same wind field as in Fig. 9, with direction interval retrieval performed.

B. Comparisons to ECMWF Wind Fields

In this section, we compare QSCAT retrieved winds using both the baseline and DIRTH techniques with European Centre for Medium-Range Weather Forecasts (ECMWF) analysis wind fields. The $1° \times 1°$ ECMWF wind fields were used and interpolated spatially (but not temporally) to the QSCAT wind vector cell locations. ECMWF wind fields are produced every six hours and each 100 min QSCAT orbit was only co-located with a single ECMWF field, so that the greatest possible temporal difference is three hours and 50 min. When on occasion a particular ECMWF wind field was unavailable, the orbits temporally co-located with that field were left out of the analysis. Of course, one must take these comparisons with a grain of salt. Clearly, the time difference between QSCAT and ECMWF winds and the spatial resolution difference (25 km versus $1°$) can lead to substantial differences between the two wind fields even if both are error-free. For this reason, the magnitude of the difference between the fields indicates little about the absolute accuracy of either the QSCAT or ECMWF fields. However, comparing QSCAT with ECMWF is useful as means of comparing the relative accuracy of the DIRTH and baseline wind retrieval techniques, especially when used in conjunction with buoy comparisons and simulation results.

Fig. 6 depicts the RMS direction difference between ECMWF and the retrieved winds. The directional differences are plotted versus cross track distance for four ranges of ECMWF wind speeds. Fig. 7 depicts the RMS speed differences similarly. Wind vector cells in which the QSCAT rain contamination flag is set are omitted from the analysis.

Applying DIRTH reduces the directional differences from ECMWF significantly across the entire swath for all ranges of wind speeds. The impact is less than that observed in simulation with the maximal nadir improvements of $6°$ rather than the more dramatic $10°$ observed in simulation. The speed RMS difference values are similar for the baseline and DIRTH cases. The only substantial differences are slight advantages for DIRTH in the mid swath region for the two highest wind speed ranges.

The ECMWF comparisons differ from the simulated results in that the improvement due to DIRTH is smaller. There are several reasons for this result. Notice that the differences from ECMWF for both DIRTH and the baseline wind retrieval are much larger than the simulated error values. This result is expected because the difference from ECMWF includes not only the errors in the QSCAT retrieved winds but also errors in the ECMWF wind fields, temporal differences between ECMWF and QSCAT and any high resolution ($< 1° \times 1°$) information in the QSCAT fields. Also any systematic error in the model function will produce errors in both sets of QSCAT winds. Systematic model function error was assumed to be zero in simulation.

The addition of these extra sources of noise results in RMS directional and speed differences being greater than the simulated error values. It also results in a compression in the apparent improvement due to DIRTH. DIRTH can reduce the errors in the QSCAT retrieved winds and arguably smooth high resolution information, but it has no impact on other sources of directional difference such as time differences between ECMWF and QSCAT or error in ECMWF. Due to these additional sources of noise, the height of the nadir bumps in the baseline directional accuracy curve are much less in Fig. 6 than in Fig. 4. DIRTH dramatically reduces the nadir bumps in both figures. The additional noise sources also provide a possible explanation for the much poorer baseline directional performance and DIRTH improvement at lower wind speeds, as one would expect errors in the ECMWF wind fields and high resolution information to be more significant at lower wind speeds.

C. Comparisons to Buoy Winds

In this section, we compare QSCAT retrieved winds with co-located buoy wind measurements. Three wind retrieval methods are examined: DIRTH, TN-only, and the baseline wind retrieval algorithm. Due to varying buoy heights, the wind speed measurements are not easily comparable between buoys and scatterometer measurements. Wind speed varies significantly with distance from the ocean surface. The directions do not vary as significantly with height, so in order to simplify our analysis and discussion, we only present direction comparisons. We analyzed co-locations with National Data Buoy Center moored buoys between July 1999 and April 2000. Only co-locations within 25 km and 30 min and with buoy wind speeds within the range of 7–30 m/s were included in the analysis. Buoys within 100 km of land were excluded. QSCAT wind vector cells which were flagged for rain contamination were also omitted. In all, 9761 co-locations with 24 buoys were included in the analysis.

Fig. 8 depicts the cumulative probability of the differences between buoy wind directions and the directions from each of three cases: baseline wind retrieval, TN-only, and DIRTH. Cumulative probability functions are reported for three portions of the swath: nadir, far swath and the sweet spot. Nadir was defined to be within 300 km of the nadir track, sweet zone from 400 to 700 km and far swath from 750 to 900 km. Transitional regions were omitted in order to better demonstrate the cross-track variations in performance. The outer two cross-track locations on each side of the swath (cross track distance > 900 km) were also omitted. These cross-track locations are only observed sporadically and exhibit poor wind performance. The curves depicted in Fig. 8 are cumulative probability functions, with the x-axis representing directional difference from the buoys and the y-axis representing percentage of co-locations. For example if a curve passes through the point $(x, y) = (20, 90)$ that means 90% of the buoy co-locations had directional differences less than $20°$.

Fig. 8(a) depicts the directional difference probability functions for the far swath. TN results in a modest ambiguity removal improvement in the far swath. Direction differences from the buoys greater than $90°$ occur 2.5% of the time for the baseline wind vectors as opposed to 1.7% and 1.8% for the

DIRTH and TN-only cases. It is is important to note that the far swath had significantly fewer buoy hits (1580) than the nadir and sweet spot (3329 and 3300, respectively). Fig. 8(b) depicts the sweet spot direction difference probability functions. In the sweet spot, all three wind retrieval methods compare similarly with the buoys. Fig. 8(c) illustrates the nadir case. Clearly, the DIRTH wind retrieval scheme compares more favorably with buoys than does either of the other two cases. For the DIRTH case, the directional difference from the buoys exceeds $20°$ only 13% of the time as opposed to 20% for the other two wind retrieval methods. This difference is particularly dramatic, when one considers buoy directional errors as well as inherent differences in the measurements (point measurements versus 25-km averages) are contributing to the directional differences.

D. Example Wind Field

In addition to computing performance statistics, one can determine the effectiveness of a wind retrieval technique by observing the retrieved wind fields. Fig. 9 depicts a wind field obtained using the TN algorithm only. Fig. 10 depicts the same field with the DIR algorithm also employed. Notice that DIR cleans up the noisy wind vectors along the nadir track without impacting the abrupt change in wind direction in the northwest corner of the wind field.

IV. SUMMARY

In conclusion, end-to-end simulation studies and comparisons of QSCAT data to analytical wind fields and buoy winds indicate that the DIRTH wind processing algorithms reduce errors in the wind directions of the QSCAT wind vectors. Performance improvement is especially evident for wind vector cells near the nadir track. Comparisons to buoy winds also indicate that TN results in a modest improvement in ambiguity removal at far swath.

While, both simulation studies and analytical wind field comparisons fail to address the question of whether or not DIR is oversmoothing the data, the comparisons with buoy winds do address this question somewhat: one would expect an oversmoothed wind field not to compare favorably with buoys (clearly the higher the resolution the more closely the wind field approximates point measurements). Near nadir, the DIRTH winds compare more favorably with buoys than do winds produced without DIR. In the rest of the swath, DIR and the the standard technique compare similarly with the buoy data. These results suggest that any oversmoothing resulting from DIR is minimal. Furthermore, any oversmoothing in the nadir region is overshadowed by noise reduction, yielding significantly improved buoy comparisons. Nonetheless, the question of oversmoothing still deserves further research. There are two DIRTH parameters which can be adjusted to further optimize the noise reduction/oversmoothing tradeoff. These are the error bar probability T (currently 80%) and the median filter window size (currently 7×7). Developing methods to improve these parameters would be an interesting avenue for further research.

REFERENCES

[1] F. Naderi, M. H. Freilich, and D. G. Long, "Spaceborne radar measurement of wind velocity over the ocean—An overview of the NSCAT scatterometer system," *Proc. IEEE*, vol. 79, pp. 850–866, June 1991.

[2] W. L. Jones *et al.*, "The Seasat—A satellite scatterometer: The geophysical evaluation of remote sensed winds over the ocean," *J. Geophys. Res.*, vol. 87, pp. 3297–3317, 1982.

[3] L. C. Schroeder, P. R. Schaffner, J. L. Mitchell, and W. L. Jones, "AAFE RADSCAT 13.9-GHz measurements and analysis: Wind-speed signature of the ocean," *IEEE J. Oceanic Eng.*, vol. OE-10, pp. 346–357, Apr. 1985.

[4] F. J. Wentz, S. Peteherych, and L. A. Thomas, "A model function for ocean radar cross sections at 14.6 GHz," *J. Geophys. Res.*, vol. 89, no. C3, pp. 3689–3704, 1984.

[5] M. H. Freilich, B. A. Vanhoff, and R. S. Dunbar, "Empirical determination of a Ku-band wind model function from the seawinds scanning scatterometer," *J. Geophys. Res.*, to be published.

[6] F. J. Wentz and D. K. Smith, "A model function for the ocean normalized cross section at 14 GHz derived from NSCAT observations," *J. Geophys. Res.*, vol. 104, no. C5, pp. 11 499–11 514, May 1999.

[7] W. J. Plant, "A two-scale model of short wind-generated waves and scatterometry," *J. Geophys. Res.*, vol. 91, no. C9, pp. 10 735–10 749, September 1986.

[8] M. W. Spencer, C. Wu, and D. G. Long, "Improved resolution backscatter measurements with the SeaWinds pencil-beam scatterometer," *IEEE Trans. Geosci. Remote Sensing*, vol. 38, pp. 89–104, Jan. 2000.

[9] P. L. Combettes, "The foundations of set theoretic estimation," *Proc. IEEE*, vol. 81, pp. 182–208, Feb. 1993.

[10] D. G. Long and J. M. Mendel, "Identifiability in wind estimation from scatterometer measurements," *IEEE Trans. Geosci. Remote Sensing*, vol. 29, pp. 268–276, Feb. 1991.

[11] P. L. Combettes and H. J. Trussell, "The use of noise properties in set theoretic estimation," *IEEE Trans. Signal Processing*, vol. 39, pp. 1630–1641, July 1991.

[12] E. Walter and H. Piet-Lahanier, "Guaranteed linear and nonlinear parameter estimation from bounded-error data: A Survey," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 1, Gif-sur-Yvette, France, 1993, pp. 774–777.

[13] S. J. Shaffer, R. S. Dunbar, S. V. Hsiao, and D. G. Long, "A median-filter-based ambiguity removal algorithm for NSCAT," *IEEE Trans. Geosci. Remote Sensing*, vol. 29, pp. 167–174, Jan. 1991.

[14] H. Schultz, "A circular median filter approach for resolving directional ambiguities in wind fields retrieved from spaceborne scatterometer data," *J. Geophys. Res.*, vol. 95, no. C4, pp. 5291–5303, Apr. 1990.

[15] J. Huddleston, W. Tsai, M. Spencer, and R. West, "Modeling and simulation for seawinds-1B system design and performance evaluation," *EOS/SPIE Aerosp. Remote Sensing J.*, vol. 3221, pp. 357–364, Sept. 1997.

[16] W.-Y. Tsai, S. V. Nghiem, J. N. Huddleston, M. W. Spencer, B. W. Stiles, and R. D. West, "Polarimetric scatterometry: A promising technique for improving ocean surface wind measurements from space," *IEEE Trans. Geosci. Remote Sensing*, vol. 38, pp. 1903–1921, July 2000.

[17] E. P. W. Attema, "The active microwave instrument on-board the ERS-12 satellite," *Proc. IEEE*, vol. 79, pp. 791–799, June 1991.

Bryan W. Stiles received the B. S. degree in electrical engineering from the University of Tennessee, Knoxville, in 1992, and the Ph.D. degree from the University of Texas, Austin, in 1997, where his dissertation proved the universal approximation capability of certain dynamic artificial neural networks.

In June 1997, he became a Member of the Technical Staff, Jet Propulsion Laboratory, Pasadena, CA, where he has simulated the behavior of spaceborne scatterometers (NSCAT and SeaWinds) as well as analyzed the data returned from these instruments. He is interested in statistical techniques for determining and visualizing relationships between remote sensor data and geophysical phenomena.

Brian D. Pollard received the B.S. in electrical engineering from the University of Houston, Houston, TX, in 1992. He joined the Microwave Remote Sensing Laboratory (MIRSL), University of Massachusetts, Amherst, in 1992, where he received the Ph.D. degree in electrical engineering in 1998.

His work at MIRSL included the development of a digitally beamformed volume-imaging radar for atmospheric boundary layer (ABL) studies and compared data from that instrument to large-eddy simulations of the ABL. He joined the Jet Propulsion Laboratory, Pasadena, CA, in 1998, where he studied the design and application of scatterometry, altimetry, and interferometry for airborne and spaceborne remote sensing of the ocean surface. His interests include high-resolution scatterometry and interferometry for oceanic, terrestrial, and planetary applications.

R. Scott Dunbar is a Senior Engineer at the Jet Propulsion Laboratory, Pasadena, CA. He has been involved in the NASA Scatterometer (NSCAT) and SeaWinds projects since 1985, contributing to science algorithm development and ground system operations.