N O T I C E

THIS DOCUMENT HAS BEEN REPRODUCED FROM MICROFICHE. ALTHOUGH IT IS RECOGNIZED THAT CERTAIN PORTIONS ARE ILLEGIBLE, IT IS BEING RELEASED IN THE INTEREST OF MAKING AVAILABLE AS MUCH INFORMATION AS POSSIBLE

WINDFIELD RETRIEVAL FROM DOPPLER LIDAR OBSERVATIONS

Interim Technical Report No. 1

Contract NAS8-33389

June 1980

R. W. Lee

LASSEN ⏚ RESEARCH

MANTON CALIFORNIA 96059
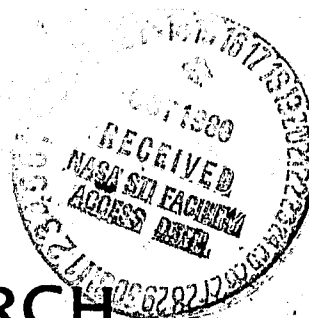
Abstract

An approach to the retrieval of a vector wind field from Doppler lidar observations is developed in general terms. The field of radial velocity measurements from each look angle is modeled by a smooth surface, the parameters of the model being determined from the data by least-squares techniques. The vector wind field and higher-order fields are obtained from the two modeled surfaces. Estimated measurement errors are taken into account, and error estimates are available for all output data sets.

Table of contents

List of figures

I. Introduction

This report is concerned with the problem of retrieving vector wind measurements in a plane from radial wind measurements made in that plane using two different look angles. There are several possible approaches to this problem, but certain characteristics of the data require a certain amount of care in the selection of the approach.

The measurements are made by an airborne Doppler lidar system. The sensitivity of this system is limited, and the windfield tracer - naturally occurring aerosols - is not always present in sufficient quantities for a satisfactory return. For this reason the radial velocity measurements will vary greatly in accuracy, causing errors in the inferred vector field which will be magnified by the less than perfect orthogonality of the two look angles. Finally, the utility of the higher-order attributes to be derived from the vector windfield - vorticity and divergence, for example - will be limited by noise and error in the inversion process.

These considerations clearly suggest an inversion algorithm which is tolerant of gross measurement errors and which minimizes the effects of random errors in the data. Achieving these goals without greatly reducing the resolution of the measurements requires care.

Section II below consists of a more precise definition of the inversion problem. Elements of this definition are presented in detail in section III. The primary inversion steps of editing and smoothing are considered in section IV. Two related topics are discussed briefly in section V, while the steps required for implementation and evaluation of the suggested algorithm are outlined in section VI. This final step - the evaluation of the solution algorithm - is just as important as the selection of the algorithm itself. Only when the error characteristics of the solution are known can one interpret the measurements with confidence.

II. Problem definition

Simply stated, the goal is to derive a two-dimensional vector field from
two scalar fields. In practice, most of the data manipulation will take place
on the separate scalar fields (see section V.B, however, for an alternate
strategy), and conversion to a vector field becomes a final, trivial operation.

Since a solution to the problem posed depends upon the definition of that
problem, some care must be exercised in characterizing the problem itself. Of
particular importance here are certain problems inherent in the measurements,
and the need for a considered definition of the properties desired in the
solution.

II.A. Problems with the data

The measurements of the two radial velocity fields may be sparse (missing
data points), and they certainly will vary greatly in quality. These two
attributes will be pre-eminent in the selection of a solution algorithm. In
addition, the measurements are made on irregular grids, with no agreement
between the measurement points in the two scalar fields. Finally, the two
scalar observations are not orthogonal, making the conditioning of the scalar-
vector transformation less than ideal.

II.B. Desired results

It is not easy to quantify the characteristics required in the solution
data set, but it is important to attempt to do so. The solution algorithm
will be optimized according to these criteria, so obviously the criteria must
be appropriate.

The general requirement is for smooth flow fields of known accuracy.
These fields must be useful in a visual sense, and the errors in the statistical
properties of these fields must be known and acceptably small.

II.C. The solution in general terms

Clearly the path to the solution is one of data smoothing and interpolation.
These steps must be accomplished in a way that is optimal given the data
characteristics and the desired goals of the solution field. In addition, the
solution must be efficient in that it makes use of all a priori information
about the measurements.

II.D. Steps required for solution implementation

In adaition to careful statement of the problem and the desired solution goals, selection of an algorithm requires a procedure for evaluation of the algorithm according to those goals. Such an evaluation will require operation on real and simulated data sets, with visual and statistical interpretation of the end products. Under these conditions the algorithm can be "tuned" and optimized for the actual characteristics expected of the raw data.

III. Expanded problem definition

III.A. Specific nature of data relevant to optimal field retrieval

As mentioned above, the data is sparse. This may be due to range
attenuation, lack of aerosols, or both. While missing data points will be a
particular problem at long range, they can occur at any range.

The quality of the measurements will vary greatly. The Doppler estimator
used for measurement of radial velocity has certain (known) error characteristics
depending upon the signal-to-noise ratio of the signal return and the spectral
width of the signal. The probability distribution of this error is sketched
in figure 1. It consists of two parts: a normally-distributed term N and a
uniformly-distributed term U. At very low signal-to-noise ratios the uniformly-
distributed term gives rise to "wild" estimates of radial velocity, which
account for most of the error variance.

In addition to estimation errors produced by the Doppler processor there
is an inherent sampling error present regardless of the signal-to-noise ratio.
Since the Doppler return has finite bandwidth any finite realization of that
return will be subject to an error variance due to sampling. At high signal-
to-noise ratios this is the dominant error term.

The locations at which radial velocities are observed do not form a
regular grid. The angles at which the observations are taken will vary
slightly due to aircraft dynamics, as shown in exaggerated form in figure 2.
The gridpoints of one scalar field will bear little or no relationship to
those of the other field. The observation angles will not differ by the desired
90 degrees, but by approximately 40 degrees. This lack of orthogonality will
magnify data errors in the vector component parallel to the aircraft track.

The spatial sampling of the laser beam presents further problems. The
beam resolution cell is long (~300 m) and narrow (~30 cm radius); each
measurement will consist of the average of a number of such cells displaced
horizontally by approximately 1 meter. Such a resolution volume will average
effectively in one spatial dimension (range), but not in the other two. The
resulting under-sampling of the spatial windfield (particularly in the vertical
dimension) will lead to an increased error variance in the data from aliased
energy in the spatial windfield. This is further complicated by the fact that
the measurements are not necessarily made in the desired plane, but at a
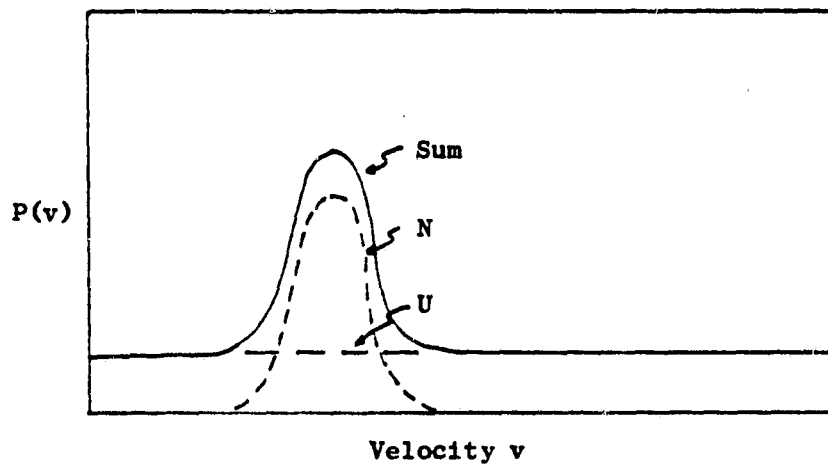horizontal angle subject to random excursions about zero. Indeed, the effect

Figure 1: Probability distribution of Doppler estimation errors for a certain signal-to-noise ratio. N and U are normally- and uniformly-distributed variance components.
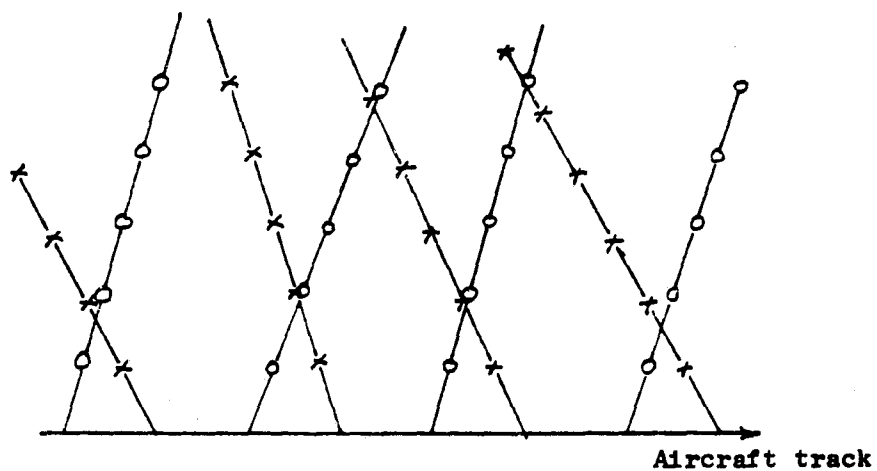

Figure 2: Exaggerated observation grid. X and O represent backward and forward look angle data points.

will be reduced by intentionally jittering the vertical axis of the system to achieve some averaging in that dimension.

If the turbulent parameters of the windfield can be estimated - and this should be possible from the data itself - then the variance due to under-sampling can be estimated.

Finally, there is a data problem due to the fact that the two scalar fields do not result from simultaneous observations. The time lag between the observations varies with range, being on the order of 30 seconds at maximum range. This is sufficient time for the windfield to translate by one or more resolution cells, and it may be necessary to advect the measurements according to the local mean wind vector for an appropriate time interval (which may vary with range) in order to achieve the required time registration of the measurements.

III.B. Inversion goals

The first use of the vector flow field will be visual, in the understanding of flow fields in the clear-air vicinity of severe storms. Thus the inferred vector field must be sufficiently smooth to allow the eye to trace parcel flow. Stated another way, the continuity between neighboring vector estimates must be reasonably high.

Aside from visual aesthetics the vector field must possess certain statistical properties. Useful vorticity and divergence fields must be obtainable from it. This in turn requires that the errors in the estimated divergence and vorticity be below some threshold (perhaps a certain percentage error) for a substantial percentage of the observations, and in addition that the errors themselves can be estimated. Obviously a vorticity estimate of $0.01/sec$ is not useful if the standard error is $0.1/sec$, or if the error is unknown. The allowable errors in vector wind and higher-order functions, expressed as percentages at some confidence level, are very important inputs to the algorithm optimization and evaluation process.

Finally, the inverted data must be as complete as possible. Gaps in the derived fields must be filled where possible by interpolation, even though this may reduce resolution in the vicinity of the missing data points. Care must be exercised in such interpolation to ensure that good data points are not contaminated by bad.

III.C. Generalized solution

In general smoothing and interpolation are achieved through the reduction in the degrees of freedom of the solution field below those of the input data, and through careful selection of those degrees of freedom retained in the solution. Only by sacrificing degrees of freedom - and here resolution - in the solution field can data reliability be enhanced and error variance estimated. Stated in spatial terms, resolution is reduced to obtain greater data stability. In spectral terms, knowledge of field components of high spatial frequency is sacrificed so that knowledge of lower spatial-frequency components will be enhanced. Clearly a balance must be struck somewhere in between the extremes: perfect knowledge of the mean wind on a 10x10 km square is of little value in the severe-storm program, as is no knowledge on a 100x100 m grid.

The general process of smoothing and interpolation is one of modeling. One selects a mathematical framework by which to model the output field, and determines the parameters of that model from the measurements.

The success of the smoothing process depends to a large extent upon the suitability of the model selected. The model must be appropriate in several respects. First it must be able to represent the natural features of the windfield adequately. It must be possible to control the spatial-frequency response of the model readily, to allow control over the smoothness of the solution. Finally, it must be mathematically tractable: a model is not useful if it takes an hour of computer time to invert a minute's worth of data.

Once a model is selected the parameters of that model may be determined from the data. Such a parameter fit may be achieved in a straightforward manner using least-squares techniques, but note that this requires use of objective error criteria. The solution is optimal in terms of these criteria, but one must make sure that the criteria are appropriate. In general these criteria will involve some compromise between spatial resolution and data stability, between smoothness or continuity and error variance.

Given that a model has been selected and the solution obtained for a certain set of error criteria, it remains to evaluate the probable errors in the various end products - wind vectors, vorticity, etc.

III.D. Solution implementation

The generalized solution outlined above can be implemented in the following sequence:

1) Obtain data. In addition to recording the raw data, this step includes obtaining all ancillary data which will be useful in data interpretation - time and location, look angles, other meteorological data, etc.

2) Establish data reliability. From the signal-to-noise ratio estimate and other parameters, estimate the probable error of the velocity estimate. This error may consist of two parts with different probability distributions.

3) Editing. The data must be put into a form that the smoothing algorithm can use. In addition to assigning a weight to the data point reflecting its probable error, and assigning to that point coordinates, spectral width, etc., an additional operation is desirable. The measurement may be compared with neighboring points (in two dimensions) as a test of measurement continuity. If the measurement is discrepant its weight may be reduced. This process will remove to a great extent the effects of "wild" estimates produced by the uniform portion of the Doppler estimator error. This operation is explored in more detail in section IV.A.

4) Smoothing. This step includes solving for the model parameters in terms of the weighted data, interpolating where required, and reducing the resolution of the measurements where data quality is low. Estimates of error variances should be carried through this process. Finally, the solution field can be sampled on any desired grid. This operation is described in detail in section IV.B.

5) Produce vector field. The two scalar fields may be combined to form the vector flow field, again carrying through the estimated errors.

6) Produce end products. Higher-order fields may be obtained through operations upon the vector flow fields, in each case carrying error estimates through the process.

7) Evaluation. The last step is to evaluate the utility of the resulting end products. If there are serious problems with them, it must be determined from error propagation which aspect of the raw data most seriously compromises the result. If that data aspect cannot be corrected, it should be determined whether or not alterations in the model can reduce the effects.

IV. Detailed implementation

IV.A. Data editing

The goal of the data editing process is to produce data of known error characteristics for the smoothing algorithm. All information available must be brought to bear in evaluating a given data point. The following list includes the most important factors.

1) Signal-to-noise ratio. From the signal amplitude estimate at the range gate of interest, in comparison with the amplitude estimates at very large ranges (where no signal is expected), an estimate of the signal-to-noise ratio can be obtained. This estimate can be used in conjunction with the (known) error characteristics of the Doppler estimator to produce a probable velocity error estimate consisting of two parts as suggested above: a normally-distributed component and a uniformly-distributed component.

2) Spectral width. The Doppler estimator produces as a matter of course an estimate of the signal spectral width. Spectral width enters into the Doppler processor error equations. Note however that useful estimates of spectral width are not produced at very low signal-to-noise ratios.

3) Continuity. Continuity may be used in two dimensions as a check upon data consistency. Continuity tests may be applied to amplitude and width estimates as well. In a typical case the weighted median value of the eight neighboring points might be compared with the point in question. Note that the median or most probable value is more useful here than the mean value, since the mean can be severely disturbed by a single bad data point.

4) Constraints. The expected characteristics of the windfield can be used as a further check upon data integrity. For example, a constraint upon velocity gradients (shear) can be used as an input for continuity tests. Limits may be set upon maximum values of velocity as a test for reasonableness. As with all constraints, care must be exercised to ensure that actual features of the wind field which were not expected are not obscured. Use of adaptive or interactive constraints can achieve this goal.

Should a given data point fail one or more tests for reliability, the weight of that point may be reduced, or in severe cases a missing data point may be declared.

Note that the editing process can be combined with the smoothing or filtering process. A first fit of the data points gives a trial solution and

a deviation for each data point. These deviations are a measure of continuity and can be used to alter the weighting given the data. A second iteration of the solution gives a revised output field.

IV.B. Data smoothing

IV.B.1 General

Data points can be considered in isolation, but since smoothing implies that each data point has an influence upon its neighborhood it is useful to consider the two-dimensional measurements as forming the height of a two-dimensional surface. The process of data smoothing then becomes one of fitting a surface of a certain character as nearly as possible to the measurements, in (for example) a least-squares sense.

Each form of data manipulation can be interpreted in terms of a surface of a certain type. For example, point data may be considered to form a surface composed of blocks centered at the measurement points, the height of each block indicating the value of the measurement at that point. That is, the data point is the altitude of the surface for that grid square. A continuous surface can be created by placing the data points at the vertices of the surface, with straight lines joining the vertices defining the surface (that is, linear interpolation between the data points, with grid rectangles broken into two triangles by a single diagonal). Surfaces formed with continuous first- or higher-order derivatives require the overlapping influence of several measurement points at each point on the surface.

IV.B.2 Continuous surfaces

Continuous surfaces may be modeled by many analytic or elementary functions. The most commonly used functions are polynomials (including splines, Hermite and other orthogonal polynomials), Fourier series, Bessel functions and spheroidal functions. The choice of a basis function depends upon:

1) Suitability for the problem. Some functions lend themselves to a particular coordinate system or situation. For example, Bessel functions are often appropriate for cylindrical coordinates, and Fourier series for band-limited functions.

2) Mathematical ease of use. Polynomials, for example, offer few difficulties in integration or differentiation, no convergence problems, etc.

3) **Parameter flexibility.** The degrees of freedom of some functions can be easily "tuned" to control the parameters of the surface. For example, the small-scale wiggles of a surface modeled by a Fourier series are easily controlled by limiting the order of the series.

IV.B.3. Surface adaptability

Since the quality of the data varies from point to point on the surface, it may be reasonable to allow the nature of the surface to vary as well. That is, in regions of high data quality the smoothness constraint on the surface may be relaxed to allow the model to represent smaller-scale features. Conversely, in regions of poor or missing data surface smoothness must be constrained even further to preserve surface continuity. This trade-off between surface smoothness and resolution may be made in an adaptive manner, with the algorithm itself sensing the need for constrained smoothness.

IV.B.4. Suitable models for flow fields

Due to its easily-controlled spatial-frequency response, a Fourier surface is attractive. However, the difficulty of incorporating data of varying quality, on a non-uniform grid, is substantial. The model of choice is a locally-defined polynomial with a basis function of limited extent. Such a model offers ease of solution using least-squares techniques, no grid problems, controllable basis size and smoothness, and continuity to any desired derivative. Suitable basis functions would include linear, quadratic or higher functions over limited (sliding) basis regions, or spline functions.

As an example, consider the lowest-order surface fit. A region of influence is defined around a point for which surface height is to be estimated. Data points in this region of influence are summed in a weighted average, the weights being derived from the error variances assigned to those points, with (for example) an additional weighting function formed by a two-dimensional Gaussian centered at the estimation point. This weighted averaging is equivalent to fitting a local plane to the data in the vicinity of the estimation point. To achieve the accuracy desired, the size of the region of influence (defined by the two-dimensional Gaussian weighting function) can expand or contract as required to enclose a suitable number of measurement points. Such an approach is easily implemented, and sliding the Gaussian region of influence around the

plane gives a continuous estimation surface. However, with this simple approach
shear in the windfield cannot be fitted by the model at each point; the result
is a poor fit requiring increased smoothing.

Use of the next higher order surface solves the shear problem: at each
estimation point one fits the height and slope of a plane surface. Shear is
no·longer a problem - the fitting errors are limited to curvature and higher-
order derivatives.

At some point increasing the order of the model (that is, increasing the
degrees of freedom in the solution) increases noise in the surface beyond a
tolerable level. The optimum surface order remains to be determined; there
will be a compromise between higher order and reduced region of influence which
must be determined by simulation.

Splines are a particularly attractive form of polynomial basis function
since the approximating functions are easily constrained to be continuous on
the boundaries between grid points.

IV.B.5 Model fitting

Once a model has been selected and the controllable parameters defined,
it remains to determine those parameters. The most suitable solution technique
is the linear least-squares approach. The variable weights of the data points
are easily taken into account, along with additionally-imposed geometrical
weighting. One particular advantage of this approach is the availability
with the solution of an estimate of the solution error variance.

Once the parameters of the surface have been determined, that surface
may be sampled at any desired grid pattern.

V. Short topics

V.A. Spectral width and signal amplitude

Although they have received little attention thus far, the signal amplitude and spectral width are also measured by the Doppler processor. These two quantities may also be considered to form solution surfaces, and the same techniques described above may be applied to the estimation of the parameters of these surfaces: editing, smoothing and interpolation.

Additional redundancy is present in these measurements, since information from the two look angles may be combined.

Note that a portion of the apparent spectral width may be contributed by horizontal velocity shear within the target volume. Since the velocity field is being determined independently, it is possible to correct for this contribution under the assumption that the shear variance and the spectral width add incoherently.

V.B. An alternate solution strategy

While this report has treated the data from the two look angles as being independent until their combination in the vector field, another approach is possible.

The wind field model may assume a single surface as a potential field. The measurements become directional derivative estimates of this surface, and techniques for surface reconstruction from derivative information can be used.

Note however that this process is strictly valid only when the divergence of the actual wind field is zero. Thus the potential field so derived will naturally produce a zero-divergence field. Divergence may then be recovered from the measurements by a second-stage solution, solving for a divergence field from the difference between the measurements and the inferred zero-divergence field. There may be a problem here since the divergence and circulation may be locally correlated.

VI. Algorithm implementation and evaluation

VI.A. Implementation of the surface-fitting model

This section is an outline of the steps required to take this solution technique from the generalized concept described in this report to the functional stage. The primary questions to be resolved at this point are:

1) Definition of the most appropriate surface model

2) Definition of the appropriate solution technique for that model

These two questions are inextricably joined. Their solution will arise through an iterative process of evaluation and optimization.

Once a model and a solution technique are chosen, they must be "tuned up" with reference to the practical problems of the data at hand. This must be done with the data end use firmly in mind, employing a well-defined evaluation technique and a set of evaluation criteria. Such tuning will determine the appropriate editing and weighting schemes, grid sizes and spatial resolution.

The suitability of the resulting algorithms is critically dependent upon the accuracy and realism of the evaluation technique, discussed in the next section.

VI.B. Algorithm evaluation

Since the solution technique will be optimized through interaction with a process of evaluation, the technique will ultimately by optimal only in terms of that evaluation procedure. Only if the evaluation procedure reflects the realities of the data and the wind field can one expect the solution algorithm to be optimal for the data.

In addition to providing a test bed for optimizing the solution algorithm, evaluation provides two important byproducts:

1) Confidence in the results. If the user can take a real or synthetic wind field, probe it with a simulated lidar system, contaminate the data with reasonable errors, and still retrieve a useful approximation of the original wind field, then he can have some confidence in using the algorithm upon data for which there is no confirming data.

2) Error propagation. By use of simulation the errors in user products can be estimated in terms of the errors in the raw data. User products without error variance estimates are of marginal utility; this is especially true of higher-order products such as convergence.

The following items may be taken as defining the components of an evaluation program:

1) Goals. A set of target goals should be established, in probabilistic terms. For example, one might desire that the vector wind components be measured to 2 m/s 90% of the time, or that vorticity be accurate to $10^{-3}s^{-1}$ on a 1-km scale.

2) Input data sets. Both simulated and actual wind fields (taken from multiple-Doppler observations) are of value - the former for their controlled nature and the latter for their realism. Obviously the statistical properties of these fields must be accurately known.

3) Signal-to-noise ratio. For simulation purposes, realistic signal-to-noise ratios must be assigned to the data points on a random basis. This would include range variation, dropouts, Rayleigh statistics, etc. This signal-to-noise ratio will be used to assign probabilistic errors to the radial velocity simulations, so it is important that they be realistic.

4) Wild measurements. In the transition from signal-to-noise ratio to velocity error, an appropriate number of totally-random estimates must be included to reflect the component of Doppler estimator error which is uniformly distributed.

5) Geometry. The grid points and look angles should be varied in a random way with reasonable values of variance.

6) Solution. Given the velocity field as probed by the synthetic lidar system - noise and all - an estimate of the original velocity field may be obtained by using the solution algorithm under test.

7) User products. The output wind field estimate may be transformed into the desired end products: visual fields, statistics, higher order fields.

8) Evaluation. The errors and utility of the user products must be assessed through comparison with the initial data set, using the evaluation goals as criteria for success.

The results of this evaluation may suggest alterations in the model or solution technique, or may suggest that certain user products cannot be reliably obtained from data of the quality simulated. By varying the characteristics of the input data set, the sensitivity of the inversion process to data problems of a given type may be determined. These sensitivity factors may suggest certain constraints upon the experimental operation, in order to improve recovery of a given user product.

VII. Conclusion

This report has suggested an approach to the retrieval of a wind field estimate from lidar measurements. This approach seems likely to draw the maximum amount of useful information from that data. Note however that some degredation of system resolution is required.

The emphasis of this approach is upon error analysis at all stages of the solution. It is felt that user products (wind fields, divergence fields, etc.) without explicit error estimates and confidence levels are of marginal value. This is particularly true of smoothed fields. A smoothed random field cannot be distinguished from the smoothed fields reported by dual Doppler observers, and one should have no confidence in such highly mathematical products unless shown an error propagation example.

With such error analysis techniques, including a carefully planned evaluation technique, one can be confident that one will know when the derived wind field has significance. This apparently modest claim is highly important when an experiment is likely to have marginal results: it is far more preferable to have a few good wind fields of known reliability - even if they represent only a small portion of the measured fields - than to have results of doubtful validity for all the data sets.