DOT/FAA/NR-91/7

Project Report
ATC-181

Microburst Divergence Detection for Terminal Doppler Weather Radar (TDWR)

M. W. Merritt

25 September 1991

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS

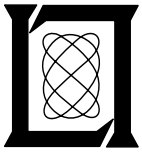

Prepared for the Federal Aviation Administration,
Washington, D.C. 20591

This document is available to the public through
the National Technical Information Service,
Springfield, VA 22161

TECHNICAL REPORT STANDARD TITLE PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
DOT/FAA/NR-91/7		

4. Title and Subtitle	5. Report Date
Microburst Divergence Detection for Terminal Doppler Weather Radar (TDWR)	25 September 1991
	6. Performing Organization Code
7. Author(s)	8. Performing Organization Report No.
Mark W. Merritt	ATC-181
9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
Lincoln Laboratory, MIT P.O. Box 73 Lexington, MA 02173-9108	11. Contract or Grant No. DTFA-01-89-Z-02033
12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
Department of Transportation Federal Aviation Administration Systems Research and Development Service Washington, DC 20591	Project Report
	14. Sponsoring Agency Code

15. Supplementary Notes

This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. The work was sponsored by the Department of the Air Force under Contract F19628-90-C-0002.

16. Abstract

The Terminal Doppler Weather Radar (TDWR) microburst surface divergence detection algorithm has been under development and evaluation at Lincoln Laboratory since 1983. The TDWR program is sponsored by the Federal Aviation Administration (FAA), and the algorithm described in this report is a primary algorithm component of the TDWR system. The divergence algorithm processes radar velocity measurements taken near the earth's surface to identify the strong divergent outflow characteristic of microburst wind shear hazards. The algorithm uses a complex set of pattern matching and validation test criteria to locate microburst outflow signatures and to filter out false alarms from various data contamination sources. The divergence algorithm is primarily responsible for the detection of most microbursts, although the complete TDWR microburst algorithm consists of more than a dozen distinct algorithmic components.

The divergence algorithm has demonstrated a very high probability of detection (POD) for strong microburst outflows, and its performance (as well as that of the complete microburst detection algorithm) was first formally assessed in the operational test and evaluation of the TDWR in Denver, CO (1988). Subsequent evaluations were performed in Kansas City, KS (1989) and Orlando, FL (1990). These evaluations have provided insight into the algorithm and system performance in a variety of meteorological and geographical environments.

This report describes the detailed operation of the divergence detection algorithm, its coupling to the remainder of the TDWR microburst algorithm, and the rationale for the various algorithmic components in the procedure. Performance is illustrated in detail by an important microburst event on July 11, 1988 (Denver), and it is statistically evaluated using a set of three active days from the 1988 measurement program.

17. Key Words		18. Distribution Statement
weather radar	Doppler radar	Document is available to the public through the National Technical Information Service, Springfield, VA 22161.
pattern recognition	TDWR	
microburst	divergence detection	

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	164	

ACKNOWLEDGMENTS

A number of people have made substantial contributions to the development and evaluation of the divergence algorithm over the past five years. The TDWR testbed operations team, particularly Mark Isaminger, Charles Curtiss, Paul Biron and Nat Fischer have provided the vast collection of TDWR testbed radar data and ground truth used to design the TDWR algorithms. Mark Isaminger has also been a constant source of ideas and observations relating to the performance and improvement of the algorithm techniques. Richard DeLaura, Robert Hallowell and Martin Eby provided the rich set of analysis and visualization software tools necessary for the radar data analysis and divergence algorithm evaluation presented in this report. Darelyn Neilley cheerfully handled the difficult task of organizing, editing and scoring the radar data, algorithm archives and ground truth data obtained from the testbed field programs. The conceptual microburst model presented in Chapter 2 was inspired by James Evans who has provided valuable insight, support and guidance throughout the algorithm development effort. Steven Campbell, Jeff Stillson, Terri Noyes and Wes Wilson have provided much assistance and insight in the algorithm development effort and in the preparation of this report. Special thanks go to Leslie Mahn for her careful editing and review of this report.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF ILLUSTRATIONS (CONTINUED)

LIST OF ILLUSTRATIONS
(CONTINUED)

LIST OF TABLES

1. INTRODUCTION

The microburst surface divergence detection algorithm is a central component of the automated processing performed by the Terminal Doppler Weather Radar (TDWR) system. This algorithm is responsible for processing the radar velocity measurements taken near the earth's surface to identify the strong divergent outflow characteristic of microburst windshear hazards. The divergence algorithm has been under development and evaluation at Lincoln Laboratory since 1983 and will be the primary algorithm component of the production TDWR systems to be deployed by the Federal Aviation Administration (FAA) in the early 1990's.

The divergence algorithm makes use of a complex set of pattern matching and validation test criteria to locate microburst outflow signatures and to filter out false alarms from various data contamination sources. These divergence signature detections are then merged with additional algorithm features detected at higher altitudes and subjected to adaptive strength and persistence tests to arrive at final microburst alerts to system users. While the complete microburst algorithm operating in the TDWR consists of more than a dozen distinct algorithmic components, the divergence algorithm is primarily responsible for the detection of most microbursts.

The development and evaluation of the divergence algorithm has been based on extensive measurements of microbursts from the TDWR testbed radar system under operational conditions, and the algorithm has demonstrated a very high probability of detection (POD) for strong microburst outflows. The detection and false-alarm performance of the divergence algorithm (and the complete microburst detection algorithm) were first formally assessed in an operational test and evaluation of the TDWR conducted in 1988 at Denver, CO. Subsequent operational evaluations performed in Kansas City (1989) and Orlando (1990) have provided insight into the algorithm and system performance in a variety of meteorological and geographical environments.

This report describes the detailed operation of the divergence detection algorithm, its coupling to the remainder of the TDWR microburst algorithm, and the rationale for the various algorithmic components in the procedure. The performance of the divergence algorithm is illustrated in detail on an important microburst event from July 11, 1988 (Denver) and is statistically evaluated using a set of three active days from the 1988 measurement program.

1.1. Description of the TDWR System Testbed

The TDWR testbed system has been developed and operated by Lincoln Laboratory to assist in the evaluation of radar and algorithm designs for the TDWR production system being procured by the FAA. The testbed provides a functional emulation of the performance-critical aspects of the TDWR and has been used to collect a very large database of operational measurements in variety of geographical and meteorological environments. The basic TDWR system consists of a doppler weather radar, a data processing subsystem and a set of user displays (Figure 1), all of which are present in the

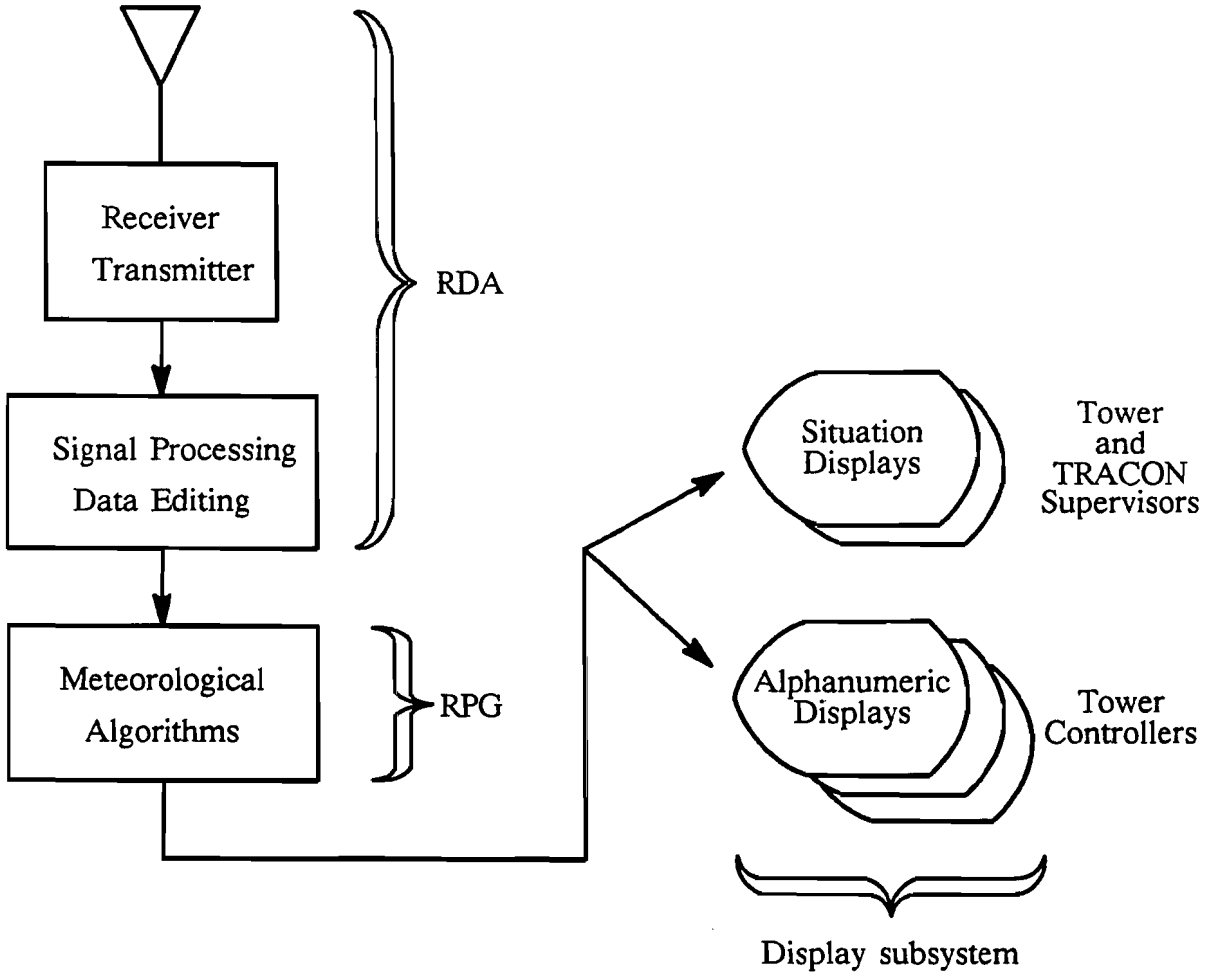

Figure 1: Block diagram of basic TDWR system components. The radar and digital processing subsystems (RDA and RPG) are located at the TDWR site, typically 15 km from the airport complex. The display subsystem is located within the airport ATC facility.

testbed system. The weather radar is a coherent pulsed-doppler system with a mechanically steered pencil beam. The TDWR production system will operate at C-band with a 0.5° antenna beamwidth. The testbed system was originally constructed at S-band with a 1.0° beamwidth (used in the Denver and Kansas City measurement programs in 1988 and 1989, respectively), but was modified to operate at the TDWR wavelength and beamwidth prior to the 1990 measurement program at Orlando.

The signal processing and basic data-editing operations are performed in the Radar Data Acquisition (RDA) subsystem. The RDA uses both sensitivity-time control (STC) and instantaneous automatic gain control (AGC) circuits to provide the wide dynamic range

necessary to prevent saturation on short-range clutter targets and heavy precipitation. A high-pass clutter filter is applied to the received signal to remove stationary ground clutter, with clutter breakthrough removed using a map of persistent ground clutter residue values. The radar pulse repetition frequency (PRF) is dynamically selected, based on the location of distant weather cells, to minimize the possibility for range aliasing of these distant echoes. Aliasing of the velocity measurements is resolved by a combination of radar waveform and data analysis techniques. Additional data processing is performed in the RDA to remove isolated moving clutter targets (e.g., airplanes and ground vehicles) and to calibrate the received intensity and velocity measurements into meteorological units.

The meteorological algorithm processing functions of the TDWR are performed in the Radar Products Generator (RPG) module. This subsystem implements the algorithms used to detect microbursts, gust fronts and significant storm regions, which are subsequently communicated to air traffic control users through the display system. Air traffic controllers are then responsible for relaying the windshear alerts to pilots over voice radio channels. The product algorithms which execute in the RPG are illustrated in Figure 2.

The TDWR radar is typically sited about 15 km away from the center of the airport complex to be protected, and it operates in a repetitive stepped-elevation scanning mode to make measurements of the reflectivity and velocity of weather systems around and above the airport complex. The scanning strategy represents a complex set of tradeoffs between the needs of the various weather processing algorithms but is required to provide an update of the surface velocity measurements over the airport complex at least once per minute.

The nominal scan strategy used by the TDWR is illustrated in Figure 3, which shows the elevation angle as a function of time for the scan sequence. At each new elevation angle the antenna scans through an azimuthal sector covering the airport complex. For most airports, this sector will be no larger than 105° in extent. Some of the scans in the sequence are full-circle scans to allow for the observation of gust fronts and storm cells over the entire airport region.

1.2. TDWR Microburst Algorithm Structure

The complete microburst detection algorithm for TDWR is made up of a number of smaller algorithm modules, as illustrated in Figure 4. The divergence regions algorithm module, operating on the velocity measurements from the surface radar scans, is the primary component of the algorithm and is responsible for the majority of the microburst detections. The remaining modules are used to identify storm features at higher altitudes in the storm cells (reflectivity, rotation, convergence and divergence aloft) and to use these storm structures to enhance the detection performance of the surface divergence algorithm. These features aloft serve primarily to reduce the thresholds required of the divergence detection function, allowing weaker divergence regions to be accepted in the presence of confirming features aloft [Campbell and Merritt, 1988]. This approach allows microbursts to be detected and warnings to be issued earlier in their lifetime (hence providing more advance warning to pilots) without incurring a significant increase in false alarms. The number of false alarms is also reduced by requiring significant storms cells aloft before issuing a microburst alarm. This test for reflectivity aloft is optionally selected as a

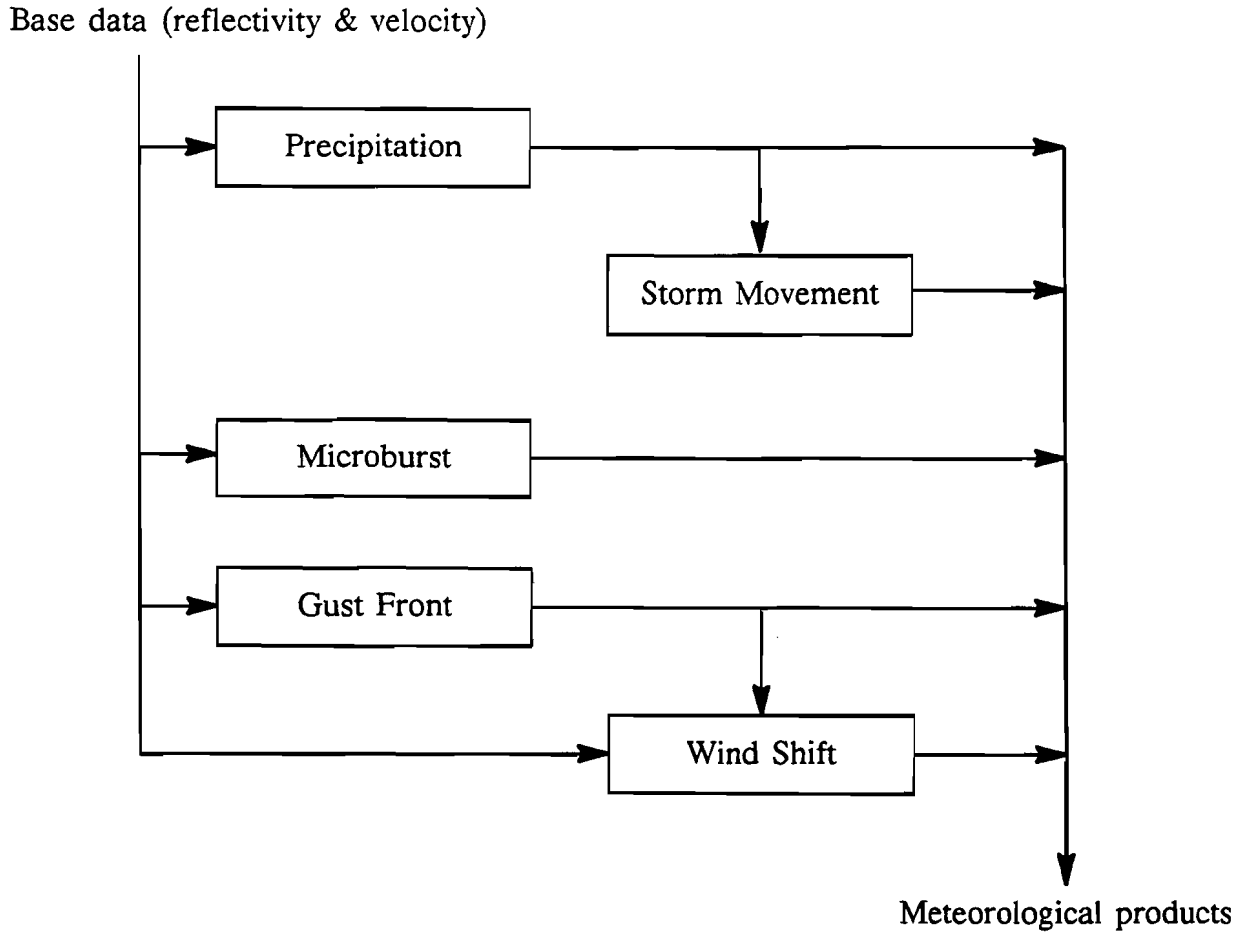

Figure 2: Block diagram of meteorological algorithm processing performed in the RPG. Base data from the RDA is supplied to the Microburst, Gust Front and Precipitation algorithms. Storm Movement and Wind Shift predictions are based on the precipitation and gust front algorithm outputs, respectively.

site–adaptable parameter for use in those meteorological environments where microbursts are reliably associated with storm cell downdrafts.

1.3. Data Quality Control in TDWR

The velocity measurements made by the TDWR may suffer from contamination in a number of ways. There are three primary sources of measurement contamination: (a) interference from ground and moving clutter targets, (b) natural chaotic variations in the windfield and (c) measurement and statistical estimator errors. The TDWR system design attempts to minimize the effects of both measurement errors and clutter interference through the use of an optimized data acquisition strategy and a series of clutter editing and decontamination algorithms. No explicit processing is applied to the base TDWR measurements to reduce the natural spatial fluctuations of the windfield, which are not

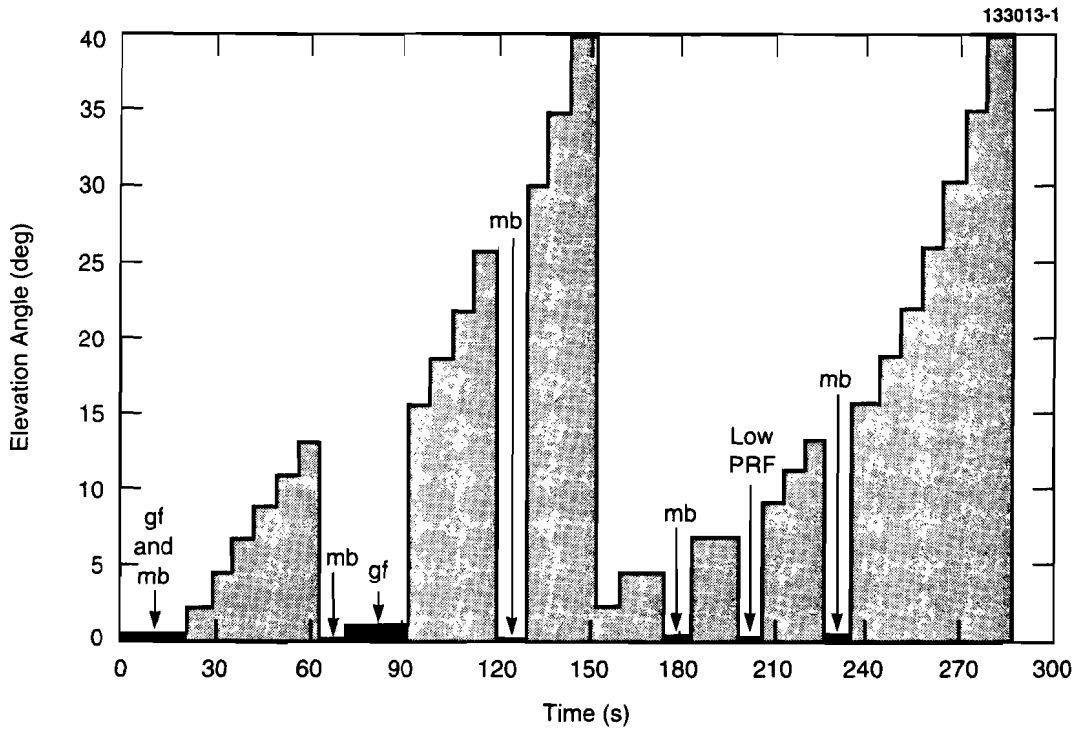

Figure 3: Antenna elevation–angle sequence used for the TDWR operational test and evaluation in 1988. The first scan in the sequence is a full–circle scan at an elevation of 0.6°. This scan is used jointly by the gust front and microburst algorithms. A second full–circle scan at an elevation of 1° is used by the gust front algorithm alone. The microburst algorithm further requires surface scans at one–minute intervals. These scans cover a 120° sector over the airport at an elevation angle of 0.4°. The PRF selection algorithm requires a long–range, low–PRF full–circle scan. The remaining scans are used by the microburst features aloft algorithms.

properly considered as "noise," though they may act to obscure the desired signatures in the data.

Stationary ground clutter (from hills, trees, buildings, etc.) is one of the dominant sources of data contamination in TDWR, and a number of powerful techniques are included in the system design to mitigate these effects [Evans and Turnbull, 1989]. A high–pass filtering operation is used to remove the bulk of the signal power near zero velocity, allowing unbiased velocity measurements to be made in all but severely contaminated regions. A time–averaged map of the clutter residue is also used to flag individual resolution cells which are still contaminated after the linear filtering stage. These filtering and editing steps leave very few velocity samples contaminated by stationary clutter, although isolated points may have been flagged as unusable by the editor.

Moving clutter targets such as birds, airplanes and automobiles will not be removed by the high–pass clutter filter or by the time–averaged clutter map (although major highways which consistently cause interference may optionally be mapped out). The TDWR includes a

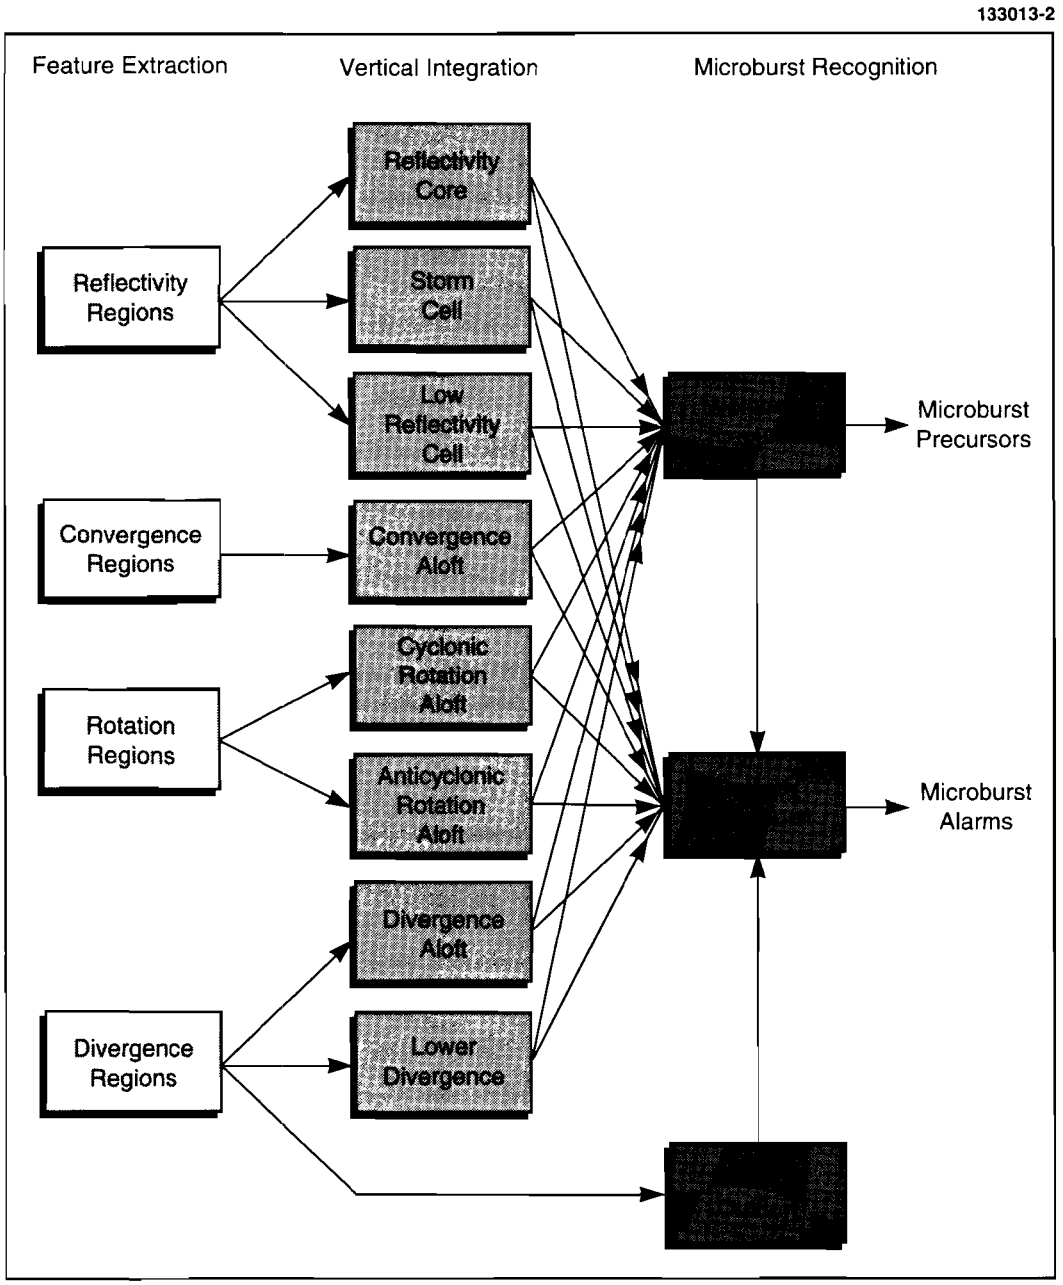

Figure 4: Block diagram of the TDWR microburst detection algorithm. The basic feature extraction modules (on the left) process the base radar measurements to identify two-dimensional regions of shear and precipitation. The middle row of modules group these two-dimensional regions vertically to form complex structures. Spatial and temporal association rules are used in the final alarm generation modules on the right to produce high-reliability alerts.

point-target rejection filter which attempts to identify these spatially small interference samples and flag them as invalid.

The velocity estimation process is stochastic and therefore includes a statistical variance dependent on the radar received signal power relative to noise and the width of the actual

velocity power spectrum in the resolution cell. The TDWR radar parameters and scanning dwell time have been chosen so as to provide a maximum velocity estimate variance of roughly 1 m/s over the range of signal power and spectrum width of primary concern. The velocity estimation process is also vulnerable to Nyquist aliasing, which is addressed through a dual-PRF scan strategy and a sophisticated data processing algorithm to place measurements into the proper Nyquist velocity interval.

Range aliasing may also cause contamination of the velocity field when distant storms are aliased into the airport coverage region. The TDWR includes an adaptive algorithm which identifies the location of distant storms and chooses the PRF values which minimize contamination of measurements in the critical airport coverage regions. Additional data editing is performed to invalidate those resolution cells which are contaminated by multiple-trip returns.

The careful selection of radar parameters, scanning strategy and signal/data processing techniques has provided the TDWR with a considerable ability to reject the major sources of interference and to maintain a moderate level of estimator variance. As a result, the base velocity measurements provided by the TDWR to the divergence detection algorithm may generally be regarded as "clean" with respect to these sources of error.

1.4. Operational Measurement and Evaluation Programs

The TDWR testbed radar system has been the primary data source for the development and evaluation of the TDWR algorithm suite and for the operational evaluation of the system by ATC users [Turnbull et al., 1989]. The testbed system was initially deployed in Memphis, TN and was fully operational in 1985. Measurement and evaluation programs have been conducted annually since 1985 in a number of geographical and meteorological environments, as listed in Table 1. The radar measurements of microbursts obtained during these field programs have been integral to the development of the microburst algorithm and have served as the basis for evaluating the algorithm performance.

Table 1:
Field measurement and evaluation programs conducted with the TDWR testbed radar system*.

Year	Location	Products Delivered to ATC users?
1985	Memphis, TN	No
1986	Huntsville, AL	No
1987	Denver, CO	No
1988	Denver, CO	Yes
1989	Kansas City, MO	Yes
1990	Orlando, FL	Yes

*Radar measurements of microbursts and other significant weather events were recorded during each program and used for algorithm development and refinement. Operational evaluations have been conducted annually since 1988, in which TDWR products were provided live (in real time) to air traffic users.

The field program in Denver during 1988 marked the first live demonstration of the TDWR system to on-duty air traffic controllers and pilots and was conducted as a formal Operational Test and Evaluation (OT&E) of the TDWR system. Based on the success of this OT&E, the FAA has awarded the production contract for 47 TDWR systems, which will begin operational deployment in the early 1990's. Additional operational demonstrations conducted with the testbed have further validated its effectiveness in various meteorological, ground clutter and air traffic environments [Evans, 1990].

During each measurement and evaluation program, the testbed system is routinely operated each afternoon and scans using the operational scan strategy appropriate to the specific airport being protected. The microburst, gust front and precipitation products are computed in real-time and provided to ATC users and pilots. Local controllers in the tower are provided with alphanumeric displays indicating microburst and/or gust front windshear alarms over the active runway and arrival/departure corridors. These alerts are read directly over voice radio channels to pilots on approach or departure. The tower supervisor and TRACON controllers are provided with a color geographical situation display (GSD), indicating the location and extent of microbursts, gust fronts and storm precipitation regions. Gust front forecast locations, based on propagating the observed gust front motion into the future, are also provided on the GSD.

During the measurement programs all base radar observations are routinely recorded, and all TDWR system alarms and products are archived for subsequent analysis. Human observers monitor the weather situation and visually monitor the radar measurements for microbursts. Detailed logs are kept by the real-time observers to note the presence of microbursts for post-mission analysis.

A number of support sensors are typically deployed along with the TDWR testbed radar, including surface mesoscale network, instrumentation recordings of the airport low-level windshear alerts system (LLWAS) sensors, and a second doppler radar operated by the University of North Dakota. These support sensors provide verification of the windshear events observed by the TDWR testbed radar and allow detailed off-line analysis of selected weather events, including dual-doppler windfield analysis for microburst outflow studies.

1.5. Organization of the Remainder of the Report

Chapter 2 presents a conceptual overview of the microburst detection process using a conceptual model for a microburst outflow to examine important aspects of the detection problem. The divergence algorithm is described in detail in Chapter 3, and the formal documentation of the algorithm, as documented for the TDWR system specification, is given in Appendix 3. Chapter 4 illustrates the behavior of the divergence algorithm on a strong microburst from July 11, 1988 which has been the subject of considerable study by the windshear community. The algorithm is seen to perform fairly well on this microburst, although considerable variations in the exact detected region are seen from one scan to the next. To provide a statistical summary of the algorithm performance, the detection and false-alarm probabilities for a number of cases are presented in Chapter 5 and are compared with the corresponding quantities which have been determined for the complete microburst detection algorithm.

These scoring statistics are based solely on single-doppler radar observations, and the impact of microburst asymmetry on these performance metrics is considered in Chapter 6. The conclusion here is that asymmetry can substantially reduce the effective performance of the microburst algorithm and that improved detection rates for events with weak signatures is needed. Chapter 7 examines some possible alternative divergence algorithms based on computational methods (rather than complex decision tests) to improve the performance of the detection process. The results obtained in limited testing on these alternative techniques indicate promise and suggest avenues for further examination as summarized in Chapter 8.

2. CONCEPTUAL MODEL FOR THE DETECTION PROCESS

Many of the basic concepts employed by the TDWR divergence detection algorithm may be examined in the context of a simple idealized microburst outflow model. The analysis below considers the case of a radially symmetric outflow from a pure divergent source as shown in Figure 5. The radar-measured velocity signature for this type of outflow is computed, and the basic size and strength properties of the signature are determined as a function of microburst size, strength and distance from the radar.

2.1. Microburst Windfield Model

The velocity field for the pure divergent source model is circularly symmetric, and the windspeed at a distance r_c from the center of the outflow is described by the profile $V_m(r_c)$ and is always directed radially outwards. The exact form of this velocity profile is not critical to the observations to be made in this section, and there are several profiles which have been used in previous studies. The basic assumption made regarding this velocity profile is that the windspeed increases monotonically from the outflow center to some peak speed and then decays back towards zero at further ranges. The peak windspeed along the velocity profile is denoted as V_p, and the range (from the outflow center) at which this peak occurs is labelled R_m.

To obtain quantitative results, a specific velocity profile must be chosen for analysis. A commonly-used profile is the sinusoidal model where the outflow signature is modelled as a half-cycle of a sinusoid. This model was used to examine detection algorithm performance for the ASR-9 system in a report by Noyes, 1990. This report presented comparisons

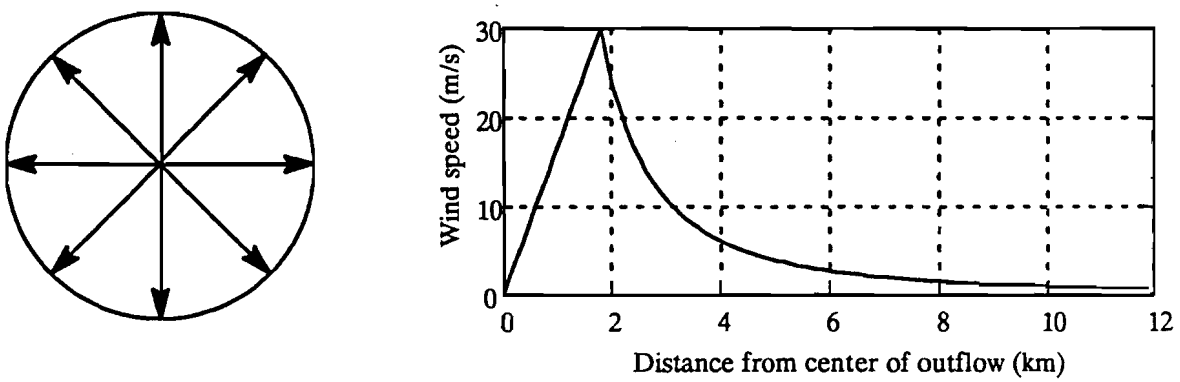

Figure 5: Conceptual microburst outflow model. (a) Windfield is circularly symmetric, with winds directed radially outward. (b) Wind speed profile as a function of distance from the center. Speed increases linearly to a maximum value then drops off at a rate of 1/r2.

between the model winds and actual measured signatures, showing good agreement. This sinusoidal model was also used to examine aircraft performance measures in [Elmore, 1989]. The second model form was described in [Hjelmfelt, 1988] and models the outflow winds as linearly increasing from the outflow center then dropping off as the square of the distance to the center. This model was obtained by comparing the profiles from several microbursts with the profile from a laboratory model of a wall jet. Figure 6 shows the normalized velocity profiles examined by Hjelmfelt, along with the wall jet profile. For both the wall jet and the observed microbursts, the windspeed increases linearly from the center to some maximum windspeed V_P (at range R_m from the center) and then drops off rapidly back to zero. The wall jet drops off at r^{-1}, but the observed microburst profiles drop off more rapidly, roughly proportional to r^{-2}.

The sinusoid and r^{-2} models are compared in Figure 7 at two different spatial scales. These examples show that the sinusoid model provides a smoother transition at the velocity peaks and a more abrupt distinction between the edge of the outflow and the ambient wind.

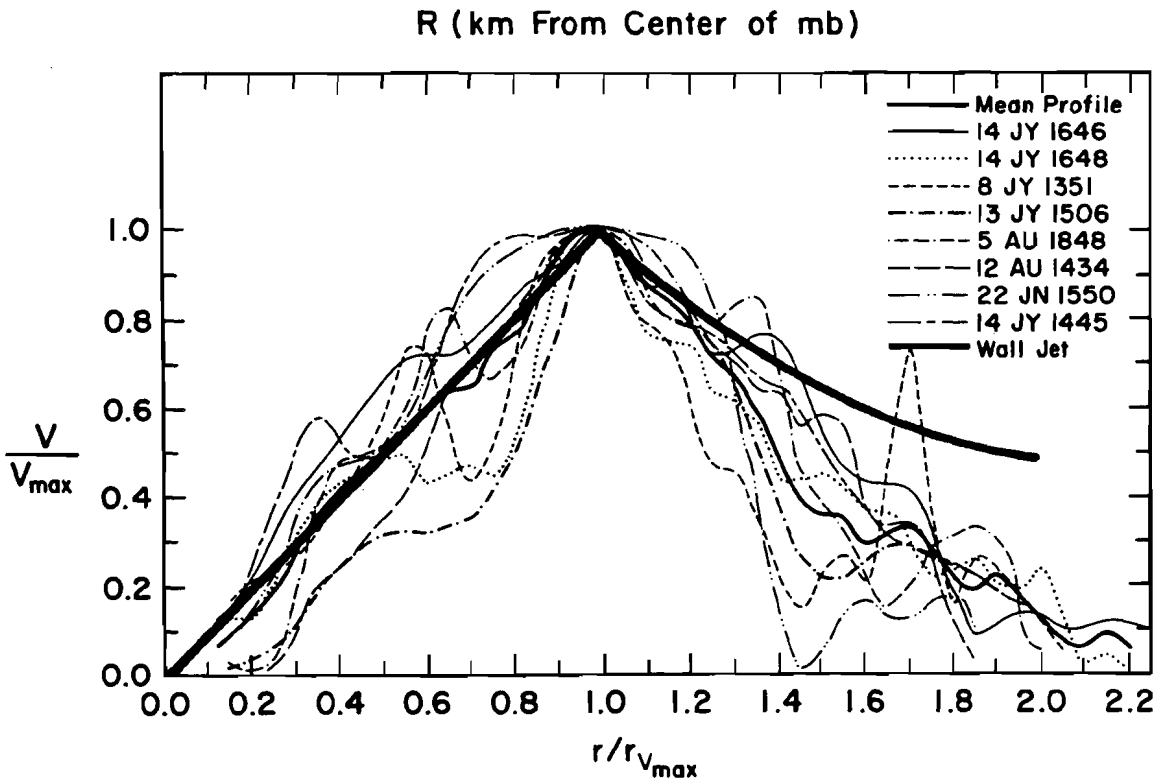

Figure 6: Microburst windspeed profiles from eight microbursts. Speeds are normalized to maximum speed and ranges are normalized to the range at which the maximum speed occurs. The solid line is the mean observed microburst profile while the heavy line is the profile for laboratory wall jet models. The wall jet model drops off at a rate of 1/r; the observed profiles drop off at a rate of approximately 1/r2. Taken from [Hjelmfelt, 1988]

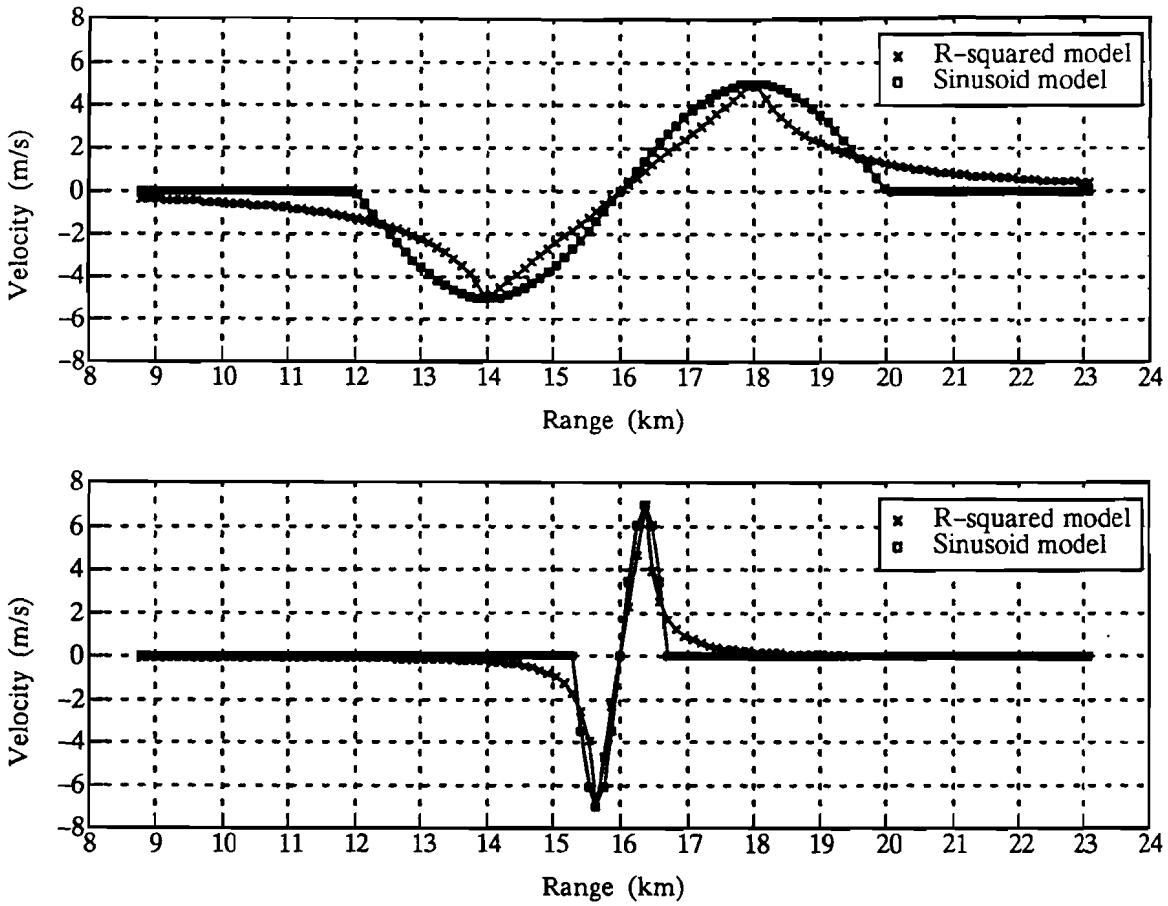

Figure 7: Comparison between sinusoid and R-squared models for two different size microbursts. The top plot shows a large, weak microburst with a 10 m/s velocity change over 4 km distance. The lower plot is a very small, strong microburst with a 14 m/s change over 0.72 km. These two examples represent the extremes of the size and strength domains for the detection process.

The r^{-2} model has very sharp discontinuities at the peaks but a more gradual and extended transition at the edges of the event.

To understand the behavior of the different smoothing methods, their effects will be considered on modelled microbursts at both small and large spatial scales. A very small, yet strong, microburst is shown in Figure 8. This profile was obtained from a very severe microburst which occurred on July 11, 1988 during the TDWR Operational Test and Evaluation (OT&E). This profile is from the first radar surface scan on which the microburst was visible. The figure shows the microburst has an initial strength of roughly 14 m/s, but across a distance of less than 750 meters. This combination of strength and small size is quite unusual and probably represents a good lower bound on the detection requirements for microbursts. Although this event is unusual, it is also very important; several aircraft actually penetrated this microburst, and the timely warnings produced by the TDWR testbed

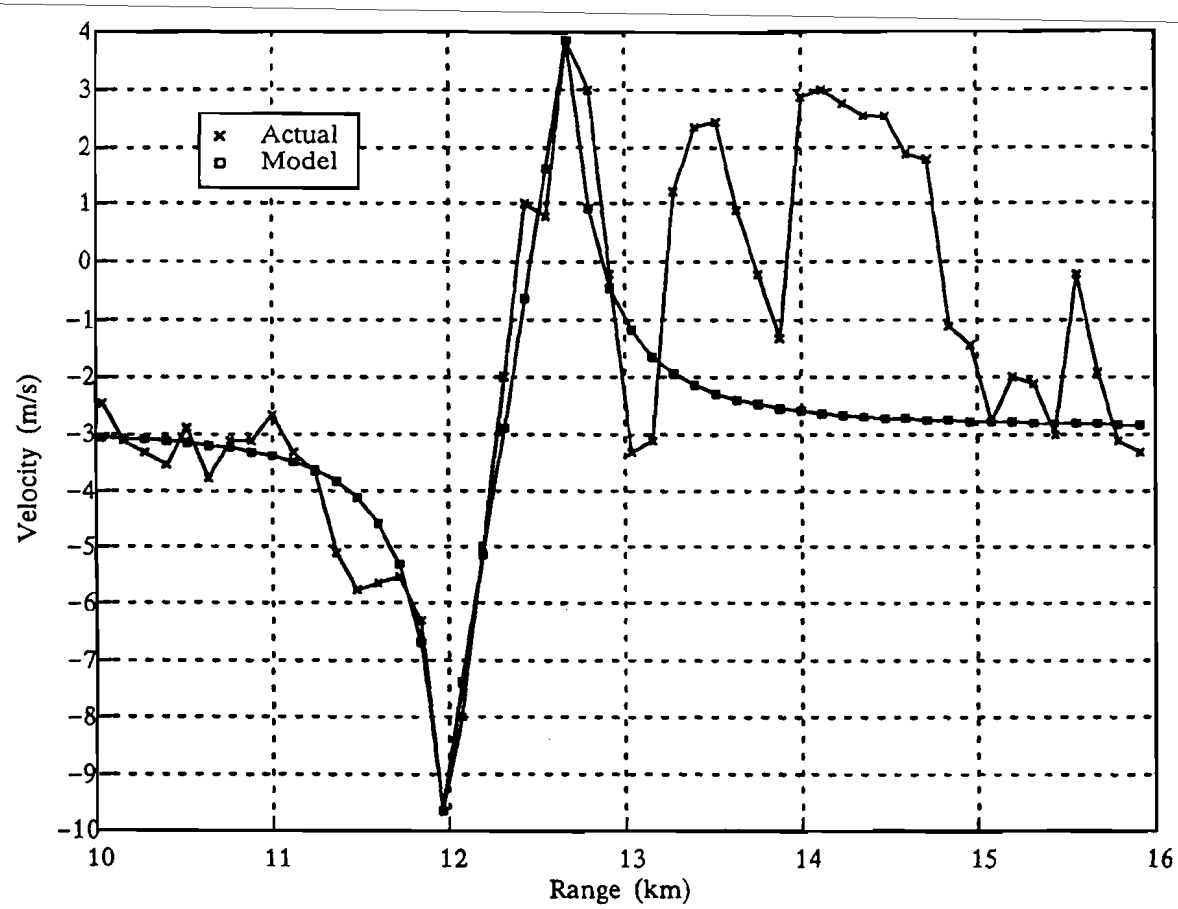

Figure 8: Comparison between measured profile of a small microburst and the profile generated by the windfield model. Microburst event is from July 11, 1988 at 22:04:01 (297.5° azimuth angle).

may well have prevented a disaster [Schlickenmaier, 1989]. Also shown in Figure 8 is the R-squared microburst model which best matches the measured velocity profile. The model fits the data quite well from ranges 10 km to 13 km, particularly along the strong shear portion. At further ranges the microburst signature is complicated by the presence of an adjacent thunderstorm outflow (extending from 14 km to roughly 18 km) not included in the microburst model. The close match between the R-squared model and the forward edge and shear region of this event lend considerable credence to the use of this model for investigating microbursts during their earliest stages of development. The sinusoidal model is perhaps more appropriate for describing more mature microbursts as described in [Noyes, 1990].

The Hjelmfelt model will be used for the analysis in this chapter, so the radial profile function is given by:

$$V_m(r_c) = V_* \frac{r_c}{R_m} \quad for\ r_c < R_m \tag{1}$$

$$= V_* \left[\frac{R_m}{r_c}\right]^2 \quad for\ r_c \geq R_m$$

The locus of points at which the windspeed is at its maximum value (a circle of radius R_m) is typically considered to be the "outline" of the microburst region. The doppler velocities measured by a radar will consist of only the component of the horizontal outflow wind in the direction of the radar beam. For a radar located at distance D_m from the center of the outflow, the radial velocity at a range r and angle θ (between the radar beam and the outflow center) may be expressed as:

$$V_r(r,\theta) = V_m(r_c)K \tag{2}$$

where:

$$K = \frac{r - D_m \cos(\theta)}{\sqrt{D_m{}^2 + r^2 - 2rD_m \cos(\theta)}} \tag{3}$$

$$r_c = \sqrt{D_m{}^2 + r^2 - 2rD_m \cos(\theta)} \tag{4}$$

An example of the radar signature obtained from this model is shown in Figure 9. The color image shows the radial velocity field (in meters per second) for a model outflow at range 12 km, with a 1.5 km radius and maximum windspeed of 15 m/s. This figure also shows a plot of the radial velocity profiles along several radials through the outflow. These profiles show a) the sharp shear pattern when the radial intersects the outflow region and b) the gradually reduced shear signature as the viewing angle moves further away from the outflow center.

The goal of the TDWR divergence detection process is to identify the shear region on each radial of measurements extending from negative peak to positive peak. Each radial intersecting the shear region results in a single shear "segment" being detected; the collection of shear segments for a sample microburst model is depicted in Figure 10. The divergence algorithm then groups these segments together to form an outline of the detected shear region. Each shear segment identified by the divergence algorithm is characterized by a length and a strength (the difference in velocity across the segment). Thresholds on the length and strength values are used to discard false segments, and a minimum number of segments will be required in a group to detect a shear region.

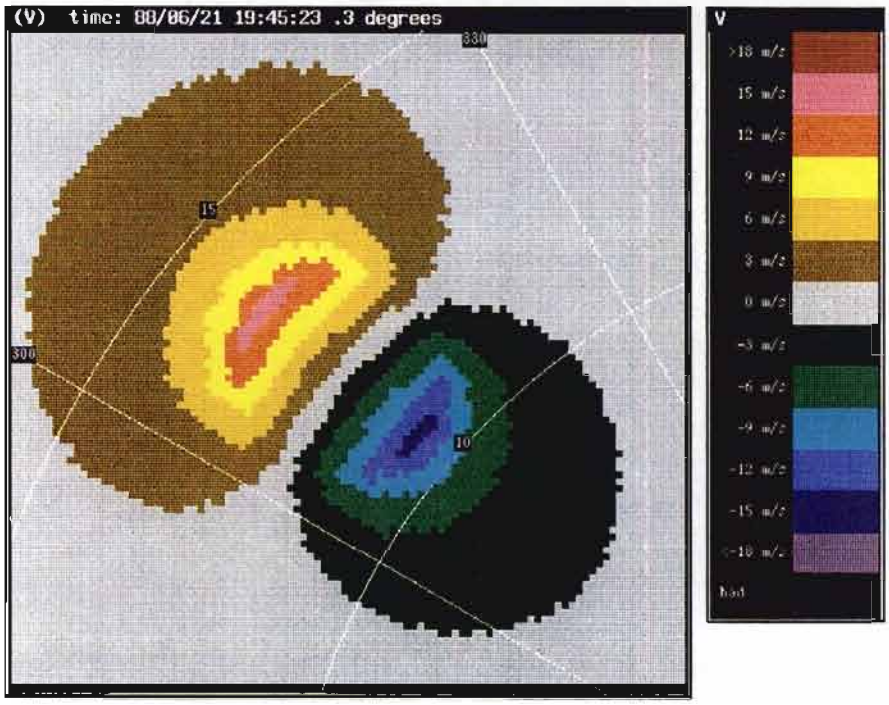

(a)

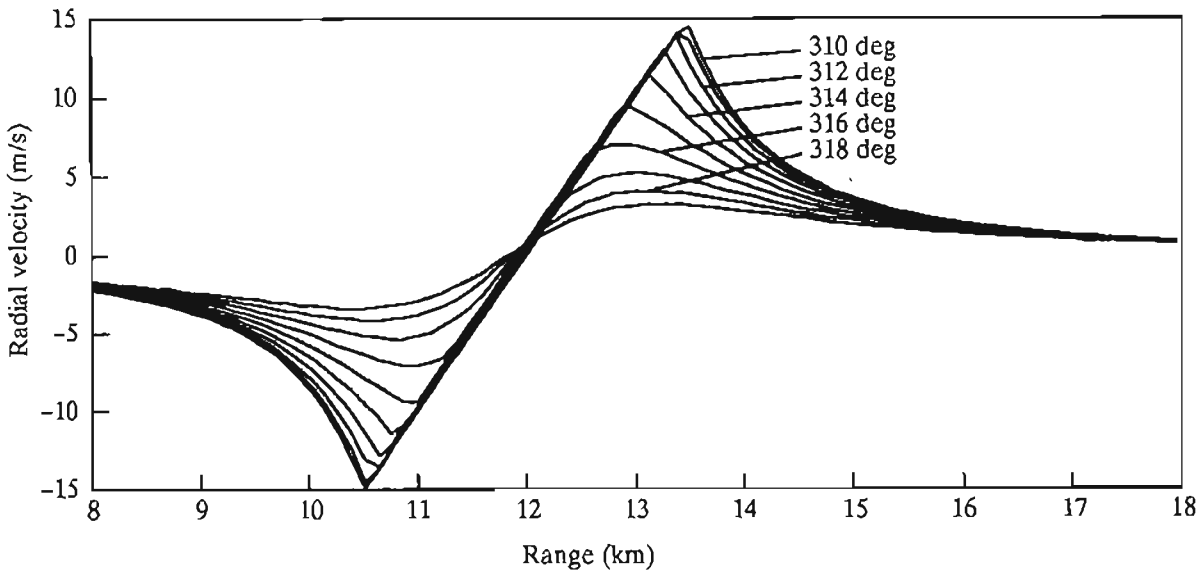

(b)

Figure 9: Radial velocity field for a model microburst. (a) Radial velocities in meters per second for microburst at range 12 km, radius 1.5 km, and peak windspeed of 15 m/s. (b) Profiles along radials through the outflow.

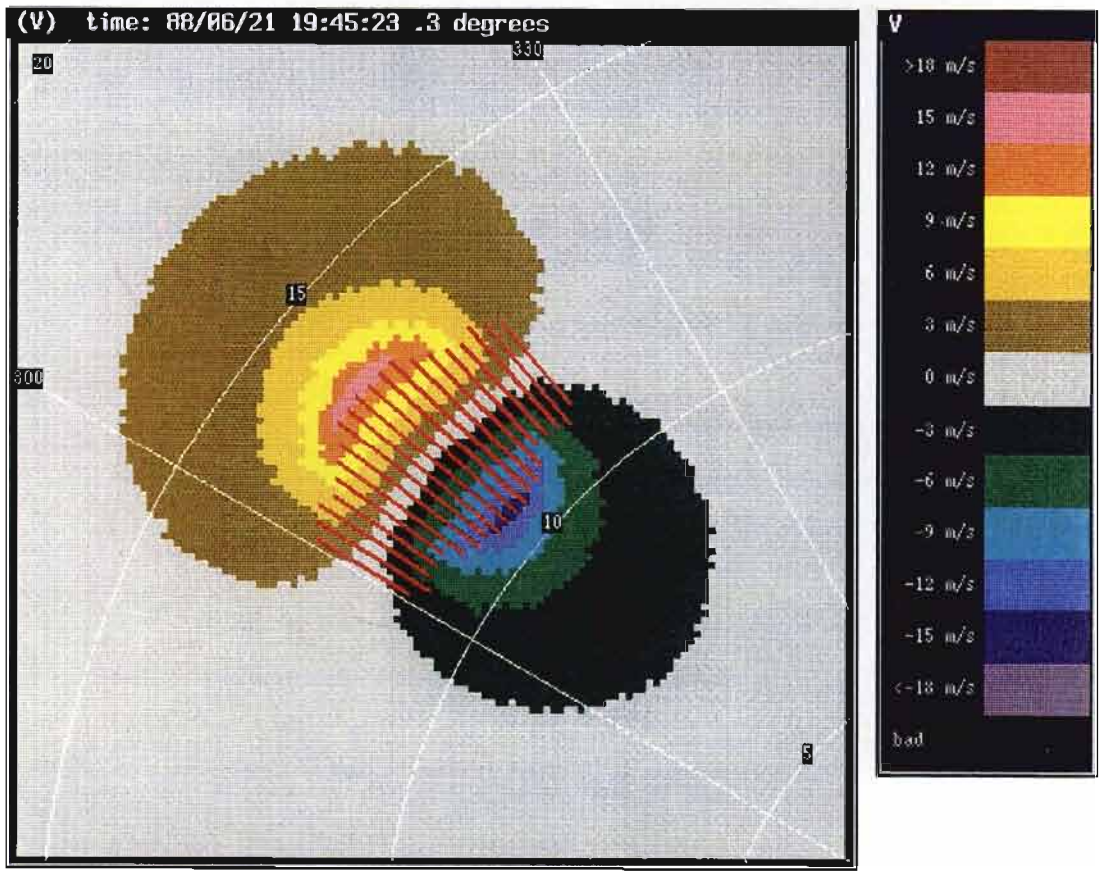

Figure 10: Shear segments detected for model microburst. Segments are grouped together based on range overlap and azimuth proximity to form a two-dimensional region outlining the microburst outflow.

The length and strength of the shear segments found for an outflow will vary with the radius of the outflow, the distance to the outflow center, the strength of the outflow, and the angle between the radar viewing direction and the center of the outflow. These variations are important to the proper selection of thresholds for the detection algorithm and may be examined in the context of this simple outflow model. For a radial which passes through the outline of the microburst, the peak radial velocities will be observed at the intersection of the radar radial and the microburst outline. The velocity difference across this shear segment

and its length are:

$$\Delta V = 2V_* \left[1 - \left(\frac{D_m}{R_m}\right)^2 \sin^2(\theta)\right]^{\frac{1}{2}} \tag{5}$$

$$\Delta R = 2R_m \left[1 - \left(\frac{D_m}{R_m}\right)^2 \sin^2(\theta)\right]^{\frac{1}{2}} \tag{6}$$

The length and strength of the segments are both reduced by the same factor as the viewing angle from the outflow center increases. This reduction factor is plotted in Figure 11 for a microburst of 2 km diameter at several angle offsets.

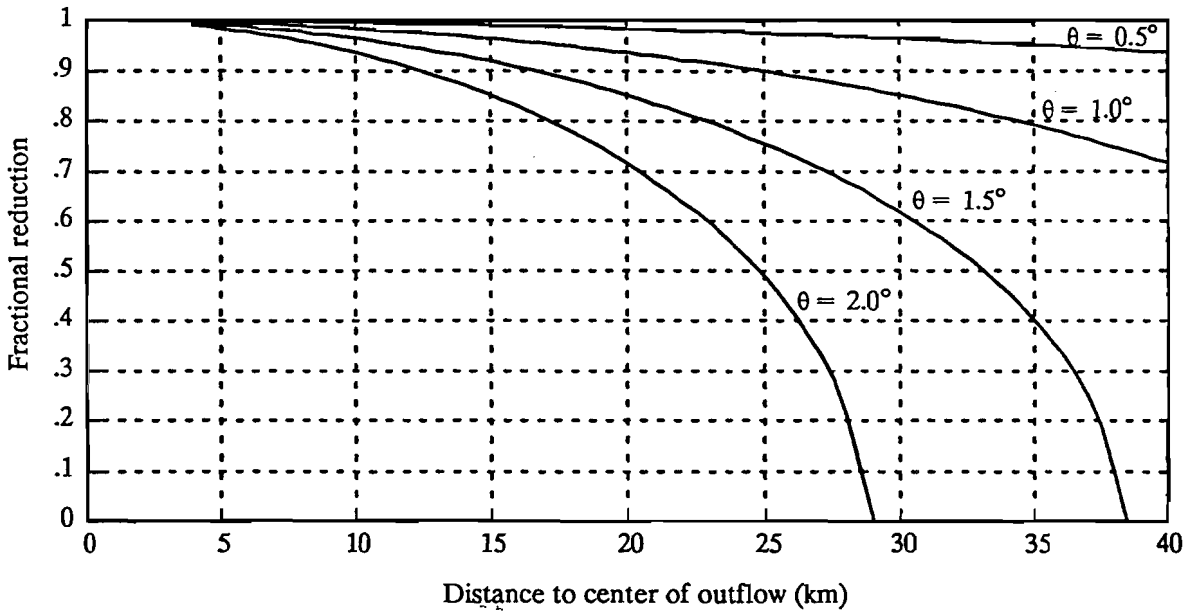

Figure 11: Reduction in shear segments length and strength when radar viewing angle is not centered on the outflow center. Microburst radius is 1 km.

The validation tests for shear segments and shear regions require each region to contain at least N_T segments and require each segment to be at least R_T km long and have a velocity differential of at least V_T m/s. Under these constraints, the conditions for detection are:

$$D_m \leq R_m \frac{\left[1-\left(\frac{V_T}{2V_s}\right)^2\right]^{\frac{1}{2}}}{\sin(\theta_{max})} \tag{7}$$

$$D_m \leq R_m \frac{\left[1-\left(\frac{R_T}{2V_s}\right)^2\right]^{\frac{1}{2}}}{\sin(\theta_{max})} \tag{8}$$

where:

$$\theta_{max} = \frac{N_T}{2} \ degrees \tag{9}$$

These two conditions simply express the requirement that the thresholds be met on each of the N_T segments obtained as the radar scans across the outline of the event. The selection of a maximum offset angle (corresponding to the weakest of the segments) is chosen, assuming a 1° azimuthal sampling interval, and is based on the worst–case alignment of the outflow center with the specific azimuth sample angles. The nominal parameter values selected for the TDWR divergence algorithm are:

$$V_T = 5 \text{ m/s} \tag{10}$$

$$R_T = 0.95 \text{ km} \tag{11}$$

$$N_T = 2 \text{ segments} \tag{12}$$

These parameter values were determined by examining a large number of cases and computing the performance statistics for each of several parameter value combinations. The number of false alarms generated by the TDWR divergence algorithm (described in detail in the next chapter) is very sensitive to the values of the above parameters. Decreasing any of these parameter values (in order to detect smaller or weaker microbursts) will result in a corresponding increase in the number of false alarms. The choice of these parameters was based on data cases observed in the Denver, CO area and may not be optimal for other geographic locations. The TDWR allows these parameters to be tuned for optimum performance at each installation, providing some degree of flexibility to match the variations in microburst characteristics across the country.

Given these nominal parameter values, the detection criteria of (7) - (9) are essentially identical and simplify to:

$$D_m < 33\ R_m \tag{13}$$

Hence, these parameter values allow detection of a 1 km radius microburst out to a range of roughly 33 km. This capability closely matches the minimum detection requirements set forth for the TDWR microburst algorithm.

2.2. Shear Region Extension Outside Outflow Boundary

For a simple microburst outflow, such as that described by the conceptual model above, the obvious "boundary" for the event is the circle corresponding to the peak outflow windspeed. This boundary, however, does not indicate the extent of the shear region which would be encountered by aircraft flying near the event.

For an aircraft penetration of the event which passes on a straight line through the outflow center, the circular outline accurately represents the extent of the shear region which would be encountered by the aircraft. A penetration along such a path results in maximum exposure to the divergent shear, and (if the shear is strong relative to the response capability of the aircraft) the peak-to-peak windspeed change across the event is indicative of the airspeed loss which the aircraft might experience. For an encounter path which passes through the outline but not through the center, the outline still indicates the extent of the shear region, but the total windspeed changes is less than that for a path directly through the outflow center. This aspect of the microburst outflow penetrations is generally understood in the TDWR community, and the system alert strategy and user training information attempt to account for the variability in perceived strength, depending on penetration path.

An observation (originally from [Campbell, 1990]) which is less well recognized is that paths near, but outside, the conventional "outline" of the microburst also contain divergent shear, albeit at a significantly reduced strength. Referring to Figure 12, it is apparent that so long as the winds on the periphery of the outline are radially directed, the longitudinal windspeed component along a straight path will change sign (indicating shear) at the point of closest approach to the microburst outline. The strength of this shear will clearly depend on the strength of the microburst, the rate of decay of the windspeed outside the event outline, and the distance between the path and the outline.

Since the divergence detection algorithm operates by finding shear segments along radials from the radar, the set of segments detected for a perfectly circular and symmetric microburst will not form a circle. The segments shown in Figure 10 illustrate this condition. Hence, for circular outflows, the shear region detected by the divergence algorithm will be elongated in the cross-range dimension and will accurately represent the locations of shear which would be experienced by aircraft flying in the direction of the radar viewing angle.

An important consequence of this observation is that aircraft flying perpendicular to the radar viewing angle at ranges just short of or just beyond the extent of the detected shear

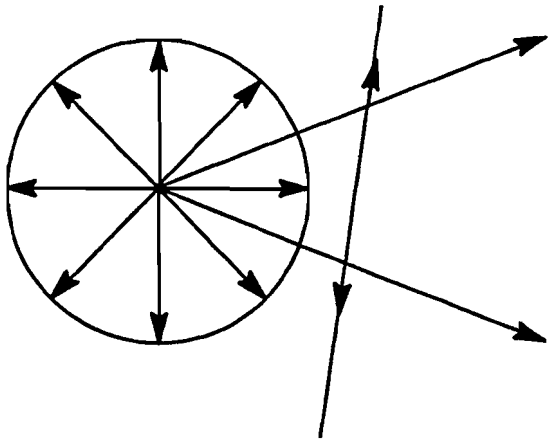

Figure 12: For a flight path outside the microburst outline, the longitudinal component of the wind will change from a headwind to a tailwind (divergent shear) as long as the winds are radially directed from the outflow center. Strength of the shear will depend on the actual rate of decay of the winds outside the event.

segments may experience significant windshear. The aspect-dependent nature of the longitudinal shear within an outflow region therefore causes "blind spots" in the radar-measured shear region, even in perfectly symmetric outflows. This factor should be taken into consideration in the interpretation of the TDWR outputs (i.e., detected boundaries may underestimate or overestimate the extent of the shear region depending on viewing angle relative to flightpath) and in the development of future detection and/or display algorithms.

2.3. Detection Issues for Realistic Signatures

This conceptual model is useful for understanding the basic geometry of microburst outflows, but it is a gross simplification of the actual signatures obtained from radar measurements. The radar images in Figure 13 are typical for an isolated microburst in the Denver, CO area. The top image in this figure is the reflectivity measurement, in units of dBz, and the lower image is the radial component of the horizontal wind velocity in m/s. The radar is located off the lower right corner of the image, as indicated by the range rings (spaced every 5 km) and azimuth lines. The white overlay near the center of the image represents the runways at Denver's Stapleton airport, where these measurements were obtained. This microburst has a peak windspeed of roughly 13 m/s and a radius of 3 km. The velocity field for the conceptual model using the same parameters is compared to the actual measurements in Figure 14. This comparison clearly demonstrates the noisy background accompanying actual signatures and the departure of the actual shear profiles from the simple symmetric model form.

The divergence algorithm, described in the following chapter, follows the basic conceptual approach of locating shear segments within the outflow region and clustering

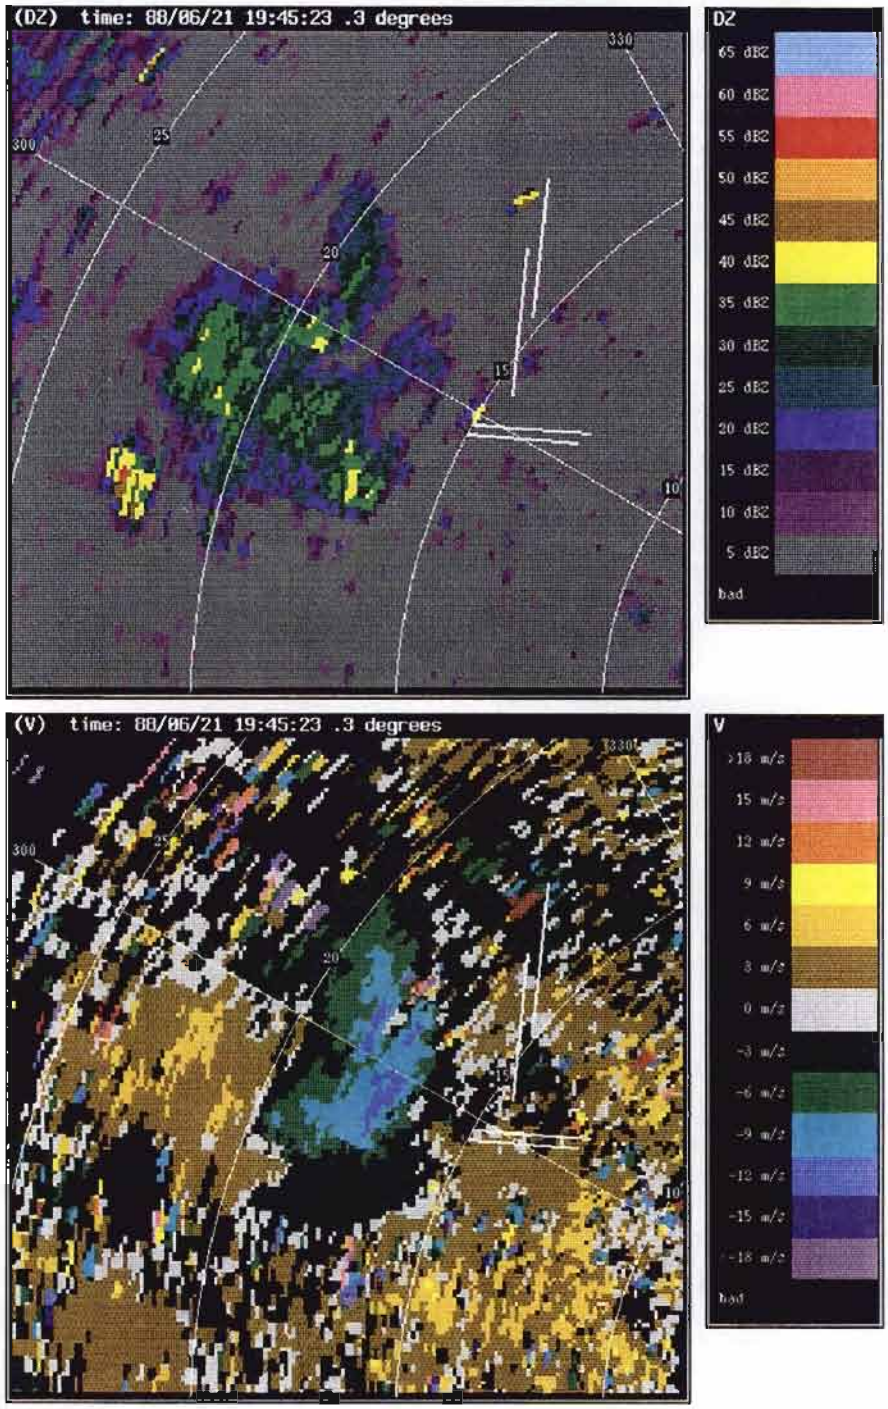

Figure 13: Example of a microburst signature as measured by the TDWR testbed radar. Top image is reflectivity in dBz; lower image is velocity in m/s.

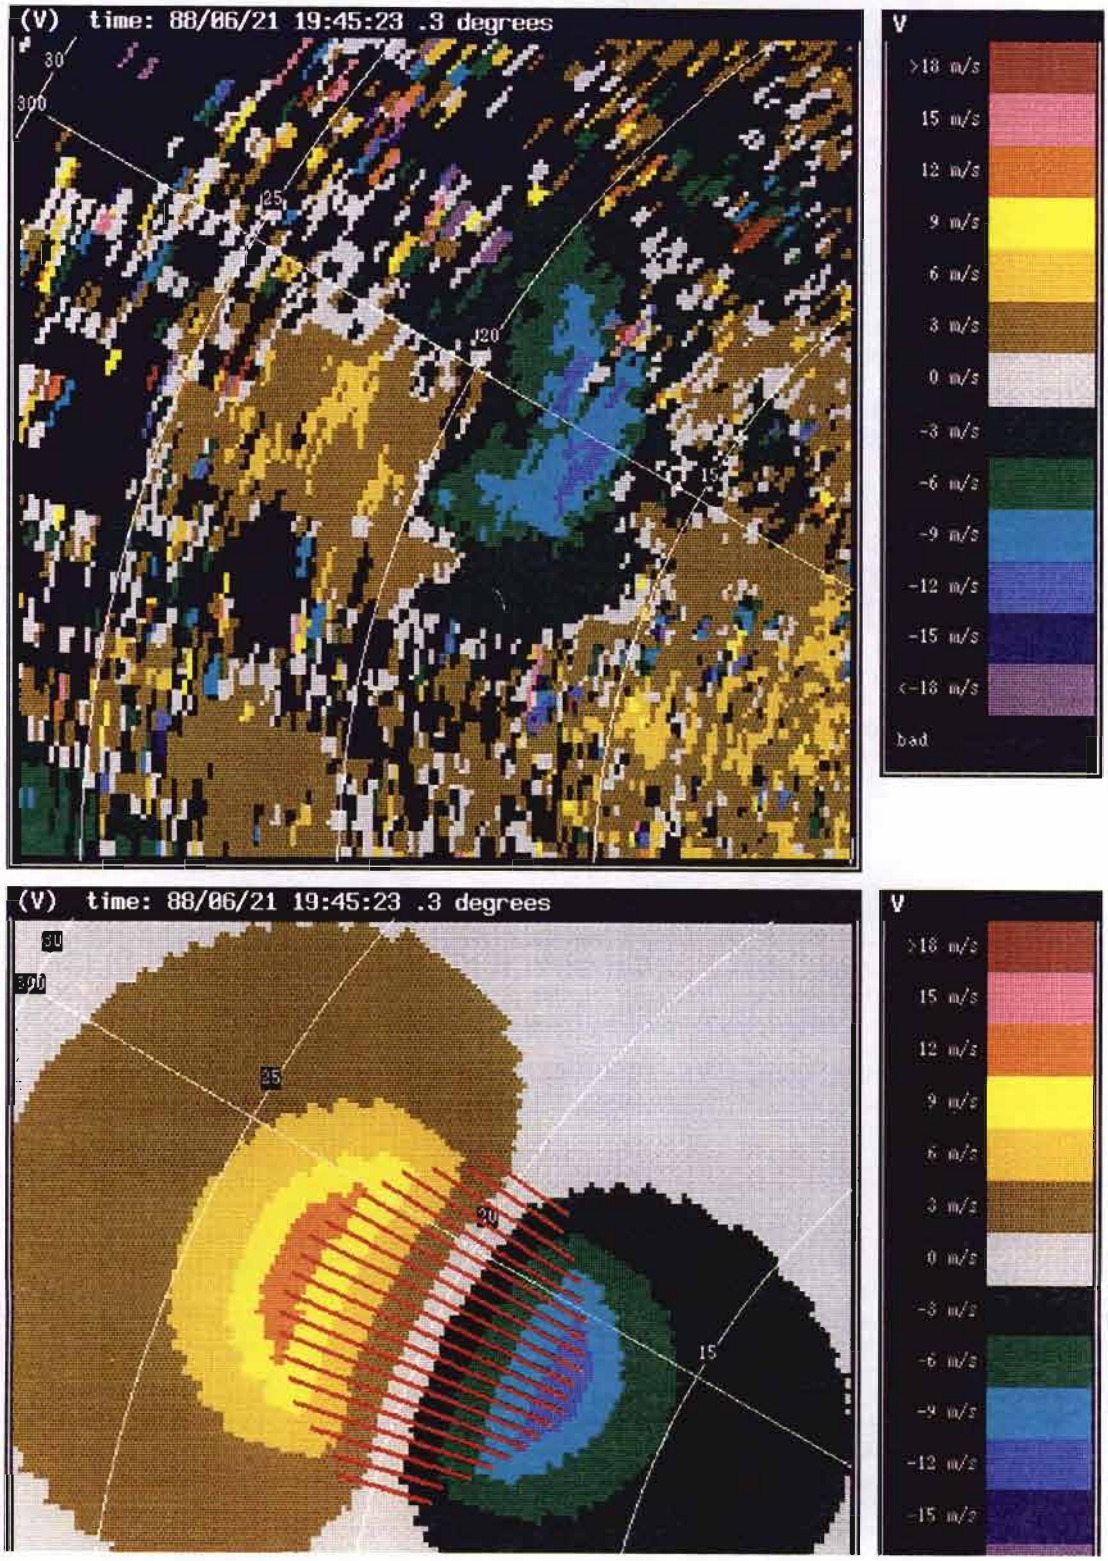

Figure 14: Comparison between actual microburst outflow measurements (top) and conceptual model field (bottom) for the same parameters.

them together based on proximity in range and azimuth. However, to provide robust performance in the presence of background noise and interference and to detect distorted and non-symmetric shear regions, the algorithm includes a number of complex tests used to distinguish true shear regions from other similar signatures.

3. DESCRIPTION OF TDWR DIVERGENCE ALGORITHM

3.1. Algorithm Specification for TDWR

The divergence algorithm is formally described for TDWR using an English-like procedural language known as Algorithm Enunciation Language (AEL). The complete formal algorithm description is presented in ATC-145. All numerical values used in the algorithm description are configured as site-adaptable parameters, meaning that they may be adjusted for optimum system performance at each TDWR radar site. The description of the algorithm provided here will use many of the same names as the AEL description, particularly in reference to site-adaptable thresholds.

3.2. Inputs and Outputs

The input to the divergence algorithm is the radial velocity data field for a single-radar tilt (a radar scan in azimuth at a constant elevation angle). The velocity field contains the radial velocity measurement at each range gate for a set of azimuths. The TDWR provides radials spaced at 1° intervals in azimuth. The velocity measurements used by the algorithm are assumed to have been preprocessed by a variety of data quality checks to remove the effects of ground clutter residue, point target interference, velocity aliasing, and low signal power, as required by the TDWR specification. Any input velocity values found to be contaminated by these quality checks are tagged as invalid by these data quality filters. The algorithm makes explicit reference to the handling of points with invalid velocity values.

The algorithm generates two sets of information as output for each tilt: a set of shear segments and a set of shear regions. The segments are the primitive one-dimensional portions of each radial found to contain divergent shear. The regions are two-dimensional areas formed by clustering the segments which meet all the thresholding criteria. Each segment is described by an azimuth angle, start and stop range, and velocity difference across the segment. Each region is described by a large number of characteristics, particularly the bounding box (minimum and maximum X and Y coordinates bounding the region) and maximum velocity differential found within the region.

The divergence regions are used by the subsequent microburst algorithm stages (which apply time continuity and strength constraints in conjunction with features aloft) to form microburst alarm regions. For these algorithm stages the region extent is described solely by the bounding box and centroid information. The shear segments are used again in the final stage of the algorithm where microburst shapes are computed for each alarm. These shapes are based on the actual locations of shear segments in the region(s) rather than on the more limited bounding box information.

3.3. Site Adaptable Parameters

The algorithm makes use of numerous thresholds and test criteria which may be adjusted at individual radar sites to provide optimum performance. These numerical thresholds are

documented as site-adaptable parameters and are the primary mechanism for tuning the performance of the algorithm to adapt to different site characteristics. The full set of site-adaptable parameters is listed in the AEL document, along with a table of nominal values used during testing.

3.4. Shear Segment Identification

The first stage of the divergence feature extraction algorithm is the identification of one-dimensional shear segments along individual radials of velocity measurements. Each segment is meant to identify a portion of the radial which contains velocity measurements exhibiting a generally increasing trend with range (i.e., which have a positive divergence).

The segment detection process operates on one radial at a time and sequentially examines each range gate along the radial. For each range gate examined, the algorithm forms an observation "window" consisting of the NUMBER(Window) gates which follow the current window. Various tests are applied to the velocity values in the window to determine whether a shear segment should be started (or terminated) at the current range gate. The window length is selected based on the range gate spacing of the radar, to correspond to roughly 0.5 kilometers distance. This window size was chosen experimentally to balance the need for locating small shear segments with the desire to filter out the small-scale fluctuations of the velocity measurements. The concept of the shear segment search window is illustrated in Figure 15.

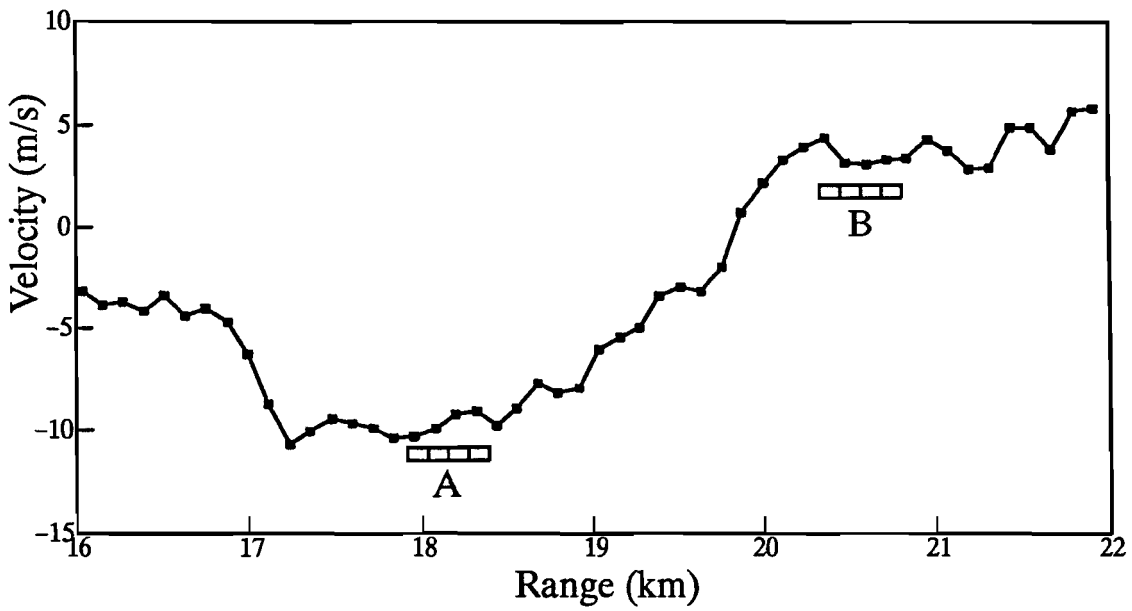

Figure 15: Example of the segment search window. The window consists of NUMBER(Window) range samples, starting at the current sample. The pattern of velocity measurements in the window determines if a segment is to be started or terminated. The window marked "A" is an example of a valid segment starting condition, while window "B" is the corresponding termination point.

As the window is moved out in range, a segment is started at the current gate when the velocity values in the window meet the following criteria:

i) All velocity values in the window are valid and

ii) All velocity values in the window are greater than the velocity value in the current gate, and

iii) The first NUMBER(Rise) velocities in the window form a monotonically increasing sequence

These conditions amount to a strict requirement for an increasing trend of velocities to begin a segment. Once a segment has been started, the window is moved out further in range until the following segment termination criteria are met:

i) The total number of velocity values in the window, which are either invalid or less than or equal to the velocity value in the current gate, exceeds NUMBER(Bad), or

ii) The velocity difference between the current gate and the value at that point with the smallest velocity greater than that at the current gate exceeds THRESHOLD(Min Pos).

The first of these conditions corresponds to a situation where the velocity values are either too noisy (and hence flagged as invalid) to continue a segment or else they display a flat or decreasing trend. The second condition indicates that an unrealistically large shear exists in the window (which might be caused by velocity aliasing, for example) and the segment should not be continued, as it would likely contain erroneously large velocity differences.

This sliding-window approach to locating the segment endpoint has two salient characteristics. First, by examining several sequential data points, it allows small-scale (one- or two-gate) perturbations in the velocity measurements to be skipped over in the search process. Second, it makes no a priori assumptions about the magnitude of the shear; the only requirement for a segment to be continued is that the values in the window be generally increasing - by any amount. Since the size and strength of microburst outflows span a considerable range, this shear-independent aspect of the segment location process is useful. However, it also makes the algorithm quite sensitive to very small changes in the velocity field and allows very weak segments to be detected. The segment validation threshold tests described below are used to remove weak segments found in this process.

The segment start and stop tests compare the value in the current gate to the remaining values to determine if an increasing trend is present. If such a trend is present, the window is moved forward. The choice of which sample point to use for the next "current gate" is important since the velocity values along a shear segment are typically not perfectly monotonic. If the next sample point chosen as the "current point" is too large (relative to the average trend of the other values in the window), then the next window (using the new starting point) may not exhibit an increasing trend and may cause the segment to be terminated early. For this reason, the next "current gate" is chosen as a balance between moving up the trend rapidly to perform as few tests as possible (by choosing a large value) and climbing slowly to prevent stopping at a local spike (by choosing a small value). To

obtain this balance, all points in the window (which are greater than the starting point) are considered, and the one with the smallest positive increase is identified. All the points are again considered, and the first point encountered whose increase is both positive and less than or equal to THRESHOLD(Next Sample) (1.5) times greater than the smallest positive increase is chosen to be the next starting point. Hence, the starting point will move forward slowly, skipping over any points which would represent a relatively large step up the trend.

3.5. Segment Validation

The segment detection process is extremely sensitive, as it incorporates no absolute length requirements and no absolute strength requirements for the increasing trends it detects. On a typical velocity field with no significant microburst windshear present, a large number of candidate segments will be identified. Therefore, candidate segments must be subjected to a variety of validation tests to reject as many as possible of those segments which do not correspond to microburst events.

The validation process is rather complex and involves an iterative cycle of segment testing and trimming, which are described in Figure 16. On each cycle, basic tests are applied to determine if the segment should be rejected. A segment will be rejected if (a) it is too short, (b) if it has too small a velocity difference across it, (c) if it has too few valid velocity points, or (d) if it does not have a consistent increasing trend along its length. In addition, the endpoints of the segment may be trimmed back to ensure that (a) the endpoints are local extrema,(b) that they do not deviate too far from the local median, and (c) that the slope of the segment near the endpoints is adequate. The trimming of the endpoints will generally shorten the segment on each cycle, and the other validation tests will be repeated.

The basic length and strength tests will reject the majority of those segments which do not correspond to actual microburst shears. The threshold values used in these tests must be chosen carefully to balance the rejection of false segments with the detection of desired ones. The nominal thresholds used in current operational testing (0.95 km minimum length and 5 m/s minimum velocity difference) have been adjusted heuristically, based on several years of operational experience, to allow the detection of outflows in their earliest stages, when they are both small and weak. The probability of detection for weak microbursts (10 – 15 m/s) appears to be quite sensitive to the choice of these length and strength threshold values.

The slope trimming test in the validation loop is designed to serve two roles. First, it shortens true shear segments which may have "tails" of weak shear. Segments of this type may be formed at the edges of outflow regions, as shown in the conceptual model profiles of Figure 9. The second role for the slope trimming is to reject altogether those segments which comprise entirely of weak shear. Such segments, if sufficiently long, could pass the basic length and strength tests but do not represent actual aircraft hazards. By trimming these segments back repeatedly, they will eventually be rejected by the basic length test. The thresholds used in the slope trimming tests correspond to a shear of 2.5 m/s per km, which has conventionally been used by the TDWR community as the minimum hazardous shear level for aircraft operations [Mahoney, et al., 1989].

Repeat until segment is accepted or rejected:

If the segment length exceeds THRESHOLD(Slope Test Min Length), then trim the start and end points back (towards center of segment) until the velocity difference over NUMBER(Slope Test) range gates (0.5 km distance) at both ends of the segment is at least THRESHOLD(Slope Test Difference) (1.25 m/s).

Trim the start and end points back (i.e., towards the center of the segment) until each is a local extrema.

Reject the modified segment if either: (a) its new length is less than THRESHOLD(Min Div Seg Len) (0.95 km) or (b) the new velocity difference across the segment is less than THRESHOLD(Min Div Seg Vel) (5 m/s) or (c) the fraction of the sample volumes in the segment which are either marked invalid or have velocity values below or above the starting or ending velocity values, respectively, exceeds THRESHOLD(Fraction Bad) (0.125).

Reject the modified segment if the block mean velocity value, averaged over NUMBER(Slope Test) sample volumes (0.5 km), is not strictly monotonically increasing along the length of the segment.

Check the start and end points to verify that each is within THRESHOLD(Median Difference) (5 m/s) of the local median velocity value computed over NUMBER(Local Median) range samples (1 km). If both points meet this criteria, accept the segment. Otherwise, trim each point not meeting the criteria back one gate, and repeat the validation loop.

Figure 16: Shear segment validation test procedure

3.6. Azimuthal Association

Those segments which survive these validation tests are then associated across radar azimuths to form two-dimensional regions of shear. Any two segments which overlap in range by at least THRESHOLD(Min Overlap) (nominally 0.0 km) and are within THRESHOLD(Angular) (2.0 degrees) in azimuth are joined together into the same region. This association process continues until all segments have been grouped into regions. These aggregates are now thresholded based on their total area, number of segments, and maximum segment strength. The total area for a region is computed as the sum of the areas of the shear segments in the region. Regions with total area less than THRESHOLD(Total Div Area), fewer than THRESHOLD(Min Div Segments) segments, or having a maximum velocity differential (across the strongest segment in the region) less than THRESHOLD(Max Div Diff) are discarded. The result of these clustering and thresholding

processes is a set of "significant" regions of divergent shear, which are the final output of the divergence regions algorithm.

3.7. Alarm Generation in the Complete Microburst Algorithm

The divergence regions detected by the algorithm are used as the primary input to the remainder of the microburst detection algorithm. Storm features from upper elevation angles (e.g., reflectivity regions, convergence, rotation, and divergence features) are used to identify significant structures in the storm cells which may be used to increase the confidence in the existence of a windshear event. The basic temporal continuity and outflow strength requirements needed to generate an alarm, based on the detection of a divergence region, are relaxed by the presence of structures aloft. The various algorithms and decision processes involved in the detection of features aloft, and in the generation of final alarm regions, is quite complex and will not be described here. A summary of the use of features aloft may be found in [Campbell, 1989], while a complete description of the algorithms is presented in [Campbell and Merritt, 1988].

The output from the complete detection algorithm is a set of final alarm regions. These alarm regions are essentially a validated subset of the initial set of divergence regions, possibly merged based on spatial proximity and features aloft. Each alarm region has an associated set of shear segments, obtained from the divergence region(s) on which the alarm region is based.

3.8. Shape Generation and Hazard Level Estimation

The alarm regions are processed by a "shape" determination stage, which fits a smooth outline to the set of shear segments which make up each alarm region. The shape is constrained to be a rectangle with semi-circular ends, referred to as a "bandaid" shape. Parameters define the maximum size and ellipticity of the shapes, and an alarm region may be broken up into multiple smaller shapes to satisfy these constraints. The shape generation algorithm is described in [F.W. Wilson, et al., 1991].

The alarm strength associated with each bandaid shape produced is determined by the collection of shear segments which were used to produce the shape. If the alarm region is described by a single bandaid region, then all segments for the alarm region are used in determining the strength. If the alarm is subdivided into multiple bandaids, each bandaid is associated with a subdivided set of segments, each of which is used to arrive at the strength value. For each shape, the strength value is taken as a percentile of the set of velocity differentials across the segments associated with the shape.

4. ALGORITHM CASE STUDY FOR JULY 11, 1988

A strong microburst occurred on the afternoon of 11 July 1988 during the TDWR operational testing period at Denver's Stapleton airport. This microburst was particularly strong, reaching a maximum velocity difference of roughly 40 m/s, and was located along an active approach path for Stapleton airport. Several aircraft encountered the microburst and experienced serious difficulties; one aircraft came within 50 ft of the ground before recovering. Because of the strength and operational exposure of this microburst, it has been studied extensively by the aviation meteorology community [Schlickenmaier, 1989]. A portion of the time history of this microburst outflow, as observed with the FL-2 radar, is presented below, along with the performance of the divergence detection algorithm on this case.

4.1. Minute-by-Minute Observations of Microburst Development

The radar images shown in Figure 17 depict the radar reflectivity and velocity measured by FL-2 at time 22:04:01 on 11 July 1988. At this time, a number of moderate-reflectivity storm cells were present in the vicinity of Stapleton airport. Roughly seven microbursts had already been observed on this day, and another 12 microbursts were to occur in the next two hours. The close-up image in Figure 18 shows an existing microburst south and east of the airport area (7 km range, 300° azimuth), and a cell with a weak divergent outflow just touching the southern extent of the airport runway complex. The maximum surface reflectivity for this storm cell is 30-35 dBz, and a small weak divergent outflow is present. The overlays on this image indicate the divergence segments (in white) and divergence regions (in red) which have passed the respective threshold tests. The approximate outline of the microburst outflows are drawn in green, based on a careful visual examination of the reflectivity and velocity measurements of both the FL-2 and UND radars.

At this time, the outflow from the storm cell is somewhat disorganized and rather weak. The strongest shear is located at a range of 12 - 14 km at azimuth angle 298°, but it extends across too small a distance for the algorithm to detect it. The algorithm does detect six shear segments on the southern edge of the cell, which result in a divergence region being identified.

The next sequential radar scan at the surface was taken about one minute later, at time 22:05:04, and is shown in Figure 19. The region of positive velocities has filled in somewhat now, and the velocity values are slightly stronger, as are the negative velocities. The shear region along the 300° azimuth line has enlarged slightly, allowing the algorithm to detect the shear region there. Note that the center portion of the outflow, where the segments were detected on the previous scan, has weakened and shifted so that no segments are detected there. A few weak segments were detected at the far southern tip of the outflow but were not sufficiently large (in area) for a divergence region detection.

At 22:06:01, the outflow has increased both in size and strength, with a distinct visual signature of well-defined positive and negative velocity regions (Figure 20). The shear region is detected with a number of segments, although some radials are missed and some

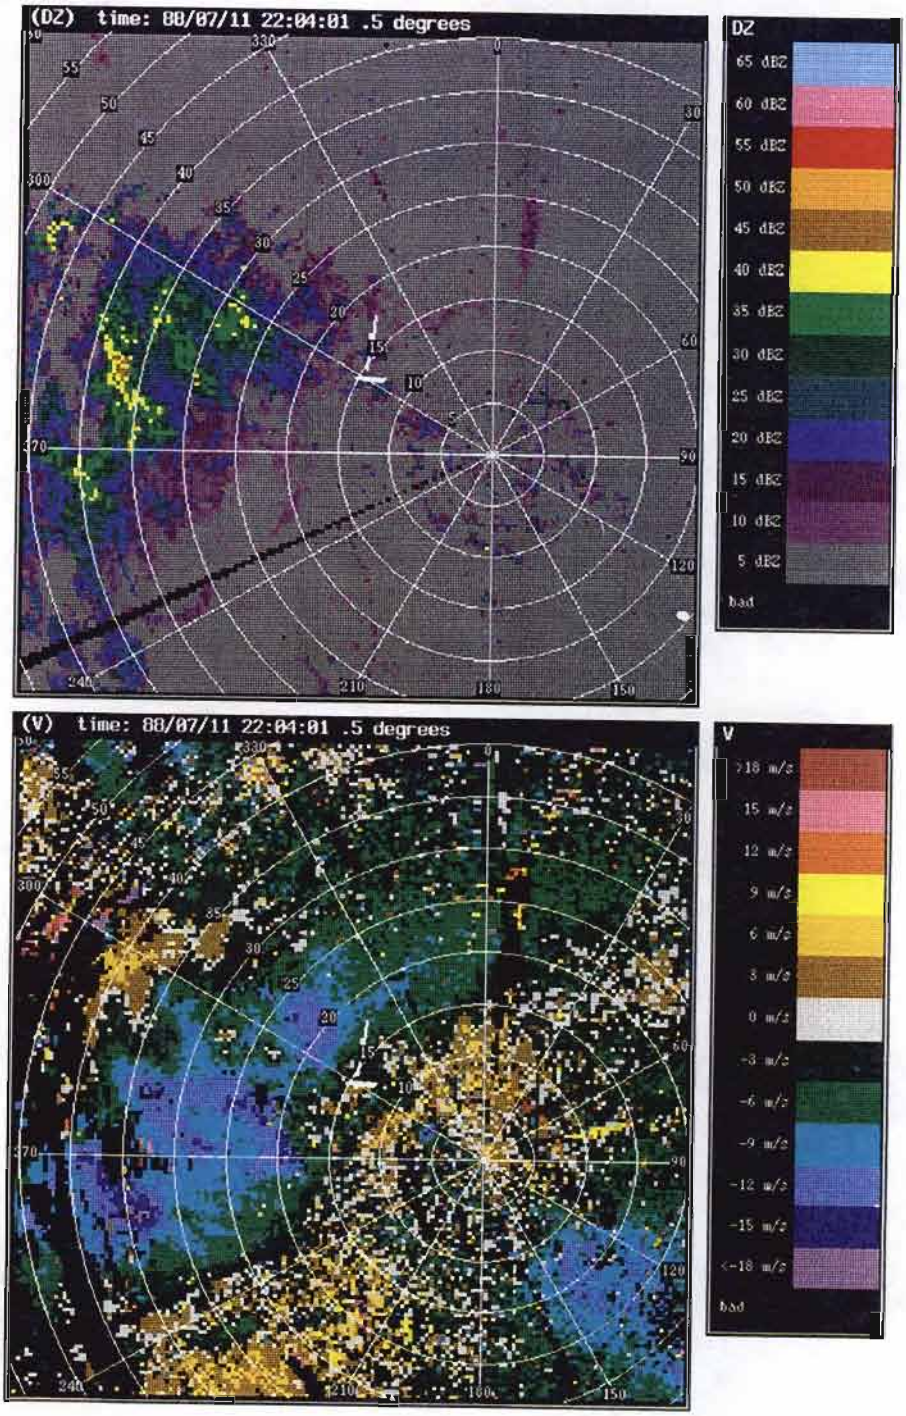

Figure 17: Precipitation and winds at start of case study for 11 July 1988. These radar images show scattered storms to the west of Stapleton airport, with a small microburst cell developing just off the east end of the east–west runways.

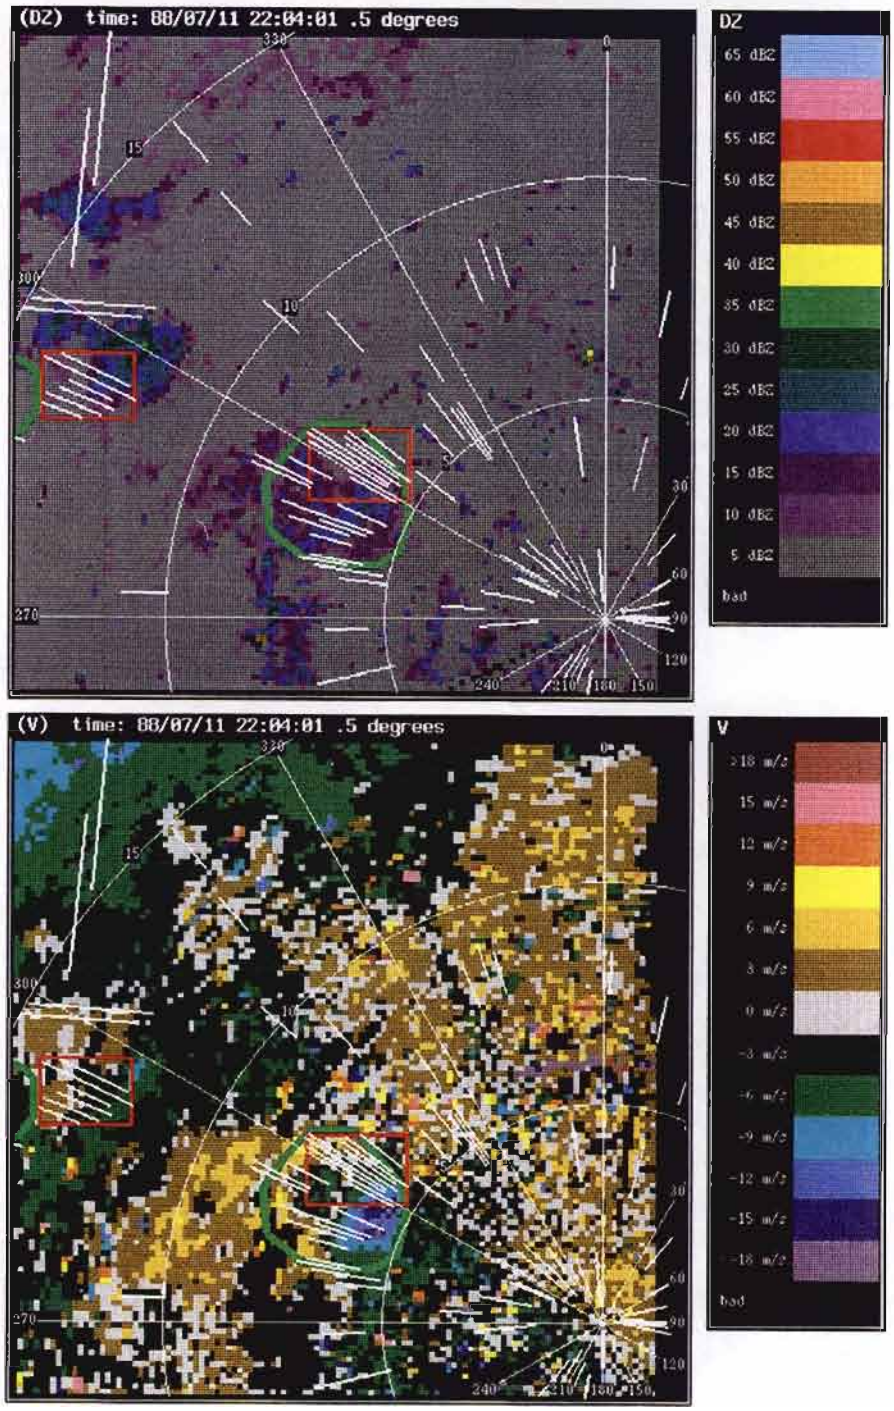

Figure 18: Close–up image of radar measurements at start of case study (22:04:01) on 11 July 1988. Divergence segments are shown in white, divergence regions in red, and ground truth regions in green.

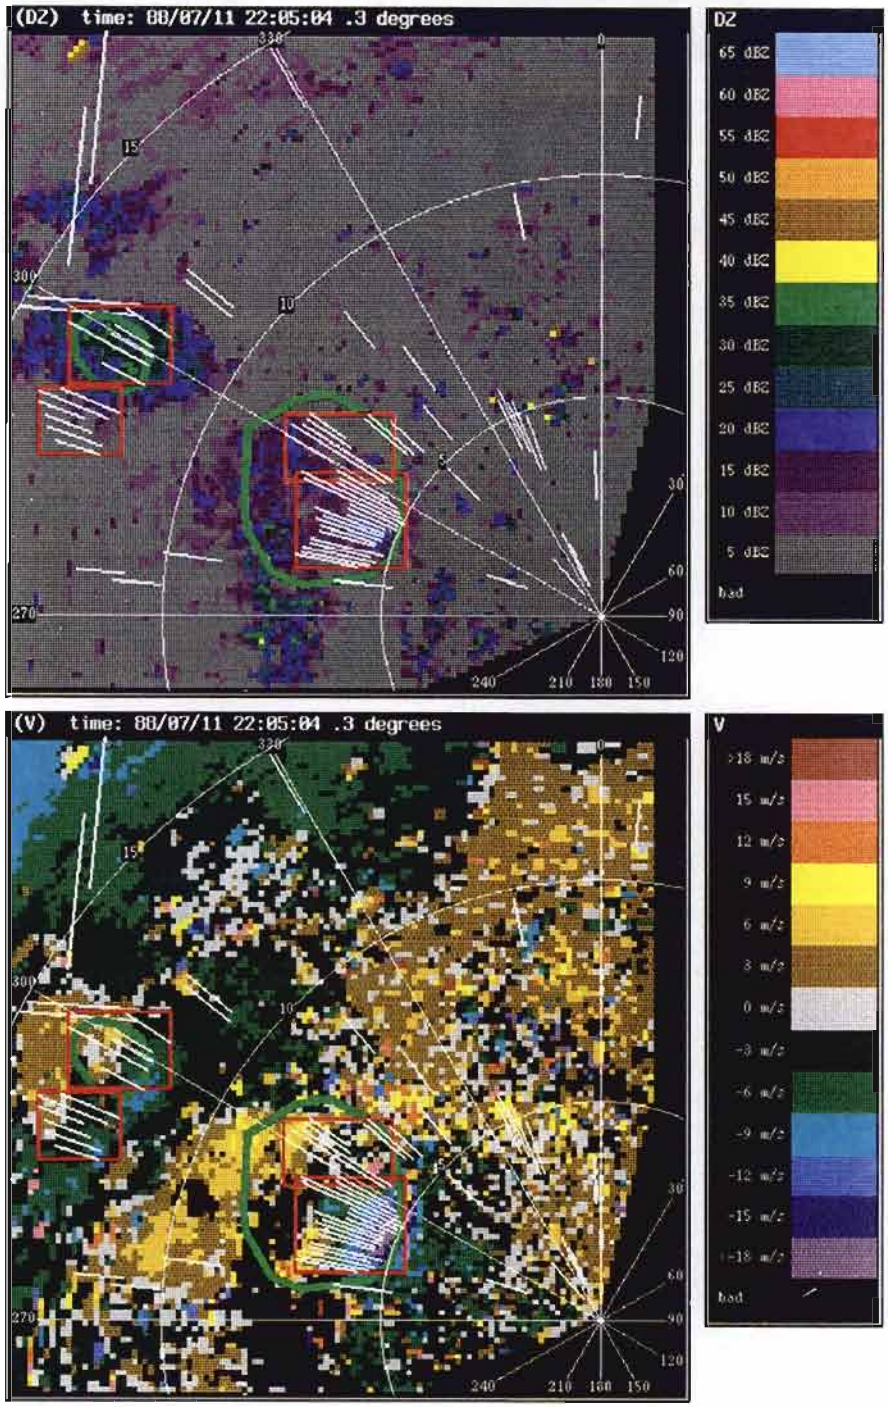

Figure 19: Radar measurements from 11 July 1988 at 22:05:04.

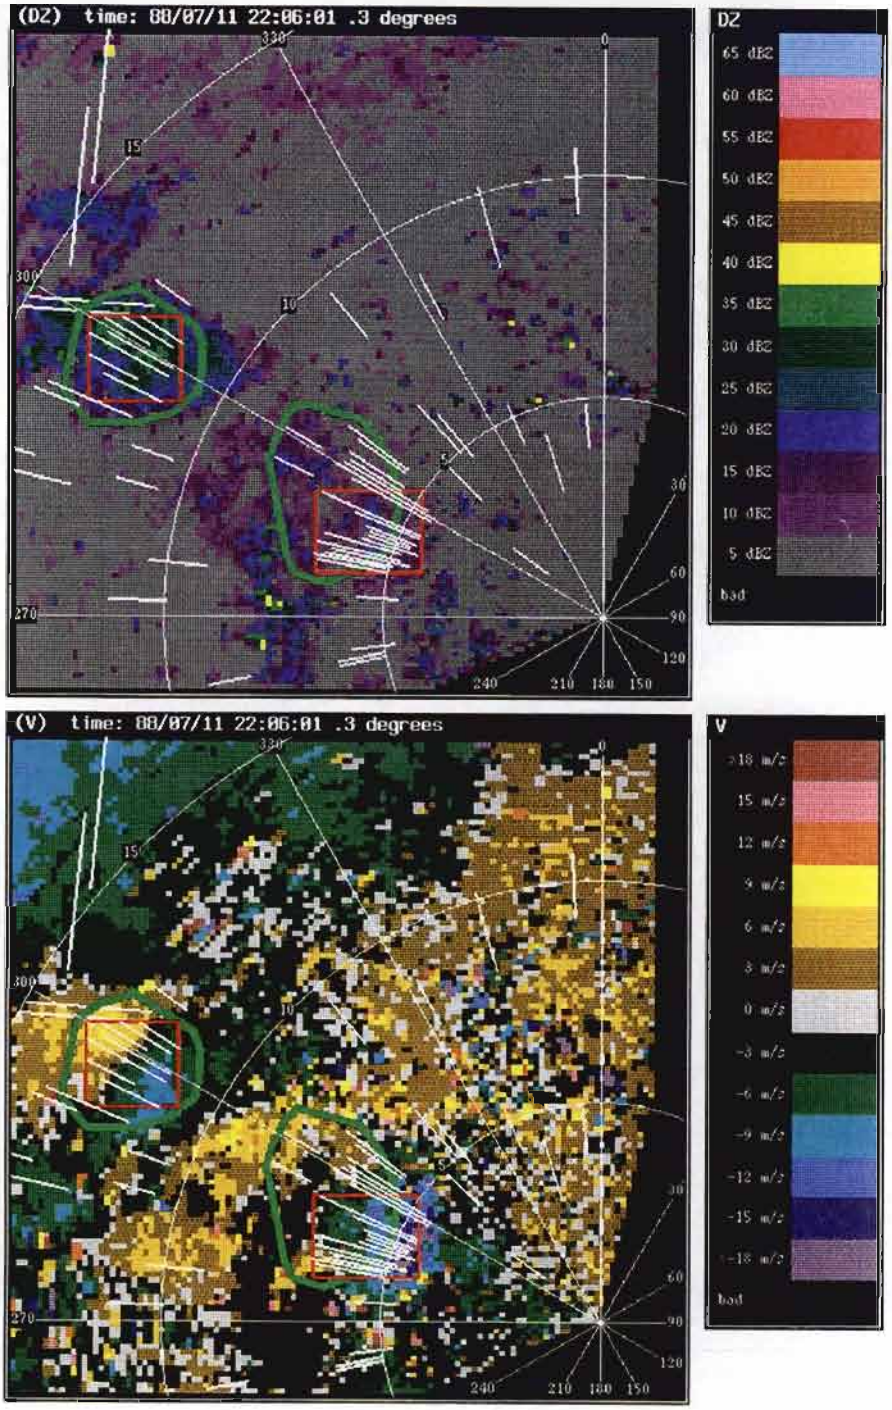

Figure 20: Radar measurements from 11 July 1988 at 22:06:01.

of the segments are shorter than the extent of the apparent shear. The region formed by these segments, however, covers the actual shear area quite well.

One minute later (22:06:58), the outflow shown in Figure 21 has again grown and is now well detected by the divergence algorithm. The strongest portion of the outflow is seen at about 300° azimuth, and the segments accurately cover the extent of this shear. A small, weak outflow has developed to the southwest of the airport microburst, and the shear segments from this secondary event have been merged into the region for the airport microburst. The segment merge criteria determined adequate overlap between the segments from the two events, and hence the regions were combined. This example illustrates the sensitivity of the segment association rules used to determine the extent of the divergence region. This example shows the possible negative consequences from the use of strictly local decisions about segment overlap in the algorithm for growing the shear region. Also, note a weak region of divergence on the northern tip of the airport microburst where three weak segments were detected. The velocity differential across these segments is roughly 9 m/s.

On the next surface scan, at 22:08:03, the measurements shown in Figure 22 indicate an extremely rapid growth in the surface outflow on the northern tip of the previous divergence region. The weak 9 m/s outflow from the previous scan has grown to a peak value of 25 m/s, and the area of the outflow has increased considerably as well. This dramatic increase in the outflow strength and extent is the result of the impact of a descending reflectivity core, and the associated downdraft, with the surface. Comparison of the reflectivity images for these two times indicates a significant expansion and increase in reflectivity at the surface which coincides with the large growth of the outflow.

At time 22:09:00 (Figure 23), the southern portion of the outflow region has weakened and begun to break up while the northern portion continues to expand and intensify. The strong outflow to the east of the east-west runways is well detected by the divergence algorithm, and the remainder of the outflow region results in two smaller divergence regions being identified. While the divergence algorithm has done a rather good job of identifying this very strong outflow region, the segment detection criteria have resulted in a segment being broken in the strongest region of the outflow. The velocity profile through the microburst center is shown in Figure 24, along with the segments found by the algorithm. The slight plateau in the velocity profile along the radial at azimuth 306.5° caused the shear segment to be broken into two pieces. Since no logic is present within the divergence algorithm to join such segments after they have been broken, the net loss across the shear region at this azimuth is underestimated by a factor of two. Fortunately, the adjacent azimuths have the same strength outflow, so this error on one radial does not result in an overall error in the estimated strength of the region. If this segment-splitting error had occurred on more radials, the maximum strength of the region could have been severely underestimated. This underestimation of the shear region strength could potentially result in the region not being detected as a microburst if the final strength were below the required threshold level. Once a region has been detected as a microburst, however, the shape generation algorithm will join any segments on the same azimuth and sum the strengths to compensate for this segment-splitting effect.

The continuing growth and expansion of the microburst outflow is seen in the next two radar scans, at times 22:10:03 (Figure 25) and 22:11:00 (Figure 26). The windfields in

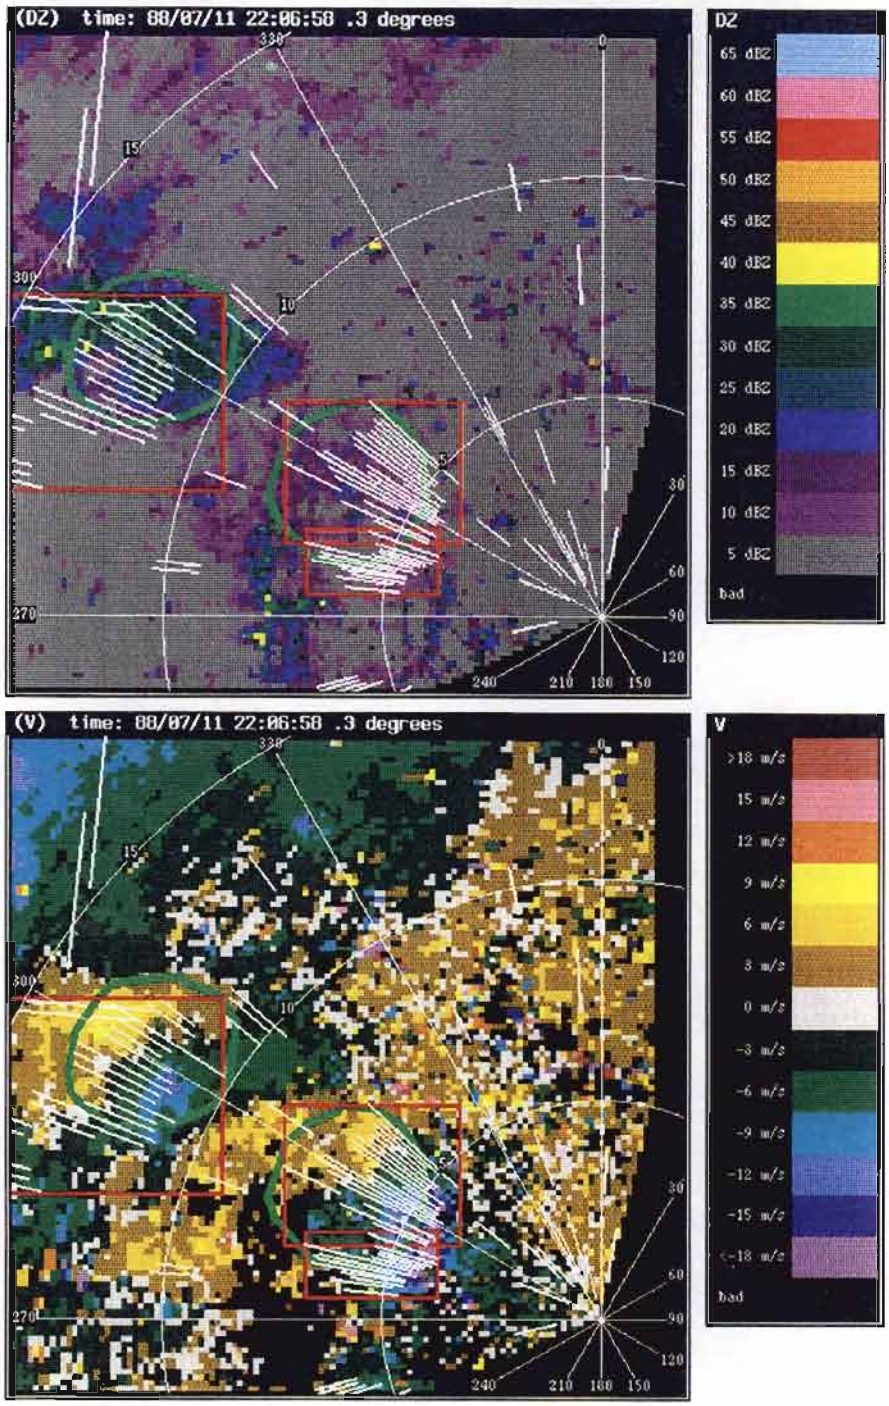

Figure 21: Radar measurements from 11 July 1988 at 22:06:58.

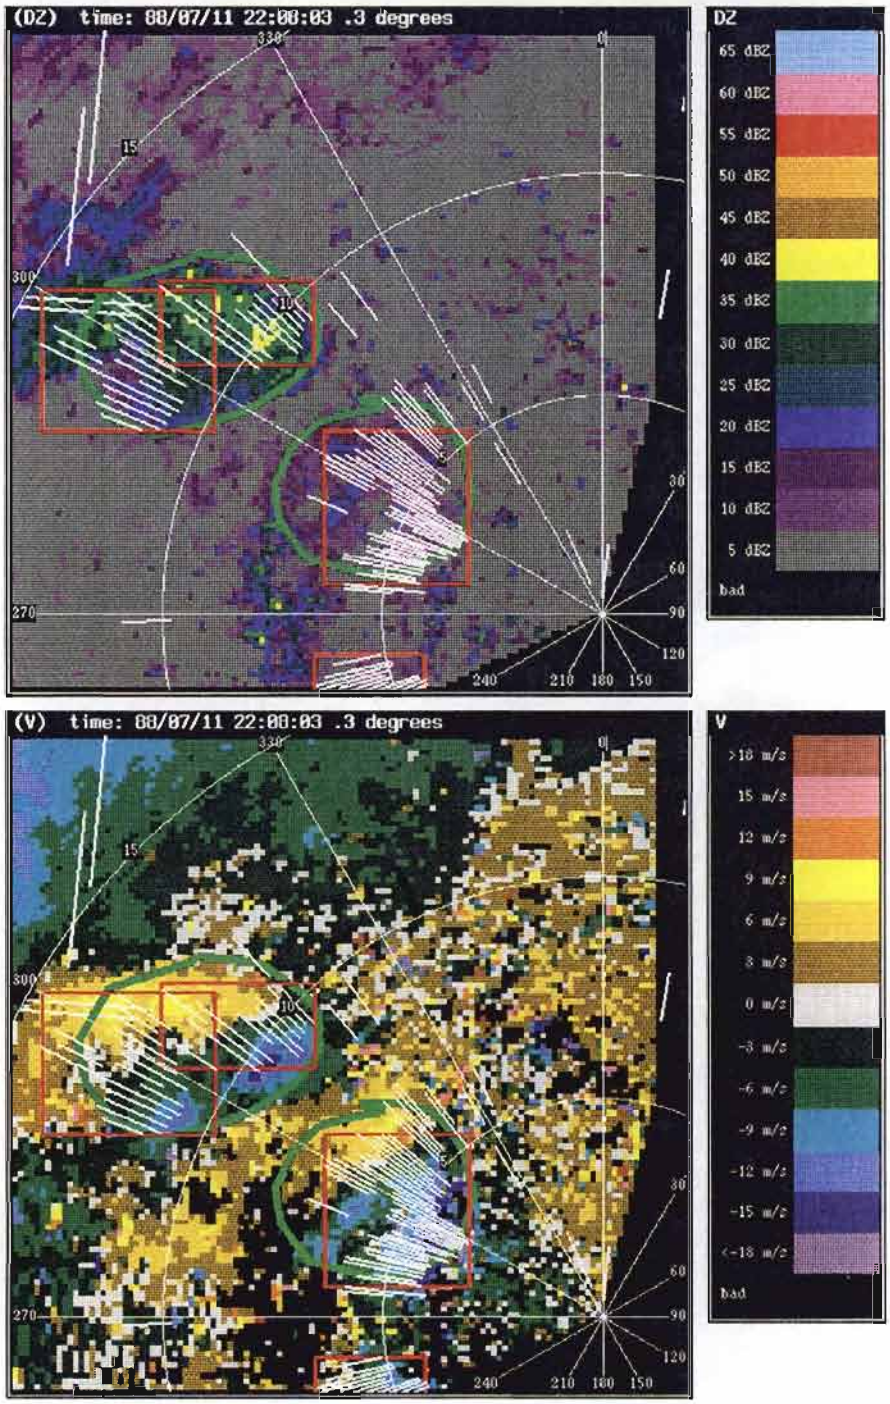

Figure 22: Radar measurements from 11 July 1988 at 22:08:03.

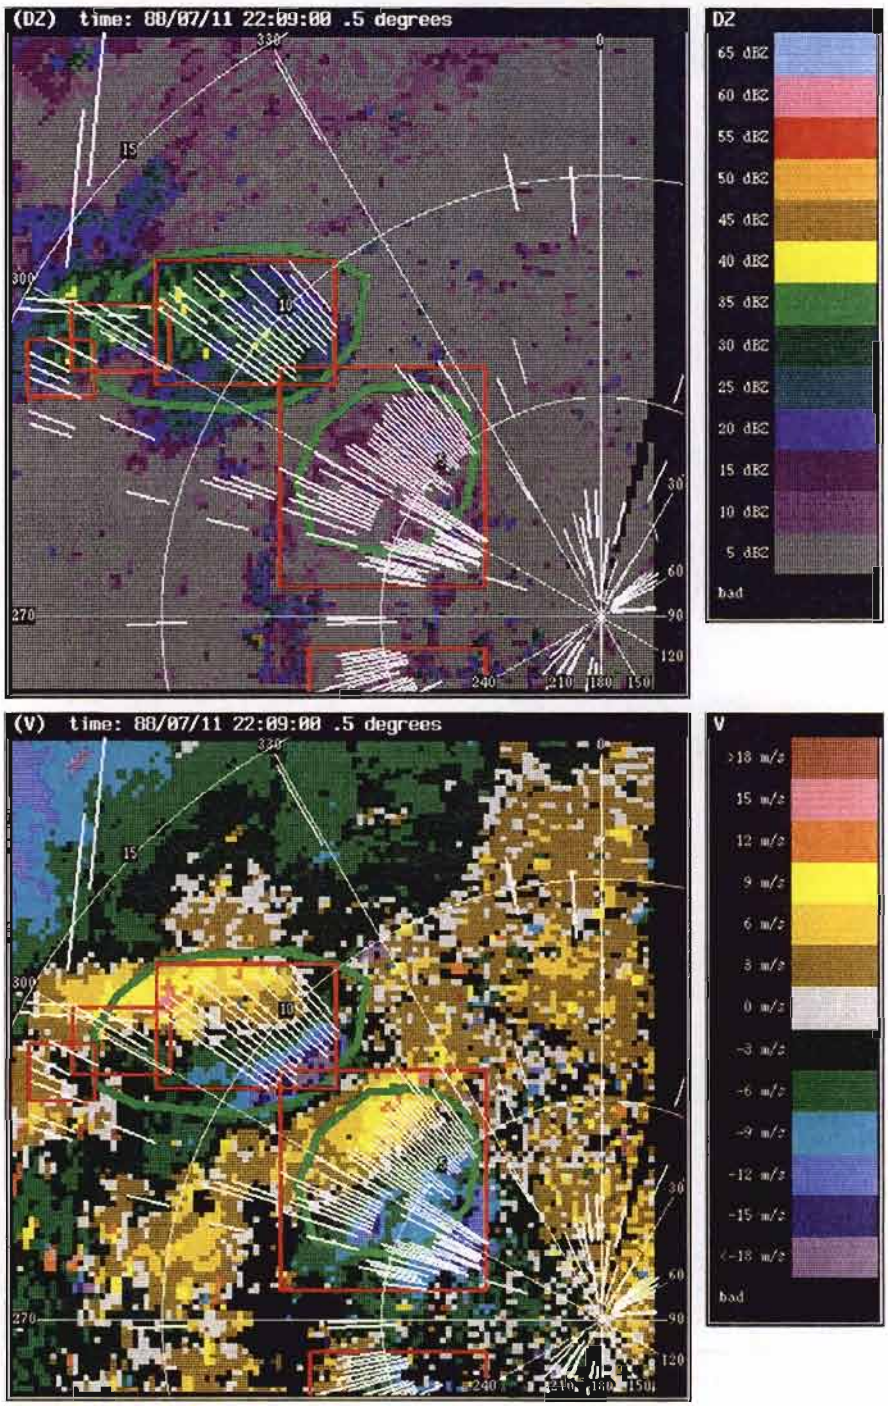

Figure 23: Radar measurements from 11 July 1988 at 22:09:00.

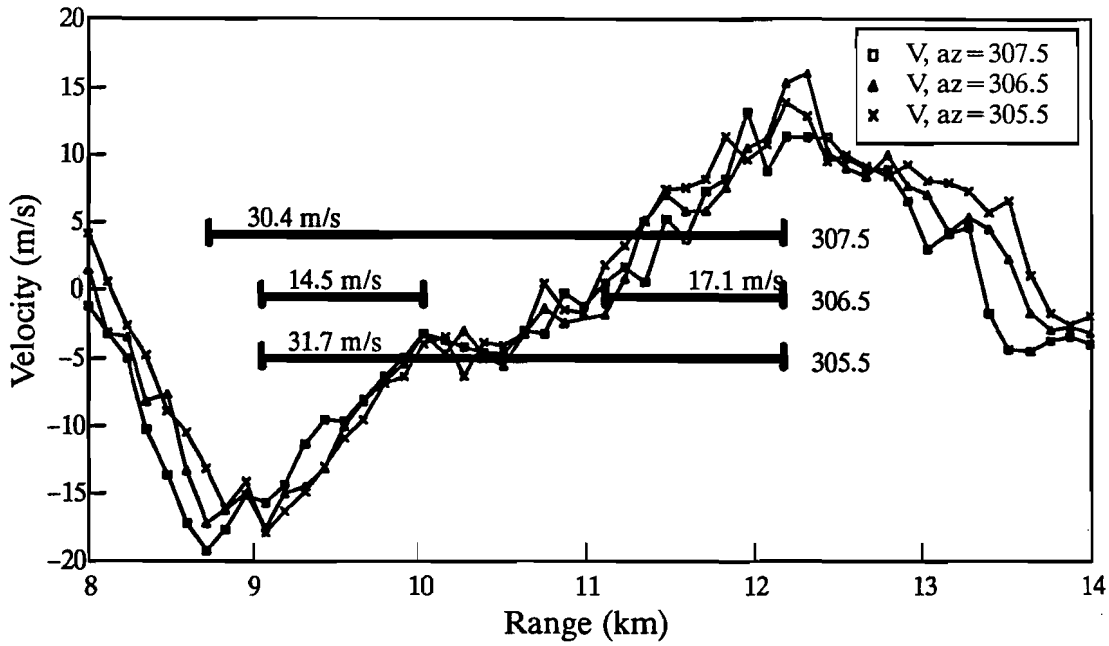

Figure 24: Velocity measurements at three azimuth angles through strongest portion of airport microburst at 22:09:00. Shear segments detected by algorithm are overlayed, indicating the segment at azimuth 306.5° is broken.

these two observations show the increasing windspeeds in the positive and negative radial velocity regions of the main outflow along with the weakening of the older outflow located just at the end of the east-west runway. The outflow from the main microburst is also seen to be colliding with the expanding outflow from the large microburst to the southeast of the runway complex, giving a very sharp convergence line at the boundary between the two outflows. The strong portion of the main outflow is completely detected by the divergence algorithm on these two scans, although the two microbursts in the airport vicinity are detected as a single, large divergence region.

4.2. Summary of Algorithm Performance for 11 July 1988 Event

The divergence algorithm was successful in detecting this microburst throughout its hazardous lifetime. A divergence region was created by the algorithm at the initial impact of the microburst and for each of the subsequent radar scans. Despite this good performance at the region-finding level, it is evident that the algorithm was much less successful when the individual shear segments are considered. On many of the radar scans, the algorithm did not find shear segments in regions of significant shear. Also, many (weak) shear segments were found in areas outside of the region of hazard.

The plots in Figure 27 compare the regions detected by the algorithm with the areas determined (by manual analysis) to be the region of significant shear. On each scan after 22:04 (the airport microburst began at 22:05) the hazardous region was almost completely enclosed by the detected divergence region, but the number of regions generated by the algorithm varies from scan to scan. This scan-to-scan variation in the clustering of

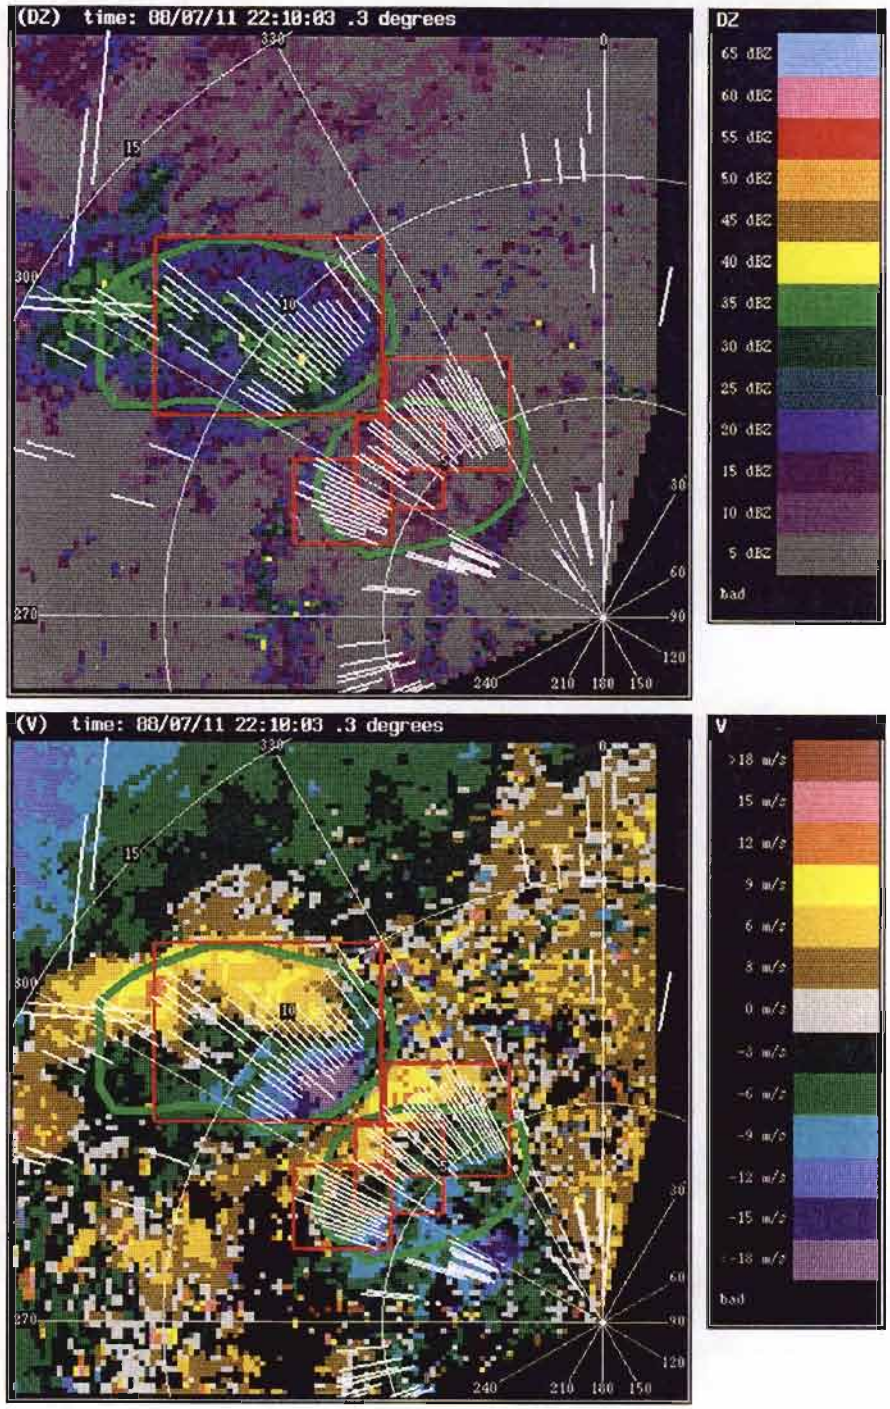

Figure 25: Radar measurements from 11 July 1988 at 22:10:03.

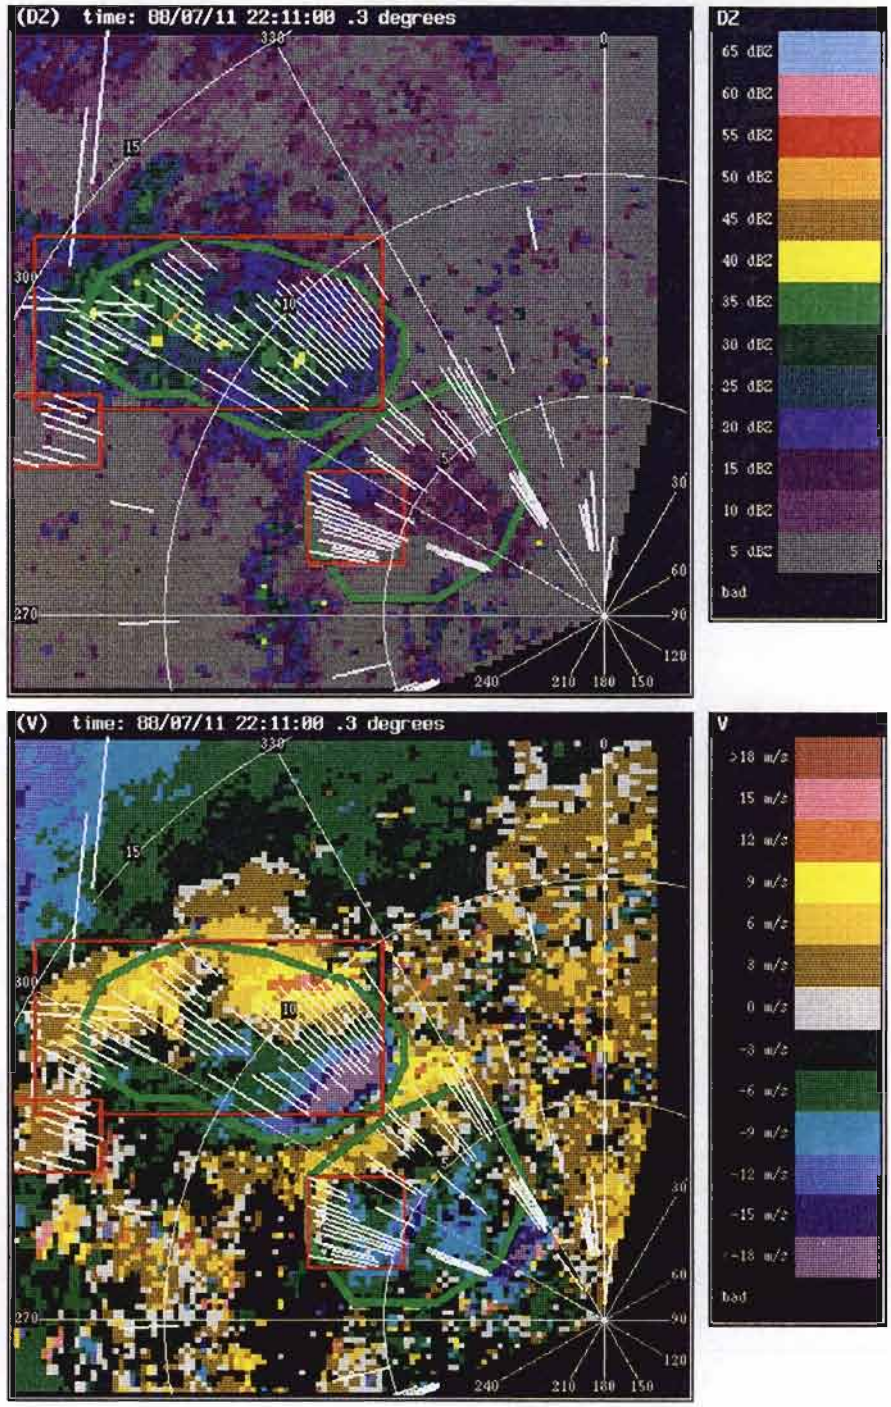

Figure 26: Radar measurements from 11 July 1988 at 22:11:00.

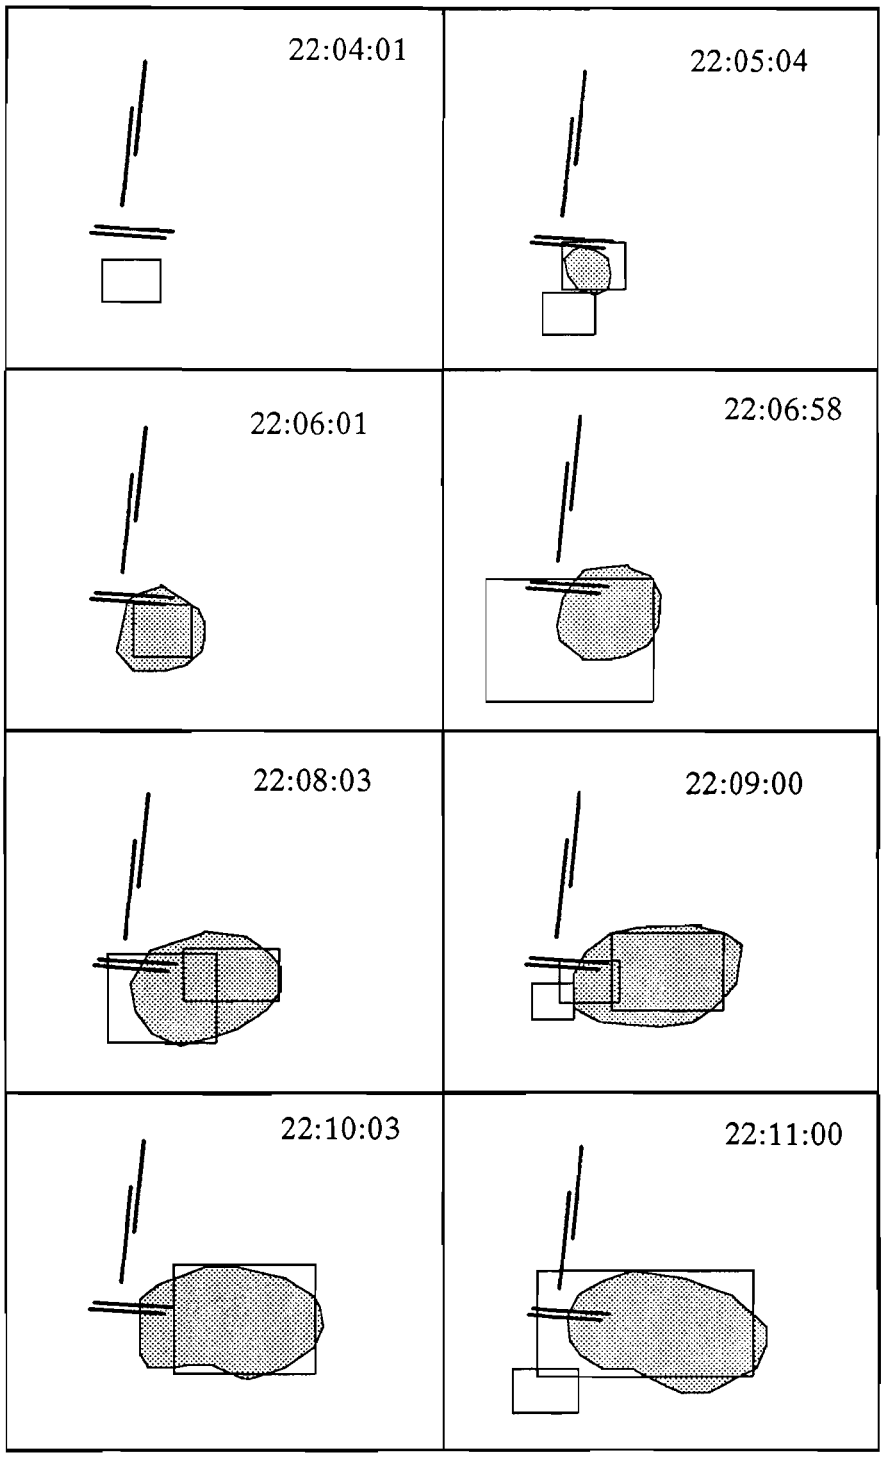

Figure 27: Summary of the detected regions for the 11 July 1988 study case. Regions are indicated by rectangles; ground truth is represented by stippled outline. While the microburst region is detected on each scan, the number of regions detected changes from scan to scan.

segments is a reflection of both (a) the inability to reliably detect all the shear segments in the hazardous region and (b) the crude nature of the clustering algorithm used to form the regions.

Although the algorithm sometimes fails to identify all of the shear segments in the target region, the detection statistics described in the following section indicate that the clustering process is usually sufficient to detect these shear regions. The large variability in the size and number of regions detected is undesireable from an operational user interface standpoint; therefore, the time association and shape generation components of the complete microburst algorithm were designed, in part, to compensate for this characteristic of the divergence algorithm.

5. ALGORITHM PERFORMANCE STATISTICS

5.1. Analysis of Divergence Detection Algorithm Performance

The divergence algorithm was applied to a large sample of radar measurements from three selected days in the 1988 experiment to obtain a statistically meaningful estimate of its detection and false-alarm performance. These days were selected because of the large number of microbursts which occurred during the data collection on those days. While the detailed comparison described here is limited to these few days of data, experience from the real-time observation of the algorithm performance (over a period of many months) verifies that the statistics developed from these days seem to be typical of the algorithm operation.

For each day selected for the analysis, a thorough subjective examination of the radar reflectivity and velocity measurements was conducted by experienced radar analysts to determine the location and strength of all actual microbursts. This ground truth information was then compared to the algorithm-generated detections to compute the number of hits, misses and false alarms. In each case, the entire data sample for the day was used in the analysis to avoid any possible biases associated with selecting particular time periods for study.

The majority of the ground truth information developed for this evaluation was based on the measurements made by the TDWR testbed radar. In some cases, data from additional sensors (the surface mesonet and a second doppler radar operated by the University of North Dakota) were available for corroboration. This support data has been used only in exceptional situations where the algorithm comparison to the data from a single radar windfield appeared to indicate a problem. The use of data from a single-doppler radar for evaluating algorithm performance has been supplemented by several studies comparing the radar-observed microburst events with those sensed by the surface mesonet [DiStefano, 1988], [DiStefano and Clark, 1990]. These comparisons have indicated that very few microbursts are unobserved by the radar, and hence the single-radar-based ground truth method is fairly accurate.

The basic criterion used in the ground truth analysis to define a microburst is the presence of a wind speed difference of at least 10 m/s over a distance of no more than 4 km. Note that the velocity difference may extend beyond the 4 km scale, so long as the required 10 m/s difference exists within some 4 km sub-region. A microburst is considered "ended" when the velocity difference (over a 4 km scale) drops (and remains) below 10 m/s for a period of at least two minutes. For each such microburst observed on a radar scan, a polygonal outline is recorded, along with the strength (total velocity differential) of the outflow.

The performance of the algorithm is characterized here using the common probability of detection (POD) and probability of false alarm (PFA) statistics, defined as follows:

$$POD = \frac{\text{Number of detected events}}{\text{Number of events}}$$

$$PFA = \frac{\text{Number of false alarms}}{(\text{Number of correct alarms} + \text{Number of false alarms})}$$

For this analysis, an event is defined as a single observation of an actual microburst by the radar on a low-elevation angle scan. Each actual microburst is typically observed on several sequential scans, and hence represents several events. Only those microbursts which fall within a 30 km radius of the radar were considered in this scoring. Since the microburst phenomena and the algorithm output are both two-dimensional regions, some criteria are required to define a "detection." For the evaluation conducted here, an event is considered detected by the algorithm if the region produced by the algorithm overlapped the actual (ground truth) microburst region by any amount. A region produced by the algorithm was considered to be a false alarm if it did not overlap a ground truth region.

To provide an operationally realistic evaluation of the algorithm, certain alarms which would be strictly classified as "false alarms" are tallied separately. Declarations which overlap actual events which appear on radar scans within two minutes (before or after the current scan) are not considered false alarms, but are counted as "early" or "late," respectively.

The results of the comparison between the algorithm detections and the ground truth for the three studied days are presented in graphical form in Appendix 1 and summarized in Table 2. This table lists the number of microburst observations on each of the case days and the number of those observations hit and missed by the algorithm. The detection performance is seen to vary across the three days, with an aggregate POD of just over 90 percent. The table also shows the number of false alarms generated by the algorithm, which is consistently high (roughly 30 percent of all detected regions are false). The detection performance is shown as a function of outflow strength in Figure 28. This figure indicates that detection performance improves rapidly with microburst strength over the range of 10

Table 2.
Performance Results for the Divergence Regions Algorithm. Based on comparison to ground truth information from Denver, 1988 Cases

	June 10	June 21	June 25	Total
Number of hits	158	217	196	571
Number of misses	28	10	5	43
Number of false alarms	55	85	114	254
Probability of detection	0.85	0.96	0.98	0.93
Probability of false alarm	0.26	0.28	0.37	0.31

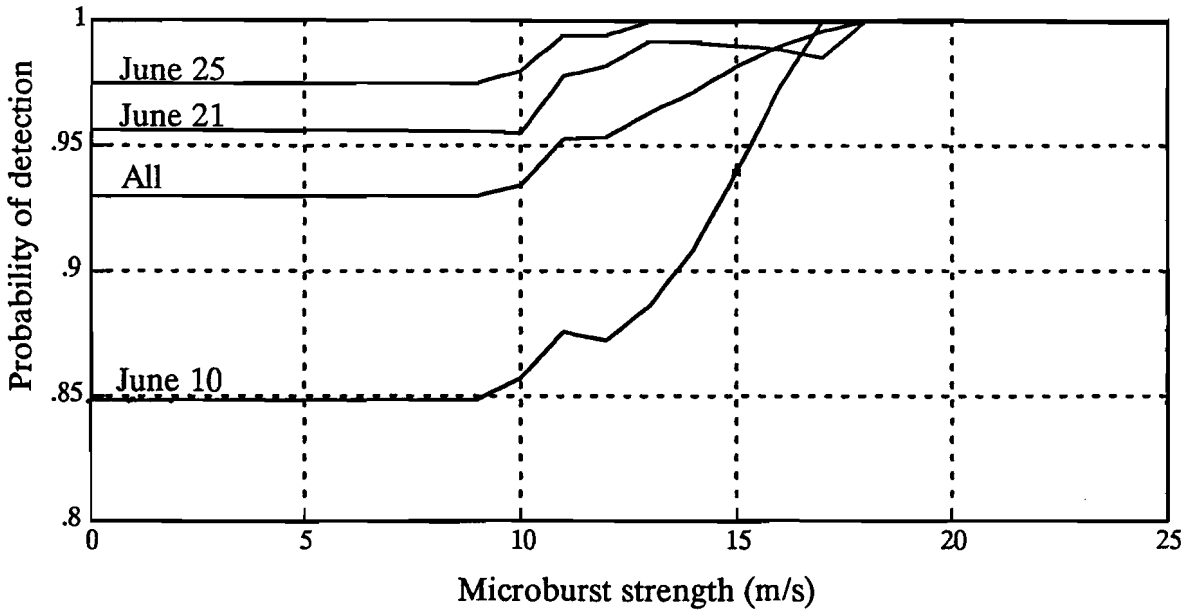

Figure 28: Probability of detection for divergence regions as a function of outflow strength. The probability of detecting all events with strength above a certain level is shown here, as a function of the strength level, for each of the case days, and in total.

to 15 m/s, and that microbursts with strengths above 15 m/s are detected at least 98 percent of the time.

5.2. Relationship to the Complete Microburst Algorithm

The performance of the complete TDWR microburst algorithm has been evaluated against a very large set of microburst observations using the same methodology described above. The results of this performance evaluation, from both the 1988 and 1989 measurement programs, are shown in Table 3, adapted from [Evans, 1990]. The detection performance for the complete algorithm is 0.90 in Denver and 0.96 in Kansas City, compared to the 0.93 value obtained for the divergence algorithm alone in Denver. These values are very comparable, given the day–to–day variation in observed performance. Note that the evaluation from Denver 1988 data presented in Table 3 includes the three case days for which the divergence algorithm was evaluated, plus two additional days of data.

The detection performance of the complete microburst algorithm is dominated by the divergence algorithm, hence the similarity in performance measurements. The temporal continuity test in the complete algorithm, which requires a divergence region to be detected at least twice in a row, dramatically reduces the false–alarm rate. The PFA for the complete algorithm is in the 0.05 – 0.07 range, whereas the divergence algorithm alone generates false alarms nearly 30 percent of the time. The second factor which reduces the likelihood of false alarm in the complete algorithm is the strength threshold, which requires a divergence region to have a strength of at least 10 m/s (in the absence of features aloft) before an alarm will be issued. The POD and PFA values for the divergence algorithm alone were based on

Table 3.
Performance Analysis for the Complete TDWR Microburst Algorithm.

	Denver (1988)	Kansas City (1989)
POD (strength > 10 m/s)	.90	.96
POD (strength > 15 m/s)	.97	.99
PFA	.05	.07

*Based on comparison to ground truth information, from the Denver (1988) and Kansas City (1989) measurement programs. Adapted from [Evans, 1990].

the strength thresholds used for region detection, which require only a 5 m/s strength. This strength thresholding helps eliminate many false alarms but also removes correct detections. The complete algorithm will reduce the strength threshold and relax the time continuity requirement when features aloft are detected in combination with a surface divergence region. This adaptive threshold adjustment is particularly useful in detecting microbursts in their earliest stages of development while still keeping the false-alarm rate at an acceptable level.

5.3. Interpretation of the Truth-Overlap Performance Statistics

The ground truth information used for the above scoring process is the result of considerable examination of a large number of data fields, and is likely to contain many small inaccuracies and perhaps some significant mistakes or omissions. The comparison process, which simply requires some overlap between the truth region and the algorithm-generated region, is very tolerant of small inaccuracies. The comparison method used in the evaluation of the divergence algorithm was performed manually, and "marginal" situations were subjectively resolved. It is unlikely that the small inaccuracies which are inherent in this form of ground truth data have any noticeable impact on the final performance statistics. The performance values which have been reported for the complete TDWR microburst algorithm are typically based on a ground truth comparison performed in an automated fashion, where "marginal" cases are handled according to strict overlap or proximity rules which often fail to characterize a complex situation properly. These statistics may be more sensitive to small variations in the ground truth outlines.

This lack of sensitivity to the details of the ground truth outlines is also an indication of a deficiency in the performance metric itself. Since any degree of overlap between an alarm and a truth region is scored as a perfect detection, there is no accounting for the degree of accuracy in the detection process. This lack of accuracy evaluation also applies to the strength information provided by the algorithm, which is not considered at all by the conventional scoring methods. These methods are very useful guides to the gross ability of the algorithm to detect hazards; if the overlap scoring metric produces a low probability of

detection, then there is clearly a problem with the performance of the system. Likewise, a high false-alarm rate by this metric is generally indicative of a problem. The use of a truth-overlap scoring method is an appropriate technique for monitoring overall system performance and for the identification of anomalous performance requiring investigation.

The truth-overlap scoring approach, however, is not fully adequate for the evaluation of finer details of the algorithm performance and is not an appropriate metric for comparing alternative algorithm implementations or parameter settings. In addition, the overlap-based scoring figures should not be interpreted as representing the system performance as would be perceived by air traffic controllers or pilots using the system. The perceptions of these operational users will be strongly influenced by the detailed correspondence between the actual winds experienced by particular aircraft and the alerts provided by the system.

5.4. An Objective Path-Based Performance Metric

A new performance evaluation method is proposed here to attempt to provide a much more refined evaluation of the correspondence between the algorithm output and the actual windfield. This new approach examines the longitudinal winds along each of a number of straight line segments (representing potential aircraft flight paths), computes a measure of the divergent wind shear present along that path, and compares that value to any algorithm regions intercepted by the path. To obtain the longitudinal winds along the paths, a dual-doppler analysis must be performed. This evaluation technique is therefore appropriate only for those measurement periods where adequate coverage from a second doppler radar is available. The requirement to perform this dual-doppler analysis also increases the computer processing cost of this method (relative to the truth outline scoring approach) but reduces the need for manual expert analysis to obtain the ground truth regions.

This path-based approach offers the ability to fully capture the complex nature of the windfield hazard and evaluate the appropriateness of the detected region in a manner which directly relates to the perception of operational users of the system.

The selection of path segments to be used in the analysis allows it to be tailored to a number of circumstances. To obtain a large statistical sample, hundreds of paths may be used covering the entire area where dual-doppler measurements are available. Alternatively, a small number of paths over a specific area can be used to focus on performance for a specific event or runway area. Path orientations may be restricted to match actual approach and departure corridors around an airport, or paths may be oriented in all directions to assess different aspects of the viewing angle dependence of the event strength.

While the path-based approach would seem to provide a much more refined evaluation of algorithm performance, it does not consider a number of significant aspects of the system operation. First, the path-based approach evaluates the degree of match between the current winds and the current algorithm outputs. It would not easily accommodate the evaluation of the timeliness of detections (i.e., how early in the microburst lifetime was the event detected) and it does not easily handle the discounting of alerts which immediately precede or follow actual hazards in time. The path-based approach does not embody any

concept of an individual microburst event, since all performance is viewed relative to the winds along specific paths.

The performance statistics which would be obtained from this sort of path-based evaluation would not be easily compared to values derived using the more common truth-overlap approach. The path-based approach should be considered as just one of (perhaps) several performance metrics useful for evaluating different aspects of the TDWR microburst algorithm behavior. It is most appropriate for examining the detailed performance of the hazard-estimation or area-delimiting aspects of the system and is less appropriate as a general-purpose performance metric. Scoring approaches based on microburst outlines, as used in previous years, are probably more appropriate for many other performance evaluation applications.

5.5. Description of the Path-Based Scoring Method

The path-based approach uses a mesh of simulated flightpaths, spread across the entire dual-doppler analysis area, for computing the performance measures. A typical configuration might use paths which are 6 km long (roughly the length of the microburst alert area around a runway), spaced 1 km apart in the X and Y dimensions, and rotated in 30° increments. This arrangement would result in roughly 5400 paths covering the analysis area.

Each of these paths would be intersected with any microburst alarms generated by the algorithm to determine the algorithm-generated strength. In the case of multiple alerts intersecting the path, the largest value would be selected. An estimated windspeed loss estimate is then calculated for each path, based on the dual-doppler windfield. This loss estimate is then associated with the path for use in evaluating the algorithm performance.

Each path in the analysis area may now be categorized as a *hit* (both alarm and loss estimates above 10 m/s), *miss* (loss > 10 with no alarm), *false warning* (alarm but loss < 10) or a *null path* (no alarm and loss < 10). The usual detection and false warning probabilities may then be calculated from the total number of hit, miss, and false warning paths. The terminology "false warning" is used here to distinguish between a warning generated in the presence of shear (although the shear is below the hazardous strength level) and a genuine false alarm generated in the absence of any true windshear. The path-based scoring method does not provide any distinction between these two different types of false outputs, and the new terminology is employed to emphasize this point.

Note that these statistics are best interpreted something like: "POD = the probability that a pilot encountering a microburst in the analysis area would receive a warning." There is no concept of microburst event or observation in this method; it makes no statement about the *number of microbursts* present, much less how many of them are detected.

The loss estimate for a path is based on the integrated, thresholded shear of the longitudinal wind. The steps involved in the computation are:

1) The windfield is used to calculate the longitudinal windspeed along the path, from which a shear profile is computed (using simple point-to-point differencing).

2) The shear profile is then thresholded against a nominal shear value (initially suggested to be 2.5 m/s per km) and set to zero where the shear values do not exceed the threshold.

3) The thresholded shear profile is then numerically integrated to obtain a windspeed loss estimate across the path.

The loss estimate described above is intended to represent (albeit crudely) the loss which might be experienced by a pilot when flying along the path. While the calculation is quite simplistic, it has three important characteristics. First, it provides at least some sort of location- and direction-sensitive hazard estimate (not true of region outlines). Second, it meets the desired criteria that the hazard estimate is monotonically increasing as a specific path is extended (not true of simple endpoint-to-endpoint velocity differences). Third, it does not include contributions from shear levels which are weak enough to be compensated for by nominal manual or automatic flight control inputs during landing or takeoff operations.

This last criteria is quite significant and has been an issue of some concern to the aviation weather community. Studies of aircraft performance in windshear have focused on the total energy loss rate (sometimes called "F-factor") experienced by an aircraft during microburst penetration. This energy loss includes contributions from two sources: (a) the divergence of the longitudinal winds along the flight path (roughly corresponding to the horizontal divergence measured by a radar) and (b) the downdraft winds along the flightpath. Aircraft control systems are capable of stabilizing aircraft flight profiles in the presence of F-factors up to some aircraft-specific limit, and windshears below this magnitude should not result in "microburst" alerts. The loss estimate described above is a form of estimate of the integrated F-factor along the flight path, although it includes only the horizontal component of the windshear impact. Only shears above a certain threshold are included in this integration in an attempt to match the response characteristics of an aircraft more closely.

While the loss estimate used here is at best a very crude indicator of actual hazard level to an aircraft, it attempts to be a more faithful representation than those metrics (e.g., maximum velocity differential) used in previous evaluations. Further study in the relationship between F-factor and aircraft performance (and the relationship between longitudinal shear and downdraft strength) should provide insight into how this loss estimate may be improved.

5.6. Example of Path-Based Scoring Applied to July 11, 1988 Case

The behavior of the path-based scoring approach can best be understood through a few examples of its application to an actual microburst case. The July 11, 1988 microburst described in Chapter 4 is used here to examine the different aspects of this scoring method.

With the path-based scoring approach, the set of paths may be chosen to focus attention on a specific area of operation. The color plots in Figure 29, Figure 30, Figure 31, Figure 32 and Figure 33 show the "true" and "detected" path strengths over the east-west

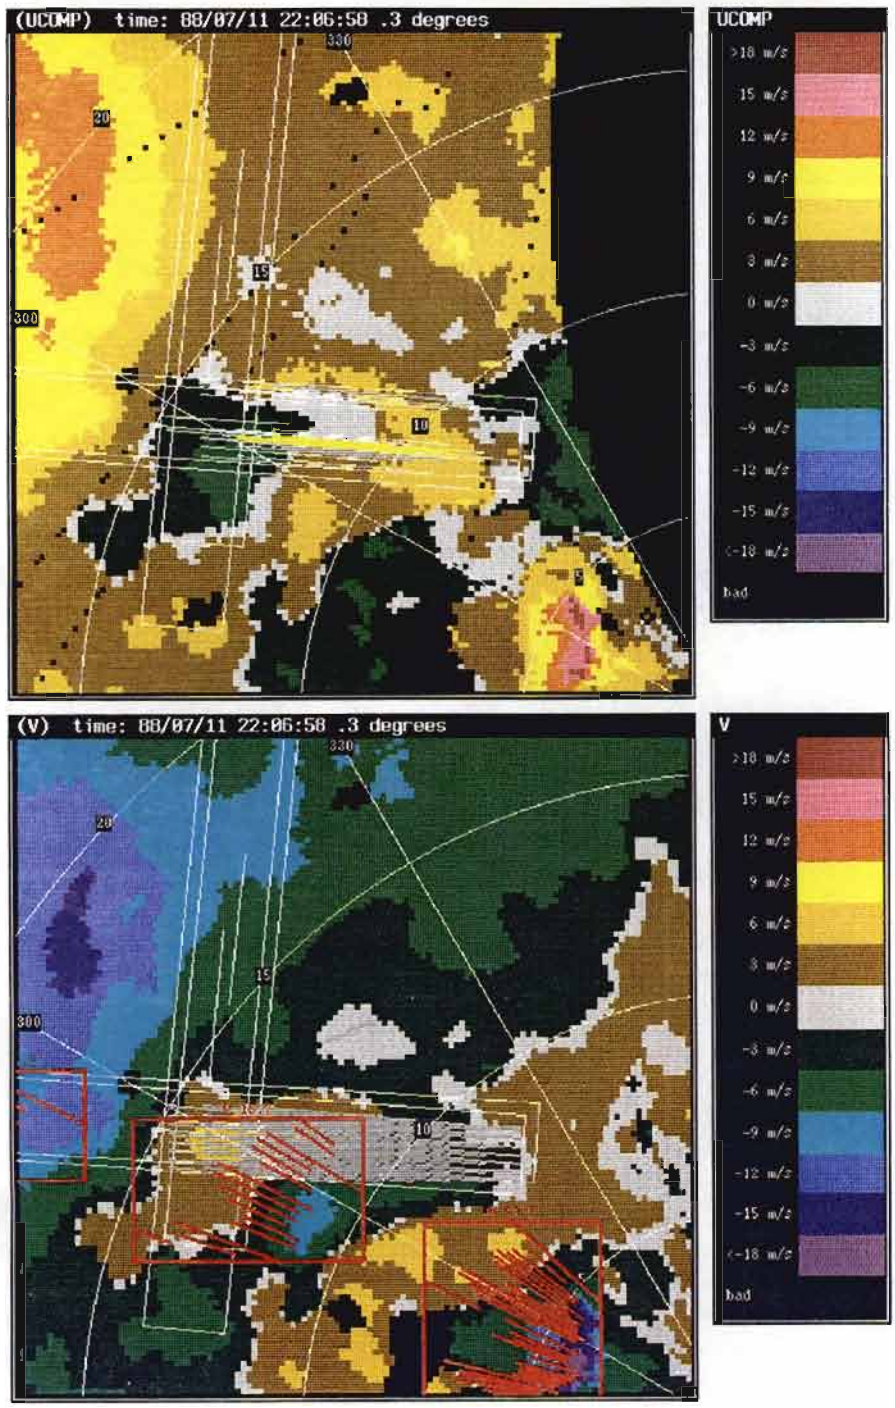

Figure 29: Path-based scoring results for July 11, 1988 microburst at 22:06:58. Upper image shows east-west component of the dual-doppler windfield; lower image shows radar-measured velocity field. Runway alert boxes, microburst alarms and scoring paths are overlaid.

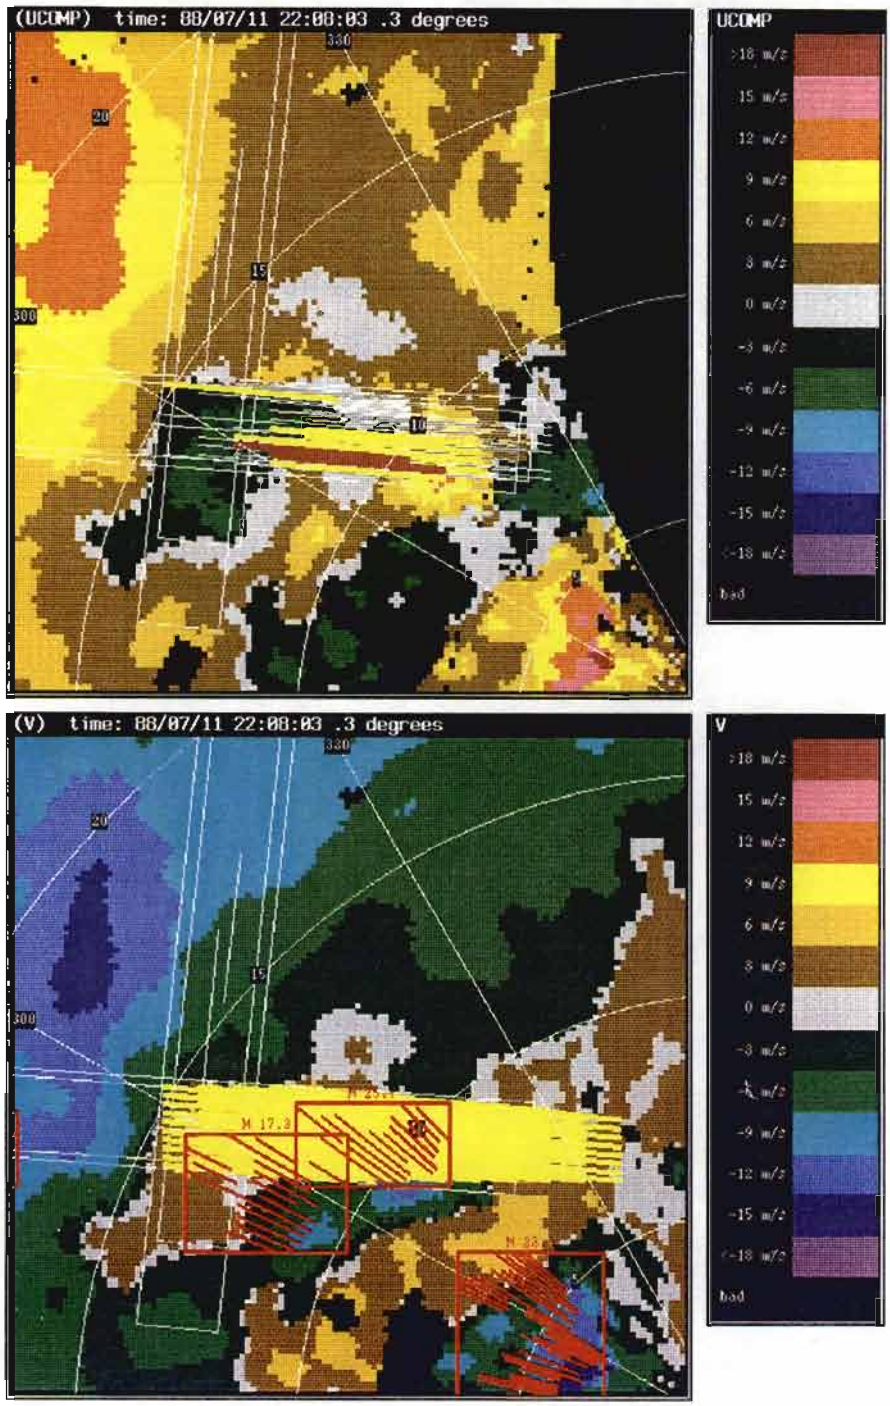

Figure 30: Path-based scoring example for 22:08:03.

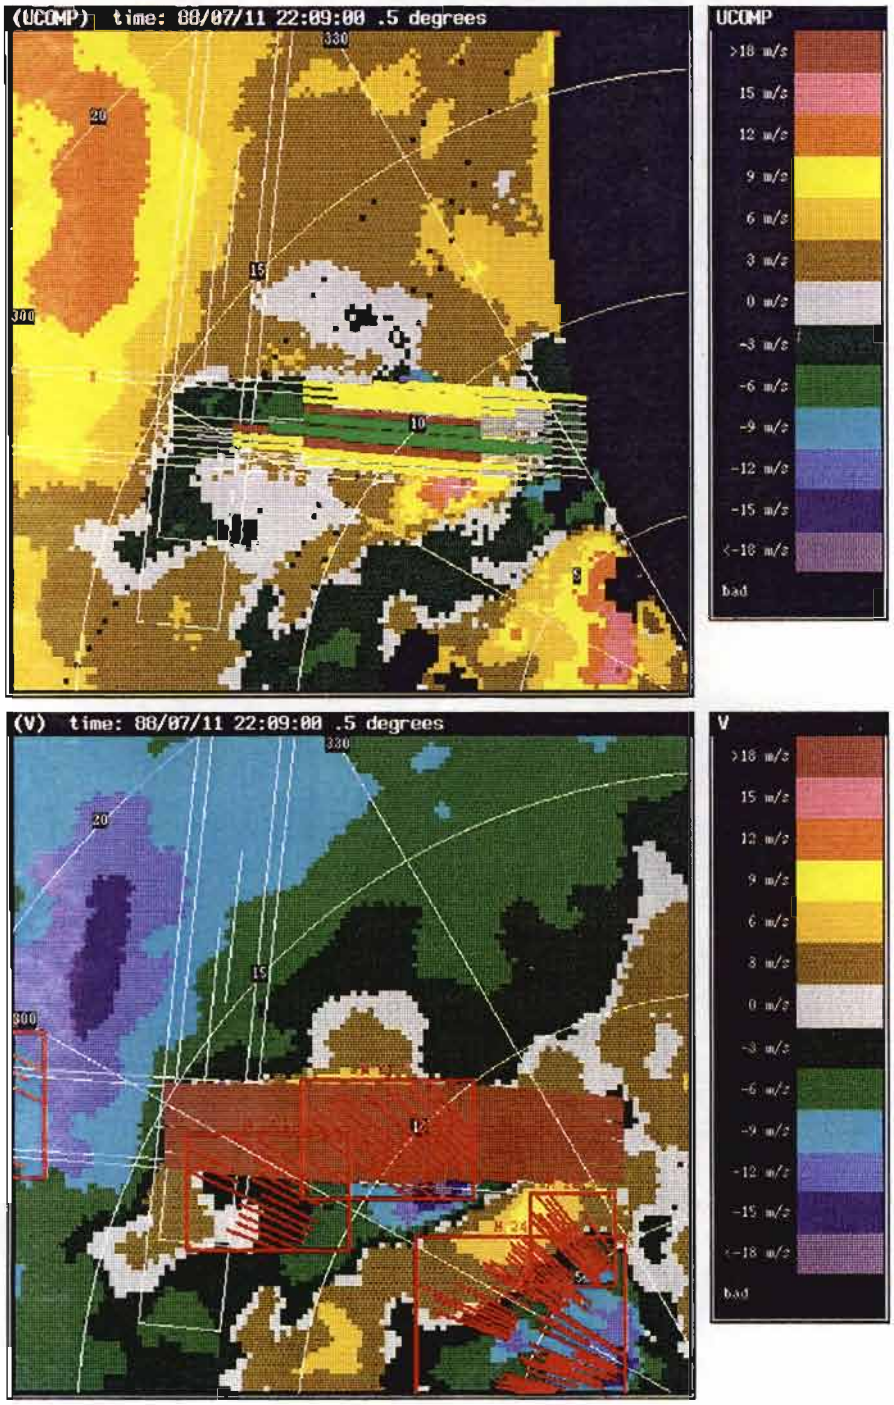

Figure 31: Path–based scoring example for 22:09:00.

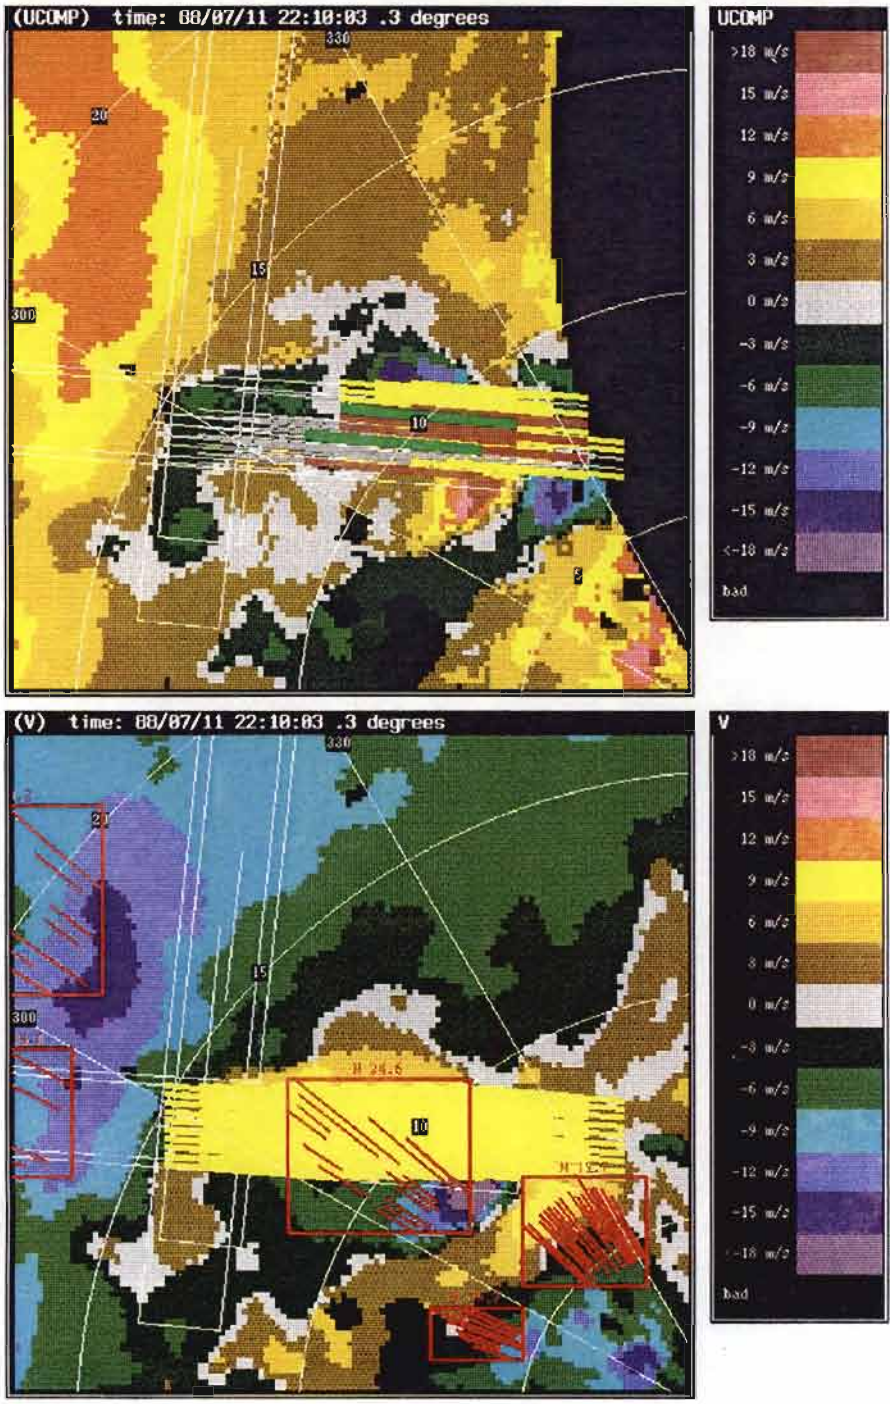

Figure 32: Path-based scoring example for 22:10:03.

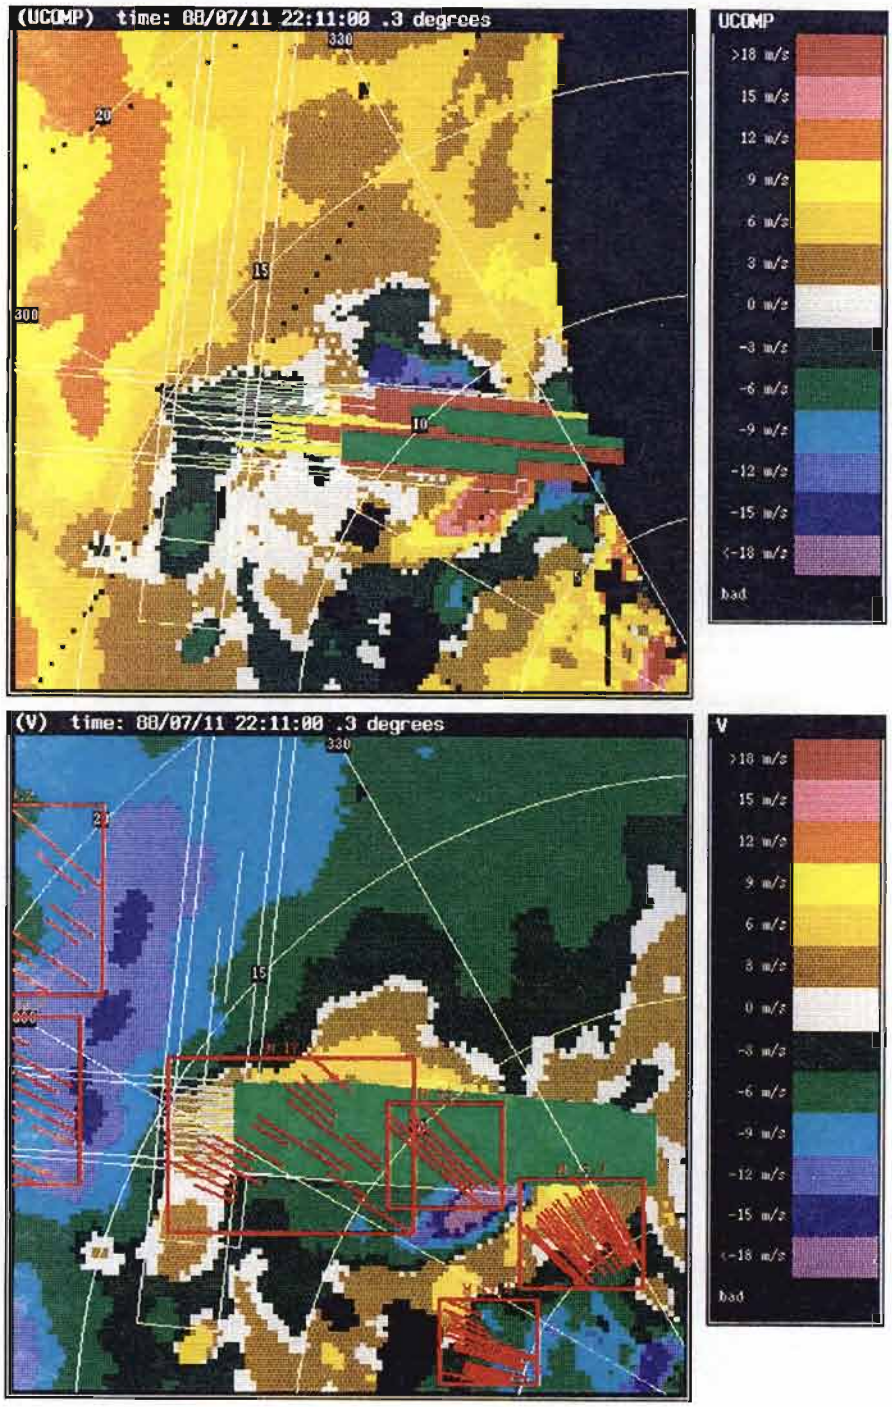

Figure 33: Path–based scoring example for 22:11:00.

runway, which was actually in use for approaches at the time of the microburst. A 10-by-10 grid of paths was used for this example, with each path 4 km long and spaced every 800 meters in the east-west direction and every 200 meters in the north-south direction. All paths are in the same direction, oriented along the runway. In each figure, the upper image shows the U component of the reconstructed dual-doppler windfield (the component of the horizontal wind field in the east-west direction, with winds to the east being positive) and the lower image shows the radar-measured surface velocity field. The airport runways and alert boxes are indicated in the white overlay on each figure. The alert boxes are obtained by drawing a rectangle around each runway, with a 1 nmi buffer on each side of the runway and a 3 nmi buffer on the approach and departure ends of the runway. The path set used in this example is chosen to fill the arrival alert box for the runway which was active at the time of the microburst.

The paths used for scoring are also shown on the color plots, color coded by their strength. The paths shown on the upper image correspond to the "true" strength determined from the dual-doppler windfield along the path. The paths in the lower image correspond to the "detected" strength, determined by the strength of the strongest microburst alarm which the path intersects. The color codes for the paths are shown in Table 4. The lower image also has the microburst alarms (and corresponding shear segments) overlaid in red.

The tabulated results for this scoring example are listed in Table 5. The table lists the number of valid paths for each tilt time (the number of paths for which the dual-doppler windfield data was valid along the path), the number of paths considered hazardous (true strength greater than 10 m/s), the total number of paths for which microburst alarms were issued (alarm strength greater than 10 m/s) and the number of hazardous paths which were also alarmed. The POD and PFW statistics are computed for each tilt and totalled for the entire case. For these five radar scans, only five paths which were considered hazardous were not alarmed by the algorithm (all on the first scan at 22:06:58). The resulting POD is 98 percent. However, 130 of the 372 paths for which alarms were generated were not in fact hazardous, yielding a PFW of 35 percent.

Table 4.
Color scale used to indicate strength of scoring paths.

Color	Strength (m/s)
White	10 - 14
Grey	15 - 19
Yellow	20 - 24
Orange	25 - 29
Green	30 - 34
Red	35 and above

Table 5.
Path-based scoring results for east-west runway case.

Tilt time: 07/11/88 22:06:58		Elevation 0.300000			
Number of valid paths	85				
Number of hazards	20	Hit	15	POD	75.0
Number of alarms	45	False	30	PFW	66.7
Tilt time: 07/11/88 22:08:03		Elevation 0.300000			
Number of valid paths	84				
Number of hazards	45	Hit	45	POD	100.0
Number of alarms	79	False	34	PFW	43.0
Tilt time: 07/11/88 22:09:00		Elevation 0.500000			
Number of valid paths	82				
Number of hazards	55	Hit	55	POD	100.0
Number of alarms	82	False	27	PFW	32.9
Tilt time: 07/11/88 22:10:03		Elevation 0.300000			
Number of valid paths	85				
Number of hazards	62	Hit	62	POD	100.0
Number of alarms	85	False	23	PFW	27.1
Tilt time: 07/11/88 22:11:00		Elevation 0.300000			
Number of valid paths	81				
Number of hazards	65	Hit	65	POD	100.0
Number of alarms	81	False	16	PFW	19.8
MISSION STATS:					
Number of valid paths	417	Number of tilts			5
Number of hazards	247	Hit	242	POD	98.0
Number of alarms	372	False	130	PFW	34.9

In this case it is clear that pilots arriving from the east onto this runway would have been adequately warned of the microburst hazards which were present. Many of the paths for which alarms were generated did not actually penetrate significant shear, but this high false warning value is an expected consequence of the path-based scoring approach. The variation in the microburst strength across the runway alert region cannot be represented by the microburst alarm (which has a single strength value for its entire area), and some intersecting paths will receive strength values higher than those which would actually be encountered.

The same case may also be scored from the perspective of aircraft approaching from the south, landing on the north-south runway. During this microburst on July 11, aircraft were not using this secondary approach pattern, but this example will illustrate the variety of evaluations possible with the path-based approach. For this second example, a 10-by-10

grid of paths was again used, but the paths were aligned with the north–south runways and spaced every 200 meters in the east–west directions and every 600 meters in the north–south direction. The true and alarmed paths for one scan from this case are shown in Figure 34, illustrating only a small number of actual hazard paths on the eastern edge of the path set. The detected shear region extends across the full width of the alert region for this runway, and many of the paths are thus alerted. The scoring summary for this case is shown in Table 6, indicating a very high probability of false warning rate (85 percent). This result is the consequence of the aspect–dependence of the microburst strength and the single strength value used for the alarm. The longitudinal shear for paths along the north–south runway are quite small, while the shears are somewhat stronger in the direction of the radar. The angle between the runway direction and the radar viewing angle is quite significant (over 45°) in this region. Also, the strength value for the alarm in this region is the maximum strength of all segments in the alarm, which extends far to the east of the runway alert box. It is likely that the alarmed strength for the north–south paths is actually based on shear segments located by the algorithm in a region well outside of the north–south alert box but is applied to the entire alarm area.

As a final example of the path–based scoring method, a set of paths in all directions was used for this case. The path set used this time was a 20–by–20 grid of 4–km–long paths, spaced every 300 meters in both directions. Nine paths are used at each grid point, spaced 20° apart in orientation angle. The resulting path strengths for one radar scan is shown in Figure 35. The dual–doppler–based path strengths shown in the upper image of this figure illustrate the aspect dependence of the outflow strength. The strongest paths are located in the east–west direction, and very few strong paths are found in the north–south direction. The summary scoring statistics for this version are listed in Table 7, showing a high POD (92 percent) and a PFW just over 60 percent.

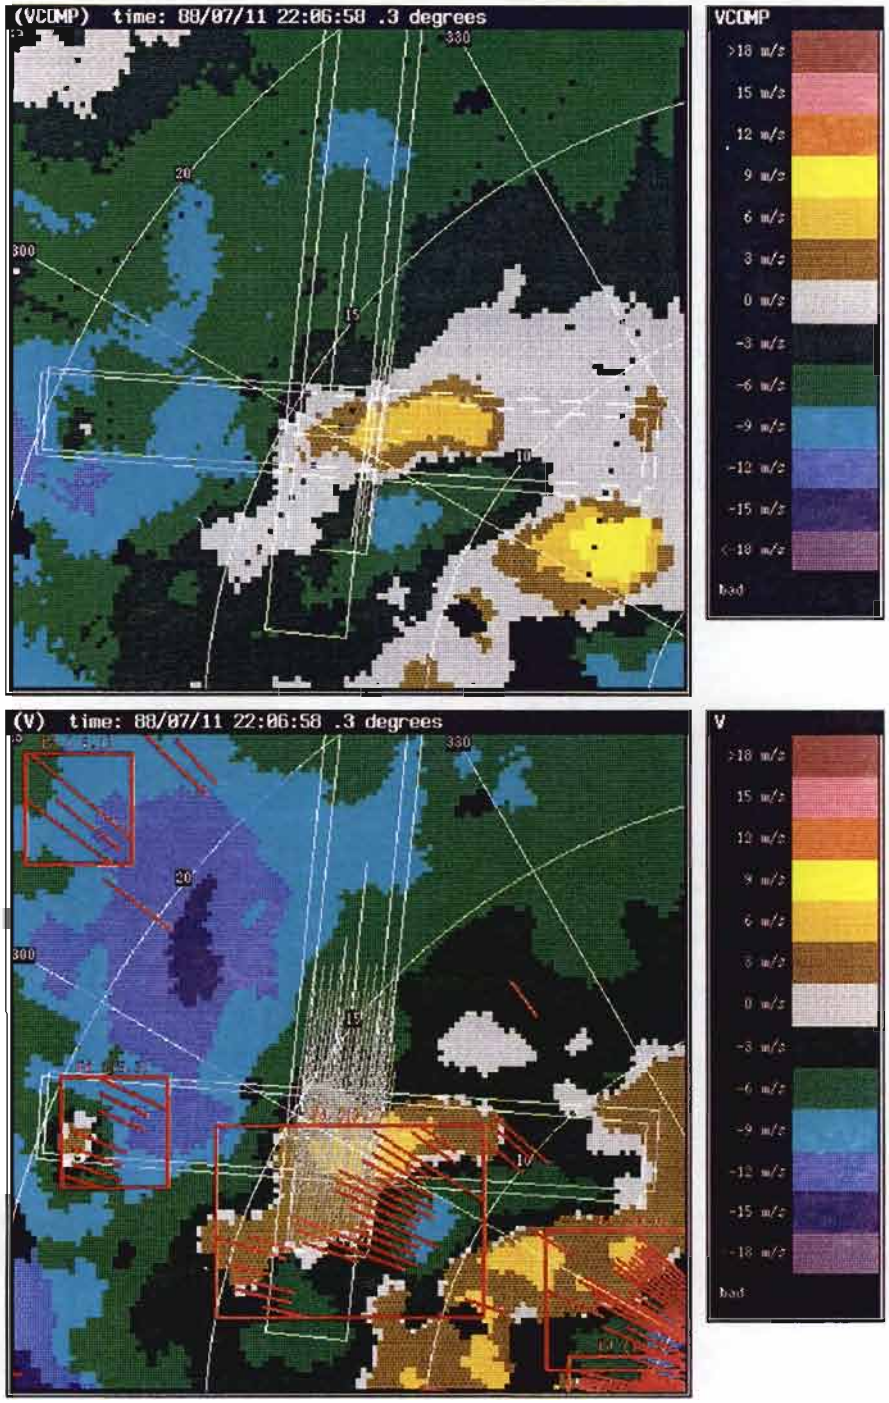

Figure 34: Path-based scoring using north-south oriented paths, for 22:06:58.

Table 6.
Path-based scoring results for north-south runway case.

Tilt time: 07/11/88 22:06:58		Elevation 0.300000			
Number of valid paths	97				
Number of hazards	5	Hit	5	POD	100.0
Number of alarms	47	False	42	PFW	89.4
Tilt time: 07/11/88 22:08:03		Elevation 0.300000			
Number of valid paths	100				
Number of hazards	12	Hit	12	POD	100.0
Number of alarms	39	False	27	PFW	69.2
Tilt time: 07/11/88 22:09:00		Elevation 0.500000			
Number of valid paths	100				
Number of hazards	6	Hit	6	POD	100.0
Number of alarms	39	False	33	PFW	84.6
Tilt time: 07/11/88 22:10:03		Elevation 0.300000			
Number of valid paths	100				
Number of hazards	4	Hit	0	POD	0.0
Number of alarms	0	False	0	PFW	0.0
Tilt time: 07/11/88 22:11:00		Elevation 0.300000			
Number of valid paths	99				
Number of hazards	6	Hit	6	POD	100.0
Number of alarms	68	False	62	PFW	91.2
MISSION STATS:					
Number of valid paths	496	Number of tilts			5
Number of hazards	33	Hit	29	POD	87.9
Number of alarms	193	False	164	PFW	85.0

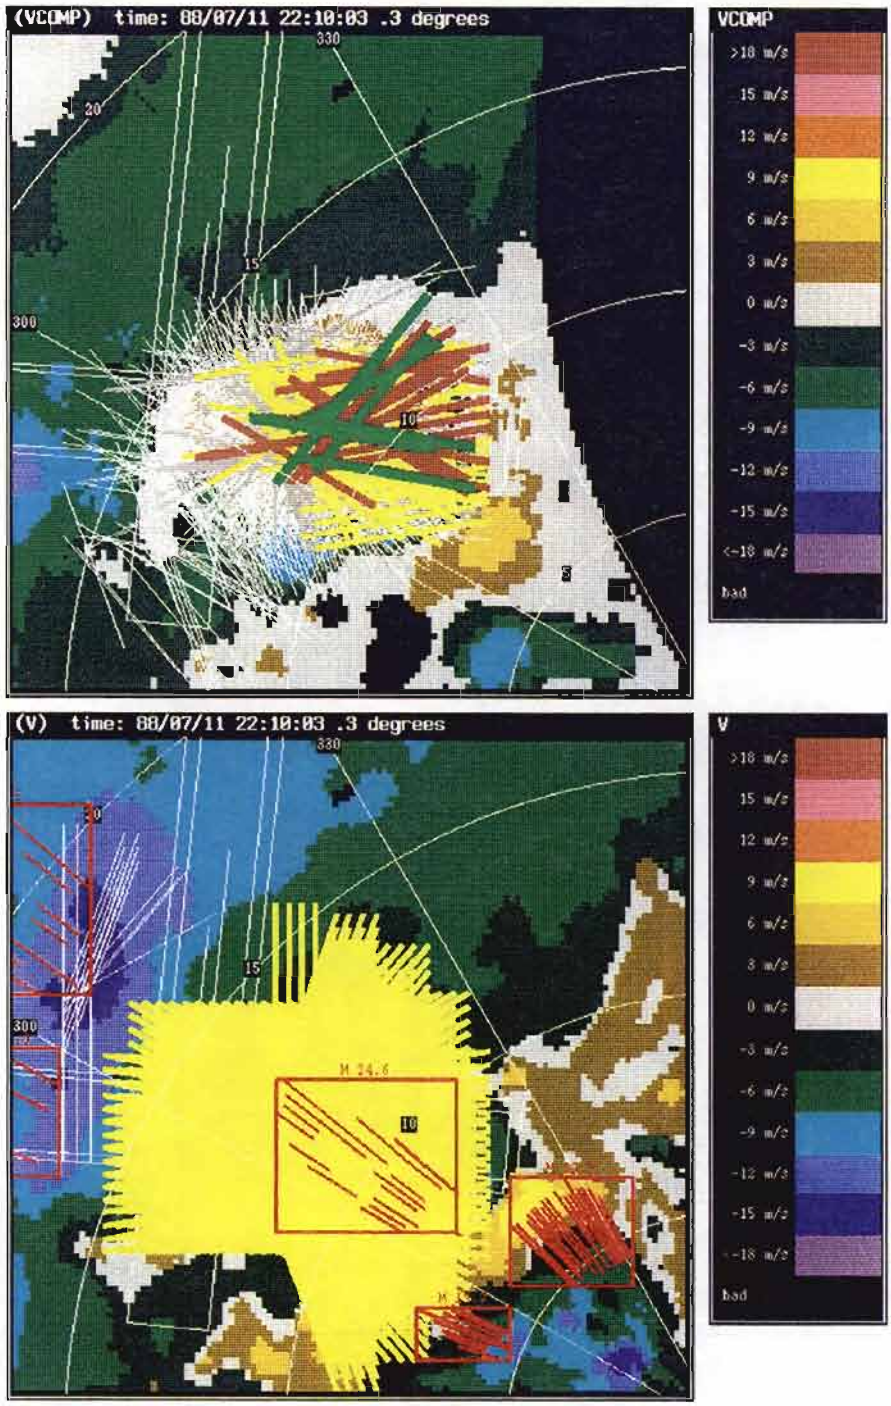

Figure 35: Scoring with paths in many directions, for the radar scan at 22:10:03.

Table 7.
Path-based scoring results for the multi-directional case.

Tilt time: 07/11/88 22:06:58		Elevation 0.300000			
Number of valid paths	3426				
Number of hazards	705	Hit	679	POD	96.3
Number of alarms	2240	False	1561	PFW	69.7
Tilt time: 07/11/88 22:08:03		Elevation 0.300000			
Number of valid paths	3530				
Number of hazards	964	Hit	912	POD	94.6
Number of alarms	1956	False	1044	PFW	53.4
Tilt time: 07/11/88 22:09:00		Elevation 0.500000			
Number of valid paths	3592				
Number of hazards	795	Hit	768	POD	96.6
Number of alarms	2105	False	1337	PFW	63.5
Tilt time: 07/11/88 22:10:03		Elevation 0.300000			
Number of valid paths	3493				
Number of hazards	815	Hit	646	POD	79.3
Number of alarms	1391	False	745	PFW	53.6
Tilt time: 07/11/88 22:11:00		Elevation 0.300000			
Number of valid paths	3416				
Number of hazards	946	Hit	892	POD	94.3
Number of alarms	2560	False	1668	PFW	65.2
MISSION STATS:					
Number of valid paths	17457	Number of tilts			5
Number of hazards	4225	Hit	3897	POD	92.2
Number of alarms	10252	False	6355	PFW	62.0

6. IMPACT OF ASYMMETRY ON TDWR PERFORMANCE

The outflow region of a microburst is often asymmetric in the sense that the airspeed loss encountered by an aircraft (or the velocity difference measured by a radar) will be stronger in some directions than in others. As a result, TDWR measurements of microburst strength (which are obtained from measurements in a single viewing direction) may not match the strength encountered by an aircraft whose flight path is not directly aligned with the radar viewing direction. At most airport locations it will be impossible to site a single TDWR radar to view along all the commonly used flight paths. The possibility that a (single radar) TDWR might significantly underestimate the actual hazard posed by an asymmetric microburst was first presented by [Wilson, et al., 1984] and has been the motivation for several subsequent studies. The characterization of microburst asymmetry and it's impact on operational systems has been the topic of several studies, including [Eilts, 1988] and [Hallowell, 1990]. In the study by Hallowell, roughly 88 radar scans through microbursts of various strengths and sizes were examined and the detailed distributions of microburst strength vs. viewing angle were computed. The results of the Hallowell study indicate that microbursts are typically asymmetric, with an average ratio of maximum strength to minimum strength of 1.9. The collection of microbursts used for this study were all obtained from measurements made in Denver, CO, and the resulting analysis provided in this chapter should be considered specific to the Denver (and high plains) meteorological environment. Ongoing studies of asymmetry characteristics in other geographical locales will be necessary to determine the applicability of these results to a broader range of environments.

Several studies have been conducted to assess the detection capability of the TDWR microburst algorithm ([Campbell, 1989] and [Evans, 1990]). A consistent result among these studies is a strong dependence of POD on the measured strength of the microburst outflow. An example of this sensitivity, taken from the study of performance in Denver during 1988, is shown in Figure 36.

Since the probability of detecting a microburst (with a TDWR) is quite sensitive to the (measured) microburst strength, the variation of strength with viewing angle may have a significant effect on the TDWR performance.

This variation in apparent microburst strength with viewing angle reduces the ability of the TDWR to detect those microbursts which exhibit significant asymmetry. While the TDWR system will detect microbursts with observed velocity differentials above 15 m/s with a high probability, the detection probability will be significantly reduced if the observed strength is significantly reduced.

6.1. Statistical Model for Microburst Asymmetry

The quantitative effect of asymmetry on the probability of detection may be assessed by a careful examination of the statistical reduction in observed microburst strength resulting from viewing angle dependence. The performance analysis presented here is based on the following assumptions:

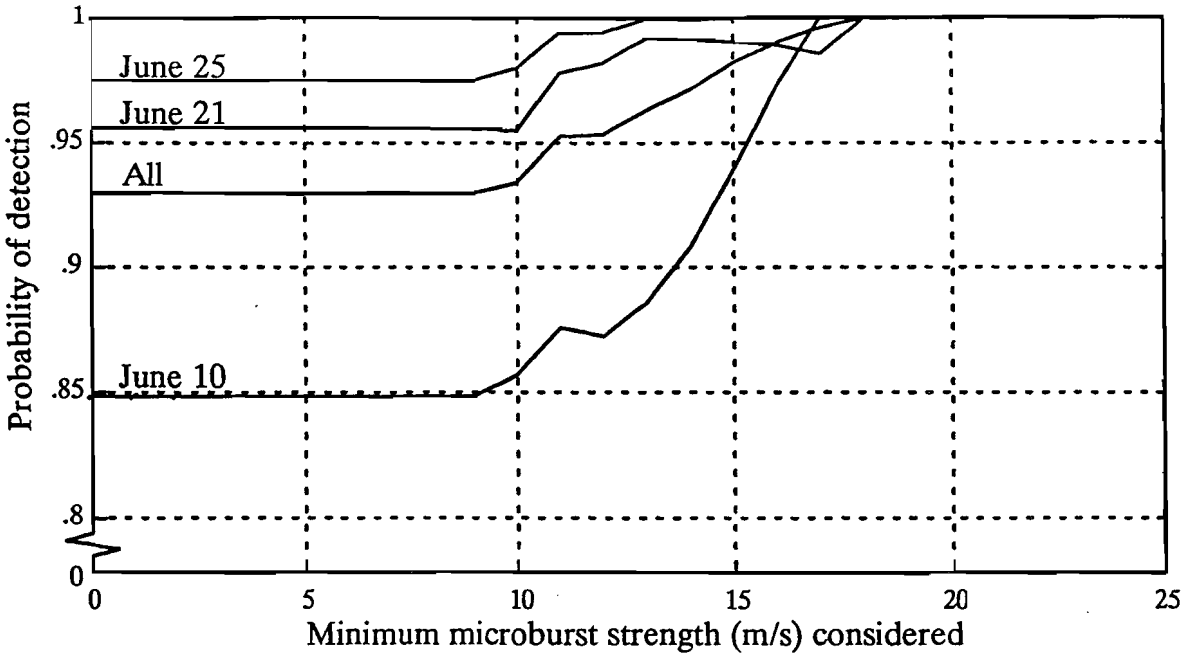

Figure 36: Probability of detecting all divergence regions above a certain strength level. The aggregate detection rate drops as the minimum included strength drops, indicating poorer detection rates for weaker strengths. Detection rates at the left of the plot indicate aggregate POD over all events for each category. Performance analysis from Denver, 1988 cases.

1) The observed strength of a microburst as a function of viewing angle may be modeled as an ellipse, i.e.:

$$R(\theta,\alpha) = \left[(\alpha^2-1)\sin^2(\theta)+1\right]^{-\frac{1}{2}} \qquad (14)$$

where: α = ratio of max/min strength
θ = viewing angle relative to direction of maximum strength
R = ratio of observed strength to maximum strength

2) The statistical distribution of the ratio of maximum to minimum strength for a microburst (over all strengths ranges) may be modeled as a single Rayleigh distribution, i.e.:

$$p_\alpha(\alpha) = \frac{2(\alpha-1)}{m^2} e^{-\left(\frac{\alpha-1}{m}\right)^2} \qquad (15)$$

where: m = mean ratio value

The validity of these two assumptions, for the Denver cases studied here, will be demonstrated below using the asymmetry data obtained by Hallowell. First, the use of these assumptions to model observed strength distributions will be addressed.

If microbursts were always observed from the worst-case direction, i.e., along the viewing direction with the minimum strength, then the observed strength would be scaled down from the true strength by a factor equal to the asymmetry ratio for that event. Using the assumed Rayleigh distribution for the asymmetry ratios, Figure 37 shows the cumulative distribution of attenuation (defined to be the ratio of observed strength to true maximum strength) which would be experienced in this worst-case condition. The curve labelled "Worst case viewing angle" shows the probability that the ratio of the measured strength (minimum strength) to the maximum strength is at least as large as the corresponding ordinate value. This curve corresponds to a Rayleigh mean parameter m = 2.0 and indicates that median value for the attenuation factor is 0.35.

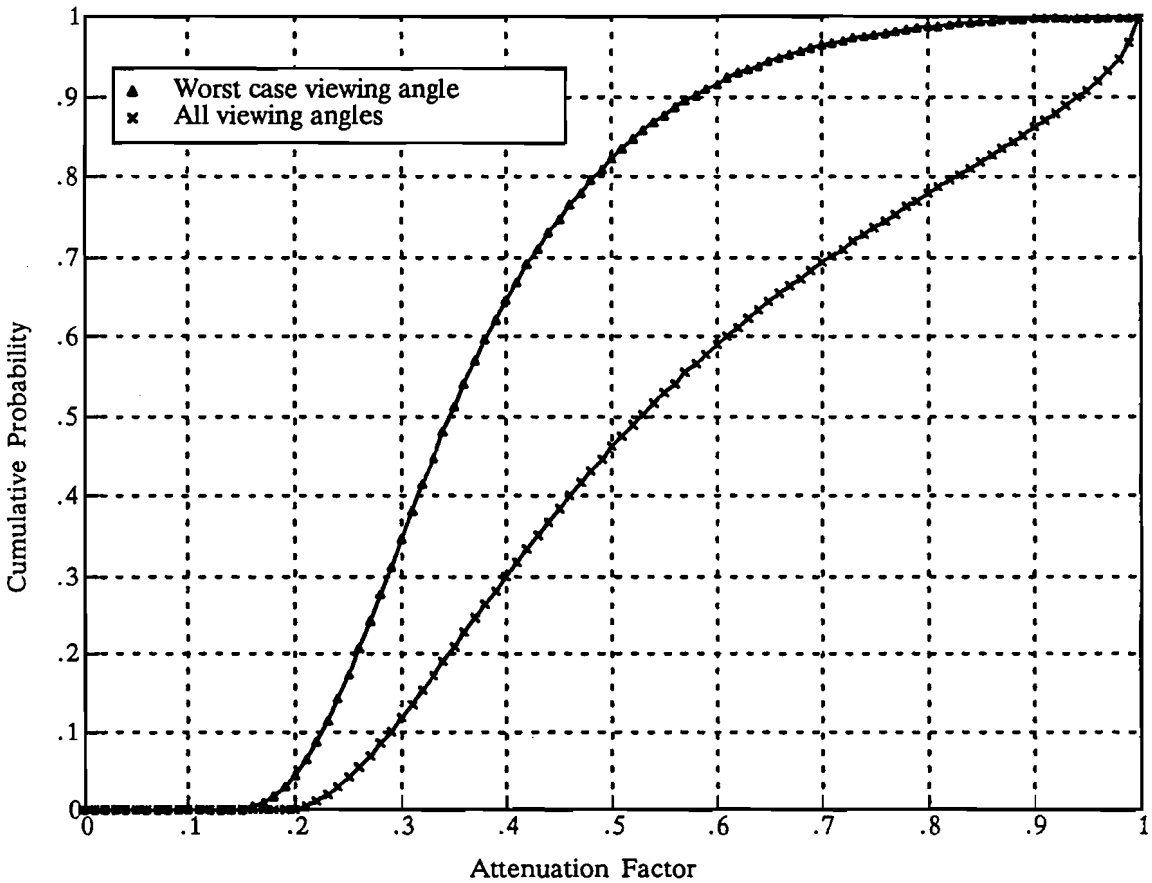

Figure 37: Cumulative probability distribution for attenuation factor. Two cases are shown: a worst case, assuming the radar is always viewing from the minimum strength direction, and a realistic case where the viewing direction is uniformly distributed.

Since the radar viewing angle is assumed to be randomly oriented with respect to the direction of maximum outflow strength (as indicated by the Hallowell study), this worst-case viewing direction case is unrealistic. If the viewing angle is assumed to be uniformly distributed and the strength is assumed to be an elliptical function of the viewing angle (as in Equation 1), the attenuation factor distribution is somewhat improved as shown by the curve labelled "All viewing angles" in Figure 37. For this more realistic case, the median attenuation factor is now 0.52.

The cumulative distribution for the attenuation factor is a very powerful tool that allows the distribution for the observed strength to be computed if the maximum strength is known. The observational study by Hallowell examined a number of microburst cases using dual-doppler radar measurements and computed the microburst strength at each of 18 viewing directions for each event. Using these observations, the distributions of both observed and maximum microburst strengths may be computed, as shown in Figure 38. The "Maximum" curve corresponds to the cumulative probability distribution of the maximum strength of an event, while the "Observed" curve depicts the cumulative probability distribution for the observed strength (over all angles and events).

Using the attenuation distribution from Figure 37 (using all viewing angles) and the distribution for maximum event strengths from Figure 38, the distribution for observed strengths may be predicted. This predicted distribution is also shown in Figure 38 and is seen to match the actual observed distribution extremely well. This excellent match between

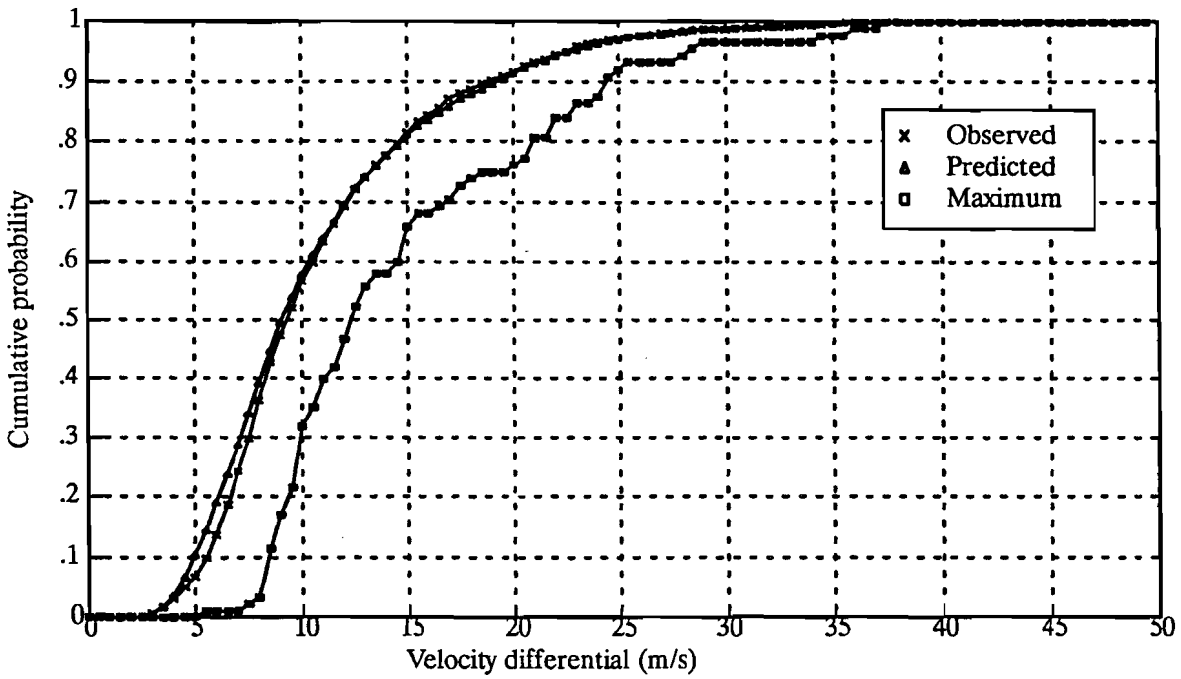

Figure 38: Cumulative distributions for observed microburst strengths in Denver cases (over all viewing angles) compared with the distribution predicted from the asymmetry model and the distribution of maximum velocities.

the actual and predicted strength distributions indicates that the asymmetry model employed (combining the elliptical strength variation with viewing angle and the Rayleigh distribution of asymmetry ratios) accurately predicts the natural relationship between the maximum event strength and the observed strengths from all angles.

6.2. Detection Performance vs Maximum Outflow Strength

Having validated the attenuation distribution model, it may now be used to examine the impact of asymmetry on TDWR detection performance. The probability of microburst detection is shown in Figure 39 as a function of radar–observed microburst strength, based on the performance analysis reported in [Campbell, Merritt and DiStefano, 1988]. To depict the effects of asymmetry, the performance as a function of actual maximum strength (as opposed to observed strength) is plotted on Figure 39 as well. The detection rate for each maximum strength value is computed by integrating the detection rate over all possible

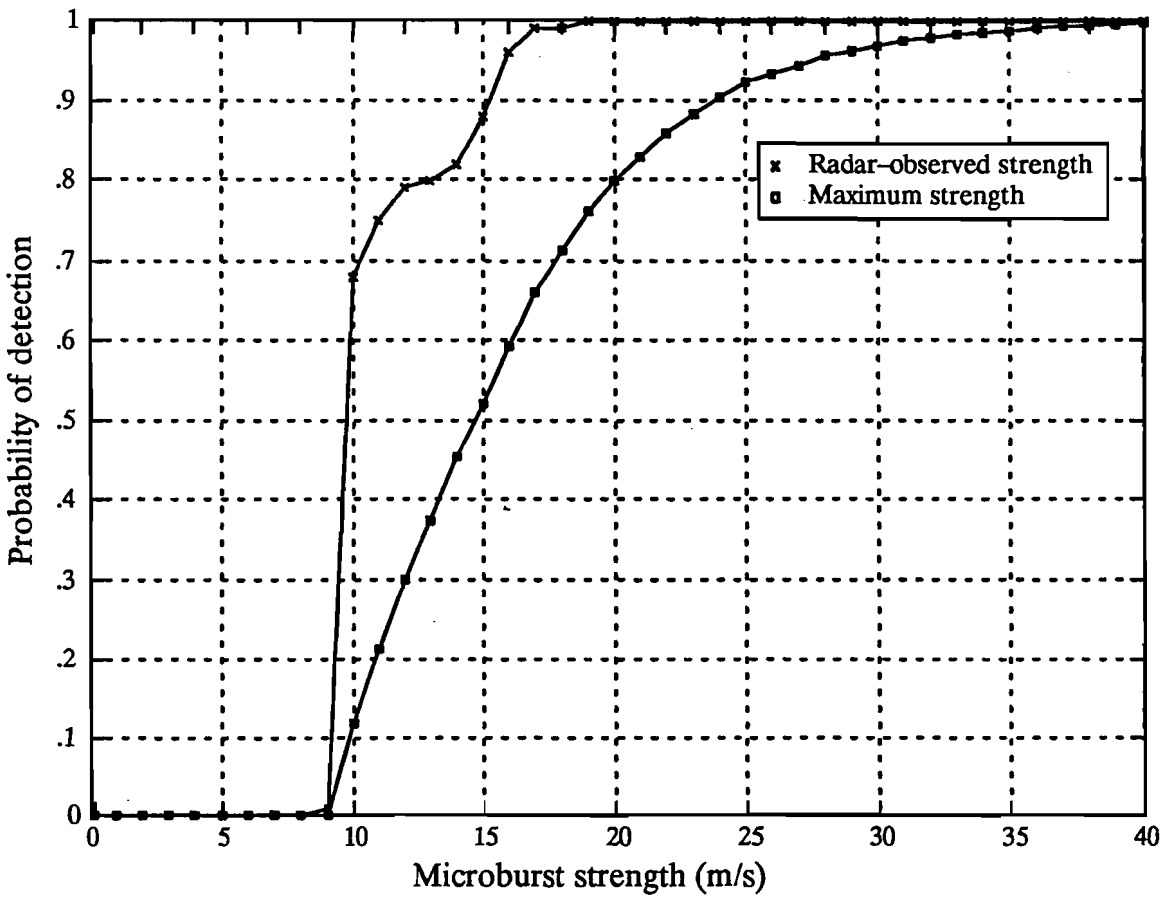

Figure 39: Probability of microburst detection (based on Denver, 1988 cases) as a function of (a) radar–observed strength and (b) as a function of maximum strength in any direction.

attenuated (observed) strengths, in proportion to the probability distribution from Figure 38.

The TDWR performance evaluations reported to date suggest that all microbursts with strengths above 15 m/s were detected with a probability of at least 90 percent. Based on the modelled distribution of attenuation from asymmetry, the revised performance curves suggest that the 90-percent-detection level is not obtained until microburst maximum strengths reach about 24 m/s.

It is important to bear in mind that this analysis of the impact of asymmetry is rather sensitive to the asymmetry characteristics obtained in the Hallowell study, which were limited to experience in the Denver area. Likewise, the detection performance curve used in Figure 39 is specific to Denver. Field programs in Huntsville, AL and Orlando, FL have indicated that the detection performance of the divergence algorithm may be worse in Denver than in other meteorological regions; hence, asymmetry impact estimates based on the Denver performance curves may be overly pessimistic.

If the divergence detection algorithm were improved so as to detect a larger fraction of the weaker microburst events, the detection of stronger asymmetric microbursts would also improve (since these stronger events are sometimes viewed by the radar as being weaker, and hence poorly detected). The plots in Figure 40 illustrate the improvement in POD which could be obtained if the divergence algorithm POD vs radar-observed strength (shown in Figure 39) were improved to yield perfect detection for all events with radar-observed strengths of (a) 10 m/s and above and (b) 7 m/s and above. As shown in the figure, these improvements in the detection of weaker events would significantly improve the detection of microbursts with maximum strengths of all levels.

6.3. Comparison Using Runway-Oriented Microburst Strengths

The analysis in the preceding section examined the probability of detection for microbursts as a function of the maximum strength of the outflow in any direction. Since the radar viewing direction is randomly oriented relative to the direction of this maximum strength, the effect of asymmetry is to reduce the detection performance of the system. The implicit assumption in this analysis is that the classification of a microburst as "hazardous" or "not hazardous" is to be based on the maximum strength of the event in any direction.

Since the hazard of interest to the TDWR system is that posed to an aircraft on final approach or takeoff, it is also relevant to consider the microburst strength in the direction of the aircraft flight (since the most significant shear which contributes to the aircraft hazard at low altitude* is the shear of the longitudinal wind along the flightpath). The aircraft flightpath is also randomly oriented with respect to the direction of the maximum outflow strength. If the radar were always sited so as to measure the winds directly along the flightpath, then the radar-observed winds would match the aircraft-experienced winds and the detection performance would be the same from either perspective. However, since it will generally not be possible to site the TDWR radars to view along all flightpaths at an airport, the runway-oriented and radar-oriented viewing directions will typically be different.

---

* The microburst downdraft velocity, which is proportional to both components of the horizontal shear, decreases to zero near the ground [Targ and Bowles, 1988].

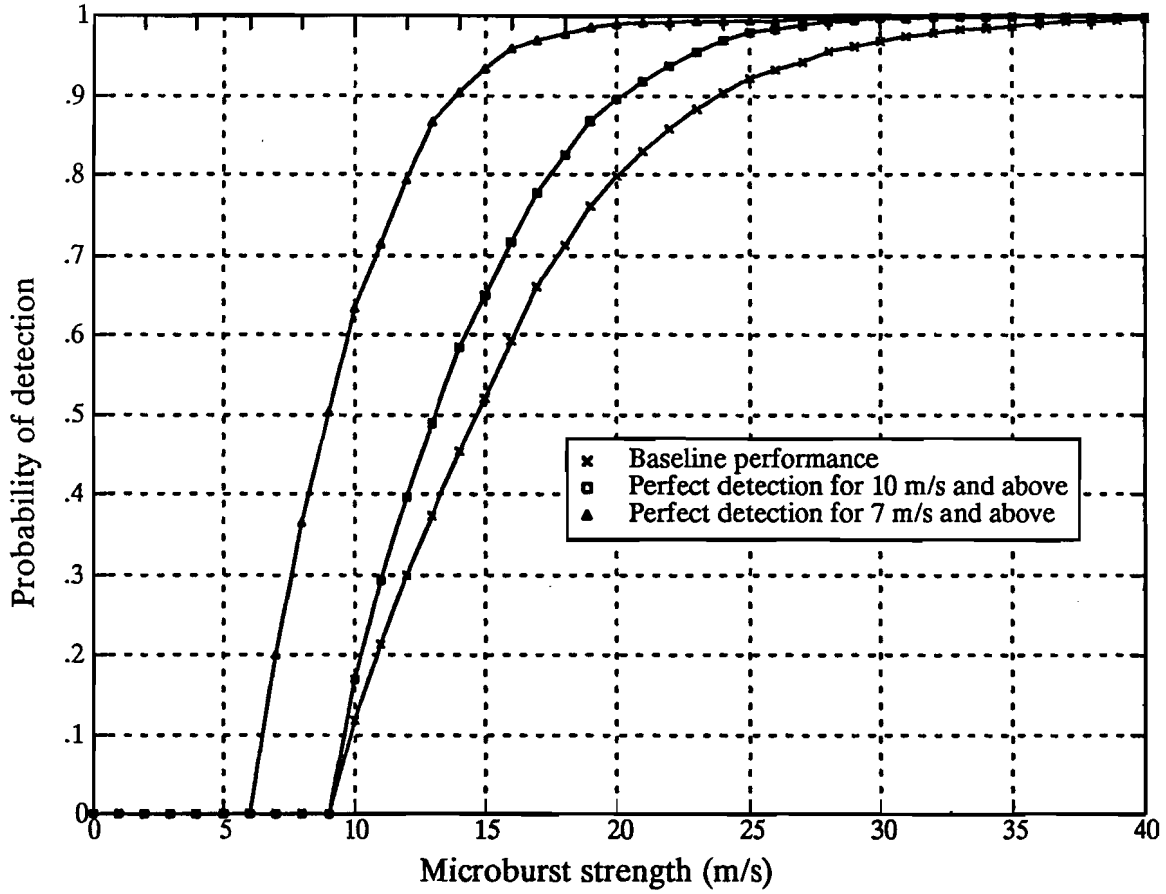

Figure 40: Probability of microburst detection as a function of maximum strength. The bottom curve indicates performance using baseline TDWR divergence algorithm detection performance vs radar-observed strength (from Denver, 1988 cases). The remaining two curves indicate improvements obtained by assuming the divergence algorithm has perfect detection for events of 10 m/s and above and for 7 m/s and above, respectively.

If the microburst maximum strength direction is randomly oriented with respect to both the radar and the runway directions, then the radar- and runway-oriented strengths will be different (for asymmetric microbursts), but will have a zero mean difference. Half of the time the radar-oriented strength will be larger than that along the runway, and half of the time the radar-oriented strength will be less than that along the runway. If the "hazardous" vs "not hazardous" criteria is determined relative to the runway-oriented strength (i.e., that experienced by the aircraft), then the impact of asymmetry will quite different than that computed in the previous section.

The curves in Figure 41 illustrate the distribution of microburst strengths which would be measured by a radar at each of several radar/runway separation angles. For each of these curves, the runway-oriented strength is assumed to be 15 m/s; the orientation of the microburst maximum strength is assumed to be uniformly distributed and the asymmetry ratio is assumed to be Rayleigh distributed as in the previous section. Regardless of the

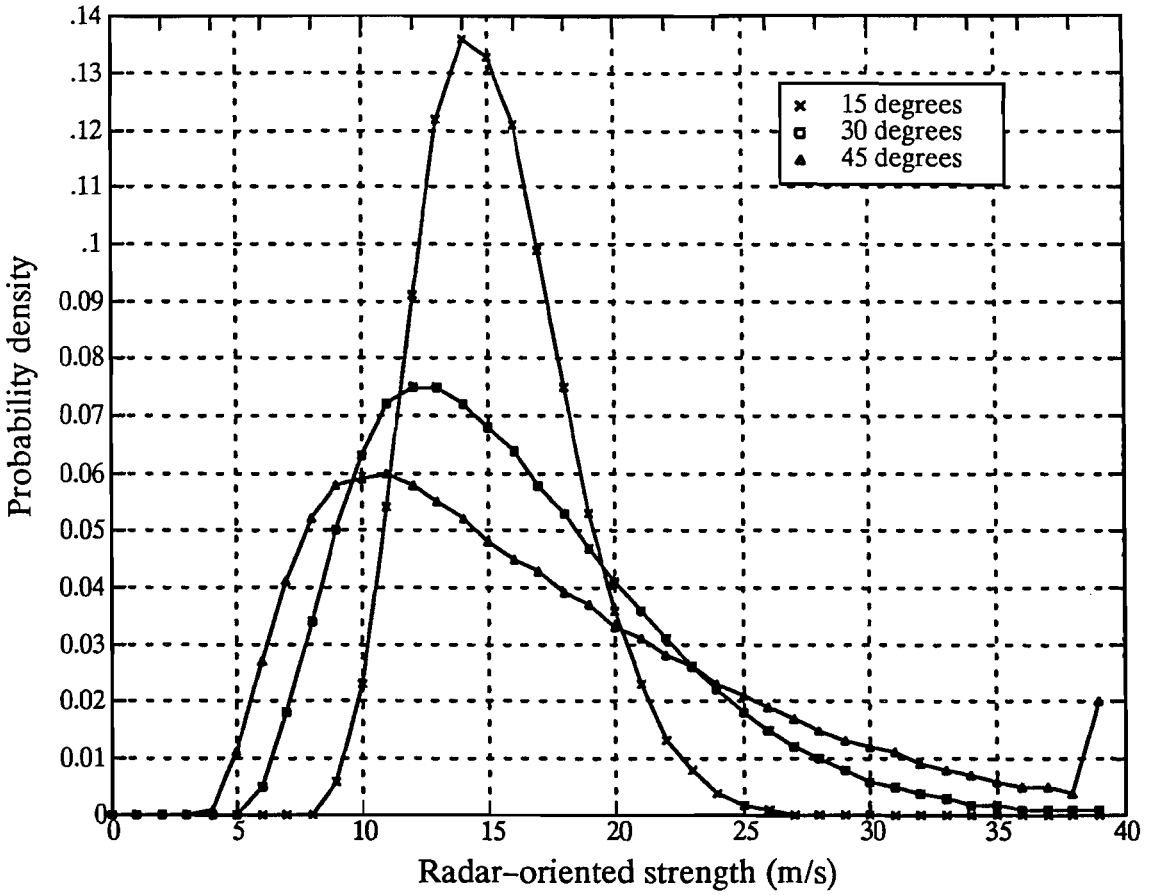

Figure 41: Probability density for radar-oriented microburst strength for a runway-oriented strength of 15 m/s and several separation angles. [Note: Discontinuous end-point for the 45-degree curve (at 39 m/s) represents cumulative probability density for all strengths of 39 m/s and greater.]

angle between the radar viewing direction and the runway, the mean radar-measured strength is always 15 m/s. While the mean radar-measured strength is always 15 m/s, the distribution of radar-measured strengths varies strongly with separation angle. As the separation angle increases, the distribution tends to broaden in width and the peak moves towards lower strengths. Since the probability of detection drops off very rapidly (faster than linear) with decreasing radar-measured strength, the effective POD, as a function of runway-oriented strength, will also decrease with increasing separation angle. This effect is demonstrated in Figure 42, which plots the POD as a function of runway-oriented strength for each of several radar-runway separation angles. The POD curves in this figure were obtained by computing the distribution of radar-oriented strengths for each runway strength level (as in Figure 41) and computing the aggregate POD using the radar-oriented POD statistics (which appear in Figure 42 as the zero-degree separation angle case). The POD curve drops significantly (for strengths above 10 m/s) as the separation angle increases beyond about 15°. Notice that for runway-oriented strengths below 10 m/s, the POD increases with separation angle. This increase in detection rate results from the fact that the

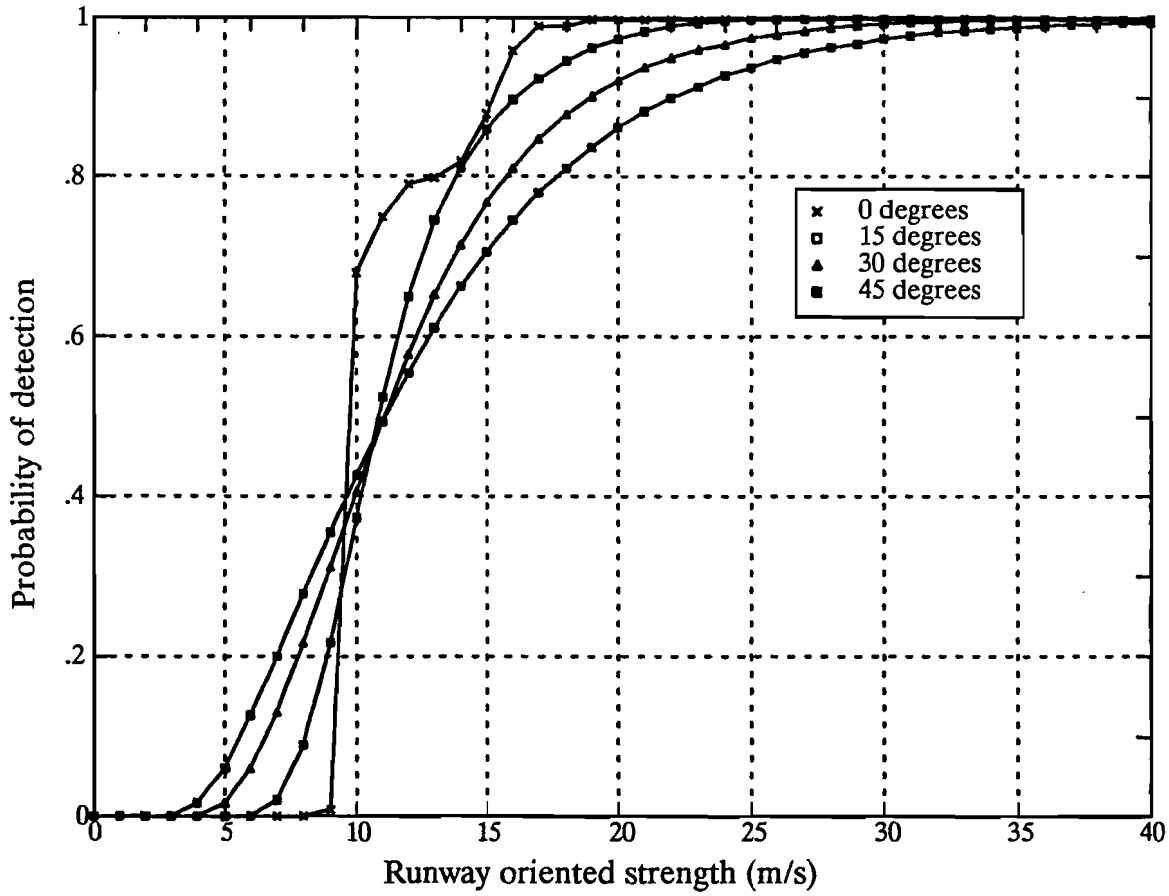

Figure 42: Probability of detection as a function of runway-oriented strength for several radar/runway separation angles, based on Denver 1988 cases. The zero-degree separation angle case corresponds to a radar oriented directly along the runway, and is simply the radar-based POD curve.

radar-oriented strength exceeds the runway-oriented strength half of the time, so that some events which are weak in the runway direction are sensed with a larger strength (and are hence better detected) in the radar direction.

Given the POD as a function of runway-oriented strength, it is possible to compute an aggregate detection rate using an observed distribution of microburst strengths. The observed distribution of radar-oriented strengths should be identical to the distribution of runway-oriented strengths if the microburst orientation is uniformly distributed. The distribution used here is taken from the Denver radar-mesonet comparison study by [DiStefano and Clark, 1990] and is shown for reference in Figure 43. The aggregate POD is simply the convolution of the POD vs strength (as in Figure 42) with the distribution of event strengths. This aggregate POD is shown as a function of separation angle in Figure 44. This figure also illustrates the effect on total POD of an increase in detection of weak outflow events. As in the previous section, two examples of performance improvement are

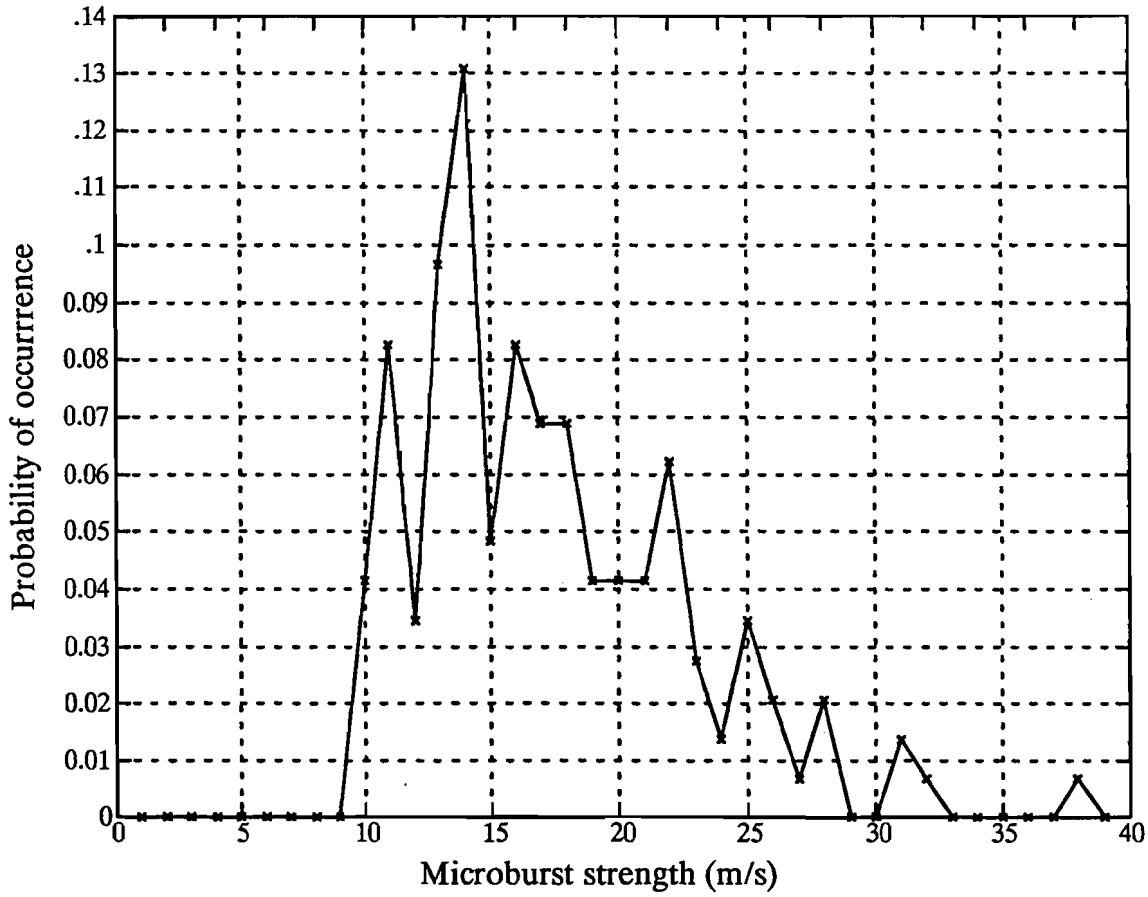

Figure 43: Observed probability density for radar-oriented microburst strengths for Denver, adapted from [DiStefano and Clark, 1990].

considered: (a) perfect detection for all events of strength 10 m/s and above and (b) perfect detection for all events of strength 7 m/s and above.

The performance impact of asymmetry is clearly more pronounced on those events with a weaker outflow strength. The detection performance statistics computed above are based on the goal of detecting all microbursts with strength above 10 m/s, as stated in the TDWR Systems Requirements Statement. The current operational goals of the TDWR are to provide informational warnings to pilots for outflows with strengths above 10 and below 15 m/s, and to begin providing hazard alerts for outflows of strengths 15 m/s and greater. If the performance impact is assessed relative to the 15 m/s alert threshold a somewhat reduced asymmetry penalty is found. The distribution of radar-oriented event strengths shown Figure 45 was obtained from the observed distribution from Denver (as in Figure 43), but renormalized to include only those events above 15 m/s strength. The aggregate POD resulting from this distribution is shown in Figure 46. As would be expected, the aggregate POD over these stronger events is considerably higher than that for all events above 10 m/s, at all radar-runway separation angles. In general the aggregate POD for events above 15 m/s

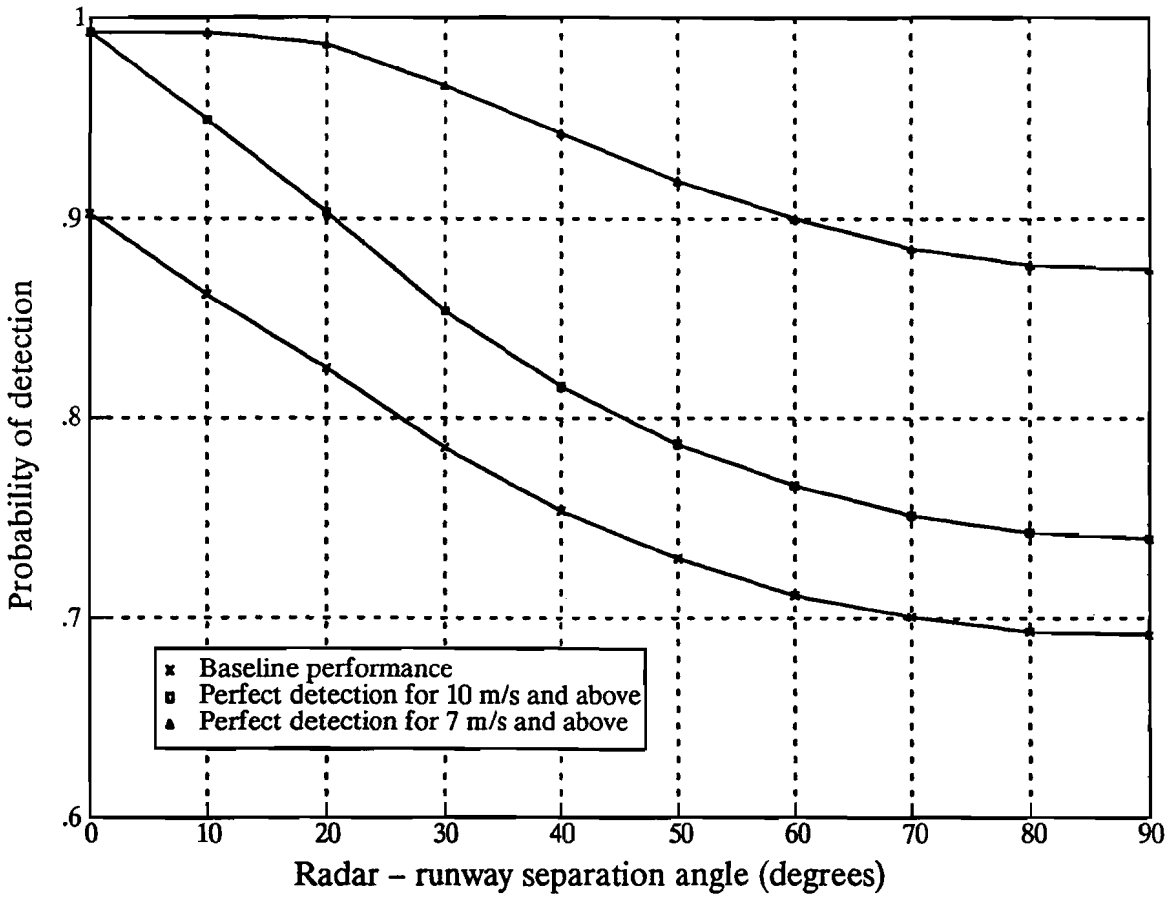

Figure 44: Aggregate probability of detection (using observed strength distribution) for all microbursts with strengths 10 m/s and above along the runway, as a function of radar–runway separation angle. The bottom curve indicates performance using baseline TDWR divergence algorithm detection performance vs radar–observed strength (from Denver, 1988 cases). The remaining two curves indicate improvements obtained by assuming divergence algorithm has perfect detection for events of 10 m/s and above and for 7 m/s and above, respectively.

is roughly 10 percent greater than the aggregate POD for events above 10 m/s, independent of the separation angle.

These total POD curves indicate the substantial effect which asymmetry can have on the overall performance of the TDWR system and the extent to which proper radar siting (relative to primary aircraft approach and departure paths) can reduce the performance loss. For radars viewing directly along flightpaths, asymmetry does not reduce the ability of the TDWR to detect windshears which present strong shears along the aircraft path. As the separation angle between the radar and the runway increases, the overall POD drops off nearly linearly until the separation angle approaches 50°. At this separation angle, the overall POD may be decreased by nearly 20 percent (assuming Denver-based divergence detection performance and probability distribution of event strengths above 10 m/s).

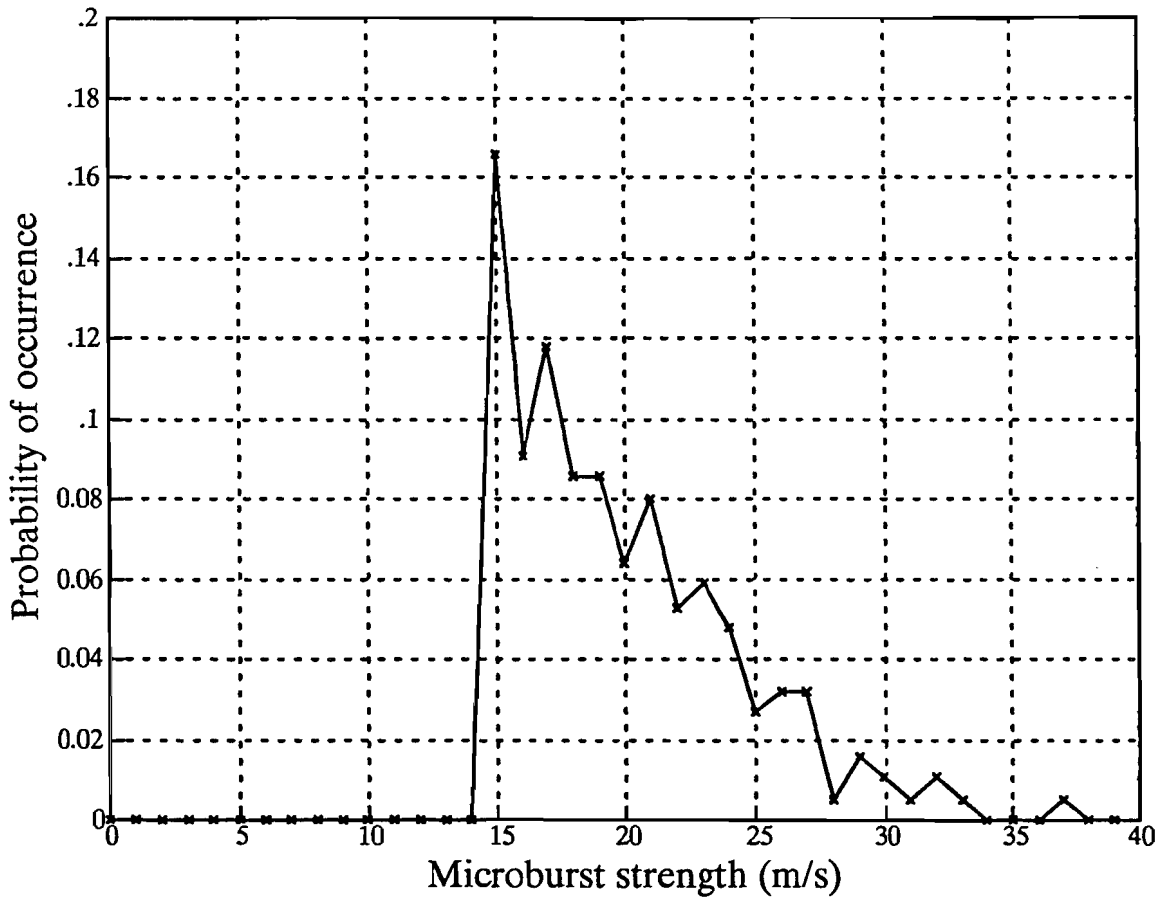

Figure 45: Observed probability density for radar-oriented microburst strengths, including only those events above 15 m/s.

Increases in the radar-runway separation angle beyond roughly 50° result in small additional losses in detection capability.

6.4. Limitations to Asymmetry Impact Analysis

The asymmetry impact analyses presented in this chapter have attempted to reflect the substantial body of information available on both the measured detection performance of the TDWR microburst algorithm (relative to radar-measured microburst strength) and the characteristics of microburst asymmetry. The results of these analyses suggest that asymmetry can seriously degrade the ability of the TDWR to detect microbursts which are strong along runways when the separation angle between the runway and the radar viewing direction exceeds 10-20°.

While this analysis appears to be qualitatively correct (i.e., that asymmetry will reduce the detection capability of the TDWR and that small radar-runway separation angles minimize the degradation), it relies heavily on several observed statistics to provide the

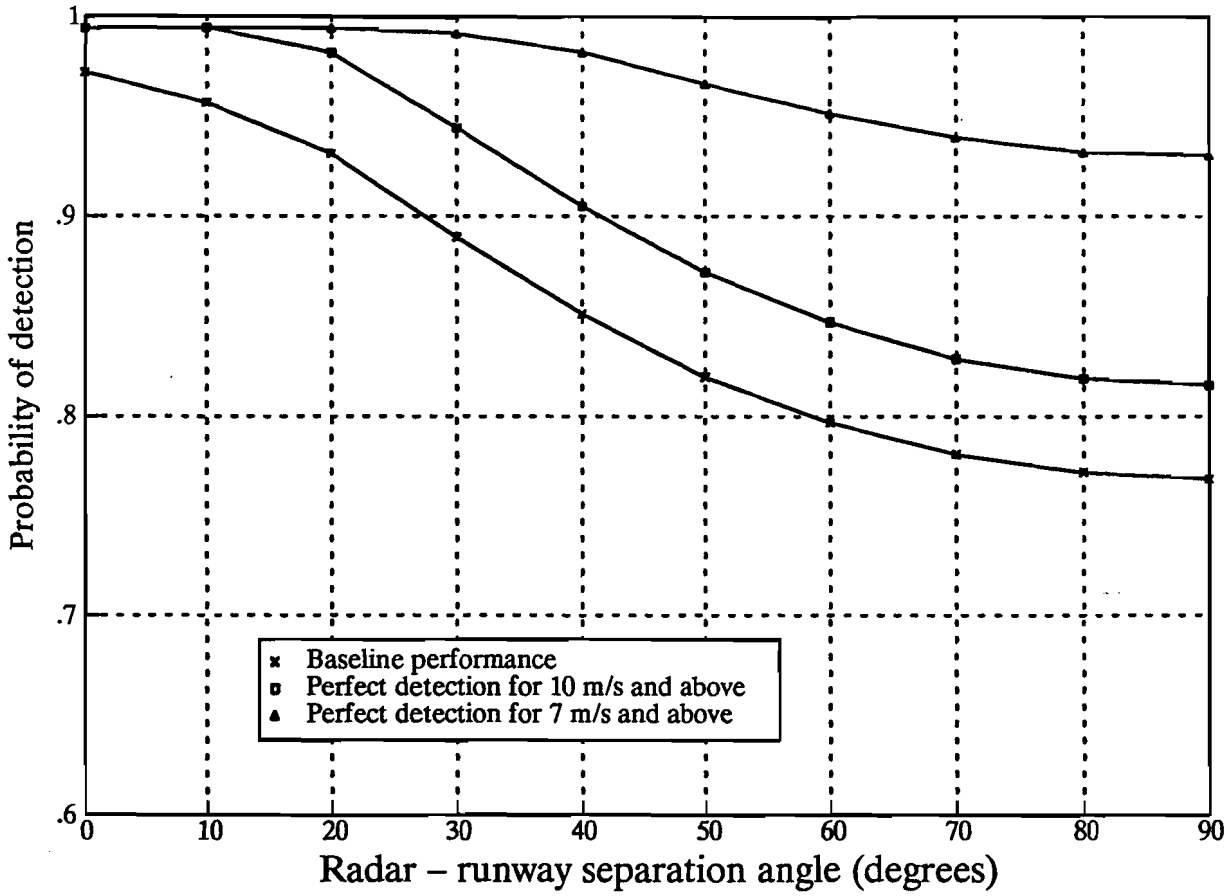

Figure 46: Aggregate probability of detection for all microbursts with strengths 15 m/s and above along the runway, as a function of radar–runway separation angle.

quantitative results obtained. The microburst asymmetry characteristics used for this analysis were obtained from an in-depth examination of dual-doppler radar observations from Denver microbursts. Similar asymmetry characteristics are not yet available from other meteorological regimes, casting doubt on the applicability of these quantitative results to the country-wide average performance. Likewise, the baseline detection performance (as a function of radar-measured strength) was based on Denver, 1988 performance results. The detection performance of the microburst algorithm was worse in Denver than in other test sites(see [Evans, 1990]), particularly in the detection of weak outflows. The net impact of asymmetry has been shown to depend strongly on the ability of the divergence algorithm to detect weak outflows, and the Denver results thus present a near worst-case assessment of asymmetry impact.

One additional caveat on the asymmetry analysis bears on the assumption that the divergence detection performance is solely a function of the observed outflow strength. It is possible that the ability of the divergence algorithm to detect an outflow depends not only on the radar-measured strength, but also on the maximum strength of the event as well. For

example, the detection rate for events of maximum strength 12 m/s, with the maximum aligned with the radar viewing direction, may be different from the detection rate for stronger events not aligned with the radar angle but which have a radar–measured strength of the same 12 m/s. Any such dependence on maximum event strength could not be identified from existing performance analysis methods and could significantly reduce the impact of asymmetry on overall performance.

Additional investigation will be required to validate and refine the analysis of asymmetry impact presented here. First, since the performance of the TDWR is known to vary with the meteorological characteristics of the environment, the asymmetry characteristics and resulting impact analysis will need to be repeated for each different environment. Clearly, those regions which experience relatively fewer asymmetric events, or fewer weak outflow events, will be impacted less by asymmetry. Second, the correlation between detection performance and maximum outflow strength (in any direction) must be investigated to determine if the projected performance degradation resulting from weaker measured outflow strengths is as severe as suggested here.

7. COMPUTATIONAL ALGORITHM ALTERNATIVES

7.1. Motivations for Improvements to the Baseline TDWR Divergence Algorithm

This chapter examines some alternative concepts for the divergence detection algorithm. These concepts are proposed to address several perceived deficiencies in the existing baseline algorithm (discussed below). The concepts presented in this chapter are oriented towards the computation of shear estimates as a basic element of the detection process as opposed to the existing pattern-based method in the baseline algorithm. The specific algorithms described in this chapter have been applied to very small samples of data for exploratory purposes but have not yet been subjected to any significant performance analysis evaluation.

The most serious deficiency of the current divergence algorithm is its poor detection performance on weak outflows. The detection rate for microburst signatures with strengths less than 15 m/s is roughly 0.75, compared to the near perfect detection rate for outflows greater than 20 m/s. Outflows of weak strength are of limited operational concern, and the reduced detection capability in this strength range has generally not been viewed as a serious threat. Unfortunately, the asymmetric nature of the microburst outflows can cause a potentially hazardous event to appear weak when viewed from a direction other than the operational flight path. The analysis in Chapter 6 indicates that a substantial number of the weak events being missed by the current microburst algorithm are probably of significant strength when penetrated from a different direction. This asymmetry effect results in the effective TDWR detection probability dropping by as much as 20 percent for microbursts with worst-case strengths above 20 m/s.

A modest improvement in the ability to detect weaker (measured) outflows would significantly reduce the likelihood that a hazardous, asymmetric microburst would go undetected by the TDWR. Such an improvement would also benefit the timeliness of the system in detecting microbursts in their earliest stages of development. For these reasons, improvements in the detection of outflows with weak signatures would appear to be warranted.

A second difficulty with the current algorithm is the selection of thresholds and parameters which optimize performance for a given environment. This tuning process is difficult to perform because of the complex interactions between the many parameters and because of the sensitive dependence of the performance on the parameter values. The behavior of the algorithm is often difficult to judge intuitively, and proper tuning thus requires considerable experimental parameter adjustments and evaluations. A more straightforward algorithm, with fewer (and less critical) parameters, would be considerably easier to adjust to the environments in meteorologically different regions of the country.

The shape estimation component of the TDWR algorithm includes complex logic aimed at smoothing the outlines of the microburst alarms without losing a significant fraction of the hazard area. This logic is needed, in part, because the existing algorithm tends to generate segments which are very irregular within the detected region. The lack of spatial

continuity in the segment detection process does not provide any smoothing influence to attempt to form a consistent boundary for the detected region; each segment is located independently, then associated together based on overlap. A detection algorithm which took more explicit advantage of the two-dimensional spatial continuity of the outflow could potentially improve both detection performance (from the added information) and provide more regular outline shapes.

The existing algorithm also fails to exploit the temporal continuity of the microburst outflow and does not use the detected region information from the previous scan to aid in the detection process on the current scan. Although the final microburst alarm generation process does use time association to reduce false alarms, this information is not fully used in the initial shear detection process. Use of previous history could allow for a more reliable detection of shear regions and could also reduce the erratic changes in detected size and shape from scan to scan.

Finally, the current algorithm does not explicitly consider the actual shear strength within the detected region. The logic tests used in the shear segment location process will identify consistently increasing velocity trends of any magnitude; only the segment endpoint trimming tests include any explicit requirement on the shear magnitude. These trimming tests ensure that segment endpoints exhibit at least a minimum shear level but do ensure that this shear level persists across the entire segment. This weak requirement on the shear level can result in segments being detected which do not actually correspond to significant shear signatures and can result in strength estimates which are inappropriate. A more direct and explicit dependence on the radial shear magnitude would provide a more sensible linkage to the actual aviation hazard factor.

The following sections of this chapter introduce the basic concept of a computational detection algorithm and present an analysis of the performance of a basic algorithm against idealized model outflow signatures. While these idealized signatures give some intuitive understanding of the tradeoffs involved in the computational approach, the remaining sections describe the additional complexities involved when applying these techniques to actual radar measurements and outline more sophisticated algorithm components for dealing with these real-world problems.

7.2. Computational Approaches to Shear Detection

Computational approaches to the detection of divergence regions may provide simpler algorithms with performance equal to or better than that obtained with the current TDWR algorithm. The term "computational algorithm" here refers to an algorithm based primarily on the numerical estimation of the radial shear as the basis for the detection of a shear region. The baseline TDWR algorithm does not explicitly calculate a shear, but rather uses logical tests within the search window to ensure that velocity measurements are generally increasing with range. The logic-based algorithm used in the TDWR incorporates a number of tests designed to reduce noise effects and handle various observed characteristics of the velocity measurements. These factors are handled in a fairly implicit manner, making the algorithm difficult to describe, difficult to understand, and difficult to adjust to new environments. The hope for a more computational-based algorithm is to provide a

technique which is more easily described in common mathematical language and which has more easily understood behavior.

One computational approach to microburst detection was investigated by [Noyes, 1990] for use in the ASR-9 windshear detection context where it was determined to perform as well (or better) than the baseline TDWR algorithm. The divergence detection problem in the ASR-9 context is very similar to that in TDWR, with two primary differences. First, the ASR-9 velocity measurements are typically much noisier than those obtained by the TDWR. TDWR velocity measurements generally have a standard deviation no greater than 1 m/s, whereas ASR-9 velocity estimates may have noise levels up to 3 m/s [Weber, 1989]. Second, the ASR-9 makes a new velocity measurement scan at the surface every 4.8 seconds compared to the once-per-minute surface updates obtained with TDWR. The algorithm described by [Noyes, 1990] takes advantage of this high update rate of the ASR-9 to reduce the number of noise-induced false alarms. This form of temporal filtering would probably not be fruitful in the TDWR because of the 60-second update rate. The discussion below extends the analysis presented by Noyes, with an emphasis on the TDWR application and presenting a new method (appropriate to the TDWR update rate) for exploiting the temporal continuity of microbursts.

The basic structure of the computational detection algorithm is shown in Figure 47. In this approach, the raw velocity measurements are first processed by a spatial filter. This filtering step is designed to reduce the noise level present in the signal to avoid false alarms from the shear estimation stage. For this analysis, three filter types have been considered: (a) no filtering, (b) mean filtering and (c) median filtering. The filtered velocity measurements are then used to estimate the shear of the radial velocity. Two common techniques for estimating this derivative are considered here: the finite difference approach and the least-squares fit method. For both the shear estimation and filtering stages (when filtering is employed) a spatial window size must be selected. The selection of this window size is a key factor to the algorithm design. A simple analysis of the noise sensitivity of the two shear estimators, as a function of window size, is presented in Appendix 2. The simulation tests described below examine this sensitivity using more realistic signal models.

Once the shear estimates have been computed, they must be thresholded against a selected shear level to produce "raw" point shear detections. A map of these point shear

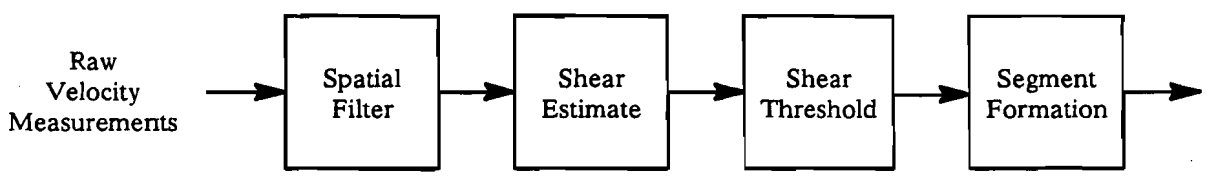

Figure 47: Basic structure of a simple computational detection algorithm. Velocity measurements are used to estimate radial shear, which is then thresholded against a fixed shear level. Contiguous segments above the threshold are identified as shear segments, which are validated using an integrated loss threshold.

detections could conceivably be used as a final display product, but further filtering is generally desired to reduce false alarms. In addition, these point detections should be clustered into two-dimensional shear regions for compatibility with the use of features aloft in the TDWR microburst algorithm. This clustering process is also necessary for the generation of a sensible estimate of the total strength of the microburst shear region.

The clustering method employed in the algorithm by [Noyes, 1990] used a two-dimensional region growing approach, common to many image processing applications. For compatibility with the existing structure of the TDWR divergence algorithm, the approach used here first locates linear segments of shear (by simply identifying contiguous gates above the shear threshold) along each radial. These shear segments are then clustered in azimuth in the same manner as in the baseline TDWR algorithm: those segments which overlap in range and are close in azimuth are joined into the same region.

The strength for each shear segment is computed as the sum of the point shear estimates along the segment, multiplied by the gate spacing, to produce a velocity value. This strength value is slightly different from that used in previous algorithms; both the TDWR and ASR-9 algorithms define the strength for a segment as the difference between the velocity measurements at the ending and starting points of the segment. The definition chosen here is designed to (1) be more consistent with the shear-based orientation of the computational algorithm, (2) provide a useful metric to indicate the inherent ability of the algorithm to detect the full shear signature, and (3) provide a more realistic estimate of the actual hazard level present.

7.3. Performance of Basic Computational Algorithm on Idealized Signatures

This simple computational algorithm performs quite well when applied to idealized microburst signatures, even in the presence of substantial levels of white noise. The three model microburst profiles shown in Figure 48 were processed by the above algorithm to determine the probability of detection and false alarm. The first case represent a strong, clear microburst signature (modelled by a sinusoid) with a 30 m/s velocity change over 3 km. Reliable detection of this signature should be possible, even in the presence of strong interference. The second case represents a small, yet strong, microburst to test the ability to detect severe microbursts in their earliest stages of development. This event has a 14 m/s velocity differential over 0.72 km, using the r^{-2} model described in Chapter 2. The third profile is a sinusoid with a 12 m/s differential over 4 km, representing a marginally weak microburst.

Each of these profiles was contaminated with a white noise signal having a standard deviation of 1 m/s. The computational algorithm was applied, and the statistics for detection and false alarm were tallied over several hundred noise realizations. The probability of detection was computed as the average (over the noise realizations) fraction of the gates in the true shear region which were detected as being in a segment. The true shear region was that region of the model profile which had a shear level above 2.5 m/s per km. The probability of false alarm was computed as the average fraction of the remaining gates which were flagged by the algorithm as being in a shear segment. An additional statistic was

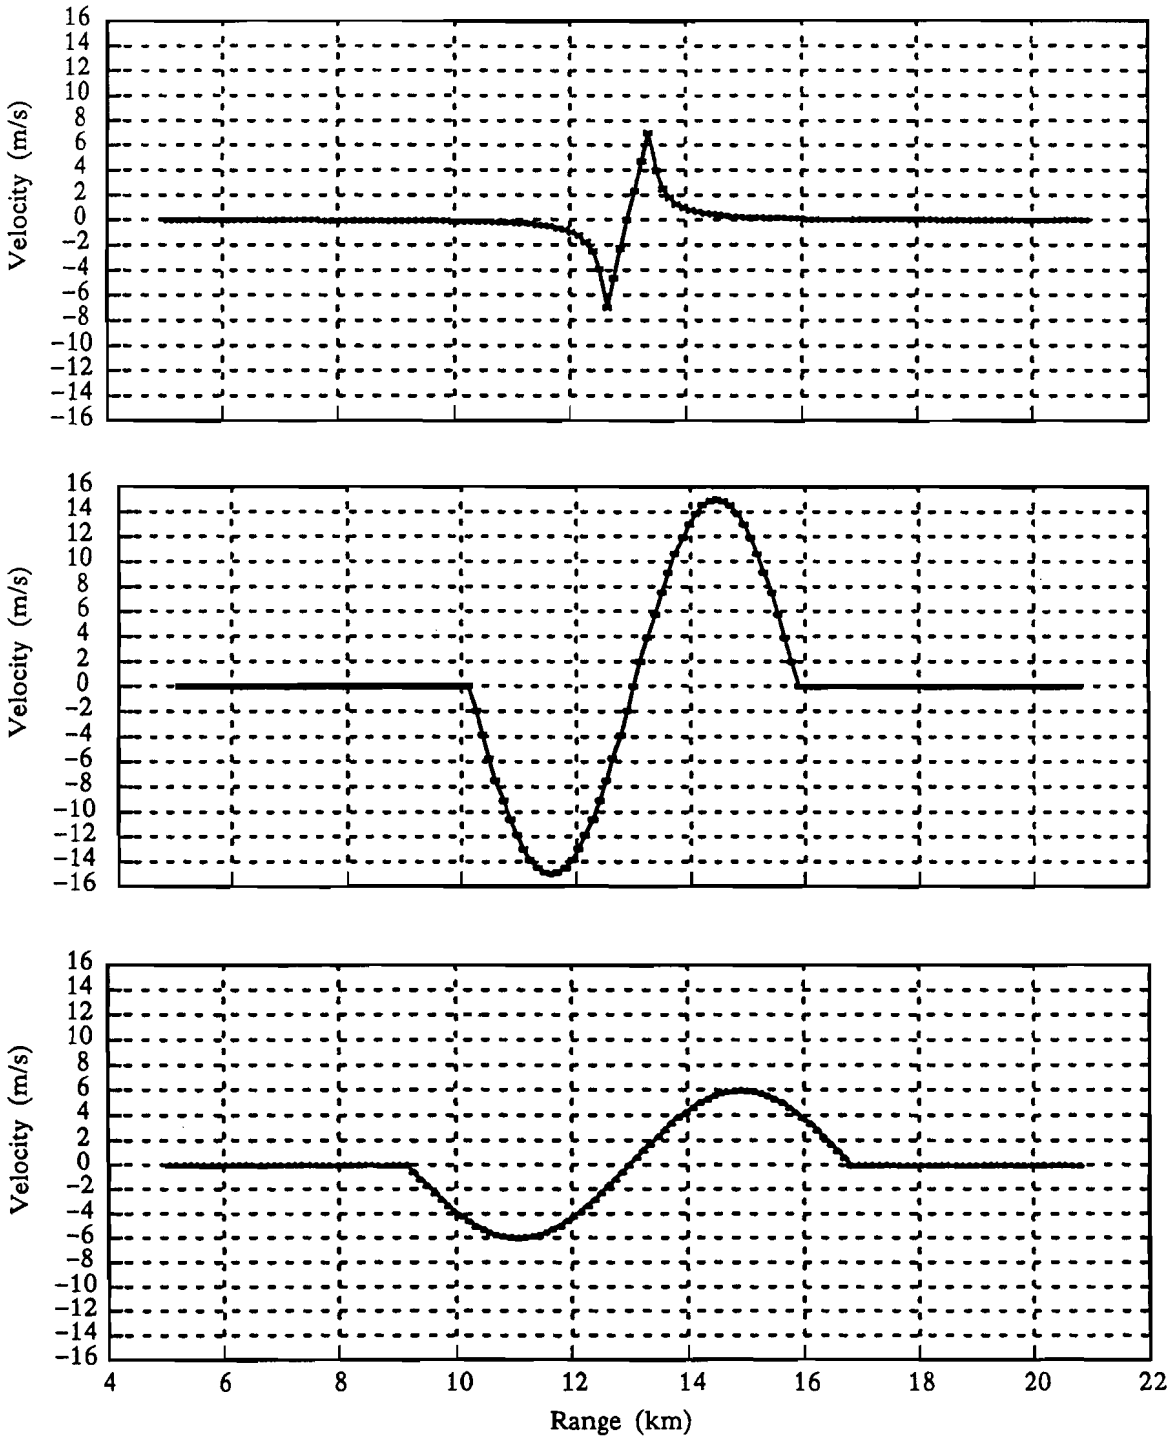

Figure 48: Three microburst velocity profiles used to evaluate filtering and shear estimation alternatives.

computed for each trial, that being the sum of the strengths for all segments which overlapped some part of the true shear region. This statistic is significant in that it quantifies the degree to which the actual shear signature was attenuated by the smoothing and shear estimation process.

These average statistics were computed for each combination of smoothing filter and shear estimator, at each of several spatial window sizes as listed in Table 8. The results of these simulations are summarized in Figure 49, Figure 50 and Figure 51 for the strong, weak and small shear models, respectively (a complete tabulation of the performance statistics is presented in Appendix 3). Each of these figures plots the probability of detection (gate by gate) versus the average detected strength of the shear region for each of 110 variations in filter algorithm, filter width, shear algorithm and shear algorithm width. False alarm rates in all cases were below two percent. This low false-alarm rate (compared to that reported by [Noyes, 1990]) is the result of the segment formation process which requires several consecutive gates to be above threshold and to form an integrated loss of at least 5 m/s. The false-alarm rates reported in [Noyes, 1990] were point shear threshold crossings prior to the segment formation (or region growing) stage of processing.

The strong shear profile is detected at least 80 percent of the time by all the algorithm variations, with most producing strength values between 27 and 29 m/s. The weak profile, however, shows a wide range of detection performance depending on the combination of spatial window sizes used. The strength detected is directly proportional to the percent of the shear gates detected, with most of the better combinations yielding an 80 - 90 percent probability of detection and a strength of 9 - 10 m/s. The behavior on the small shear signature is more complex, showing two branches in the scatter of points. About half of the parameter combinations result in poor detection rates, with strengths between 5 and 6 m/s. The strength detected in these cases appears to be independent of the detection rate. The

Table 8.
Filter algorithm, filter width, shear algorithm and shear width combinations used for performance simulations

Filter algorithm	Widths (gates)
Null	N/A
Mean	3,5,7,9,11
Median	3,5,7,9,11

Shear algorithm	Widths (gates)
Finite Difference	3,5,7,9,11
Least Squares	3,5,7,9,11

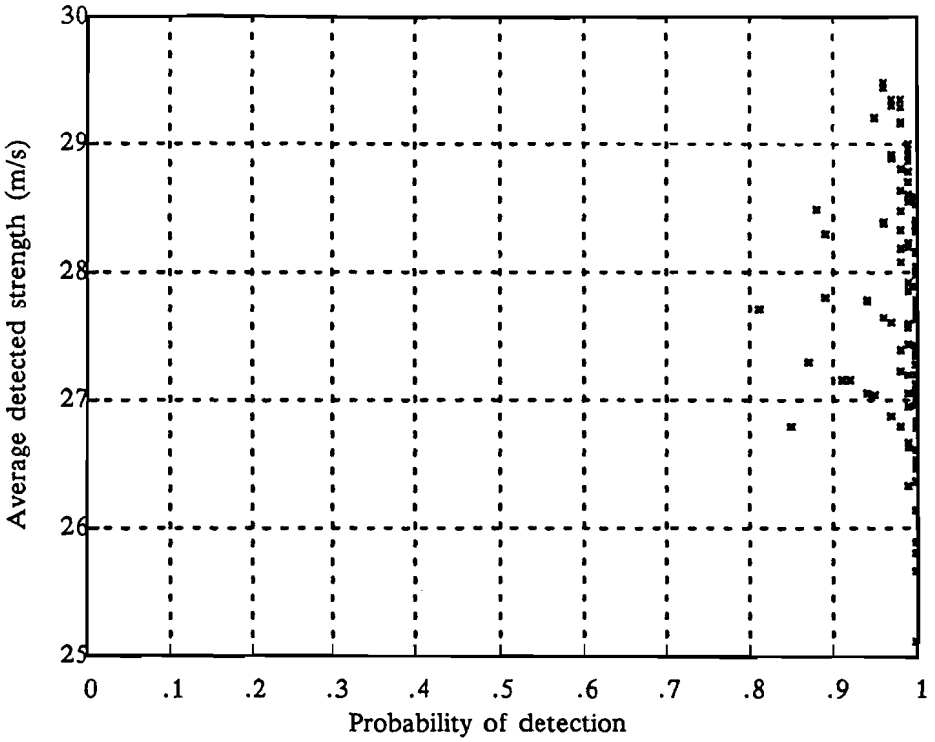

Figure 49: Detection performance scattergram for "strong" shear profile.

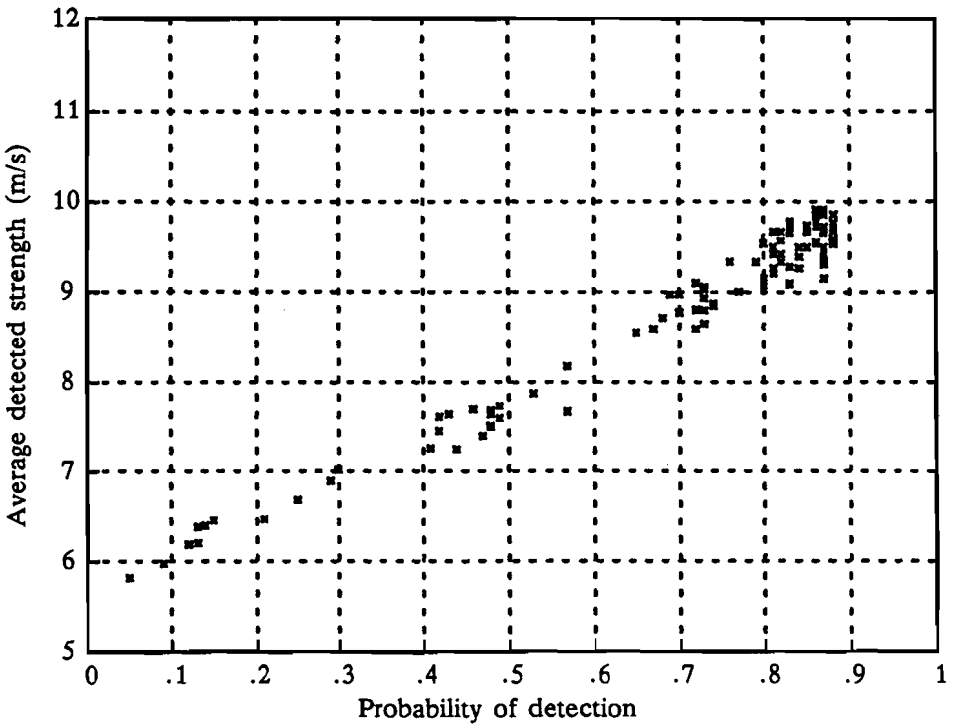

Figure 50: Detection performance scattergram for "weak" shear profile.

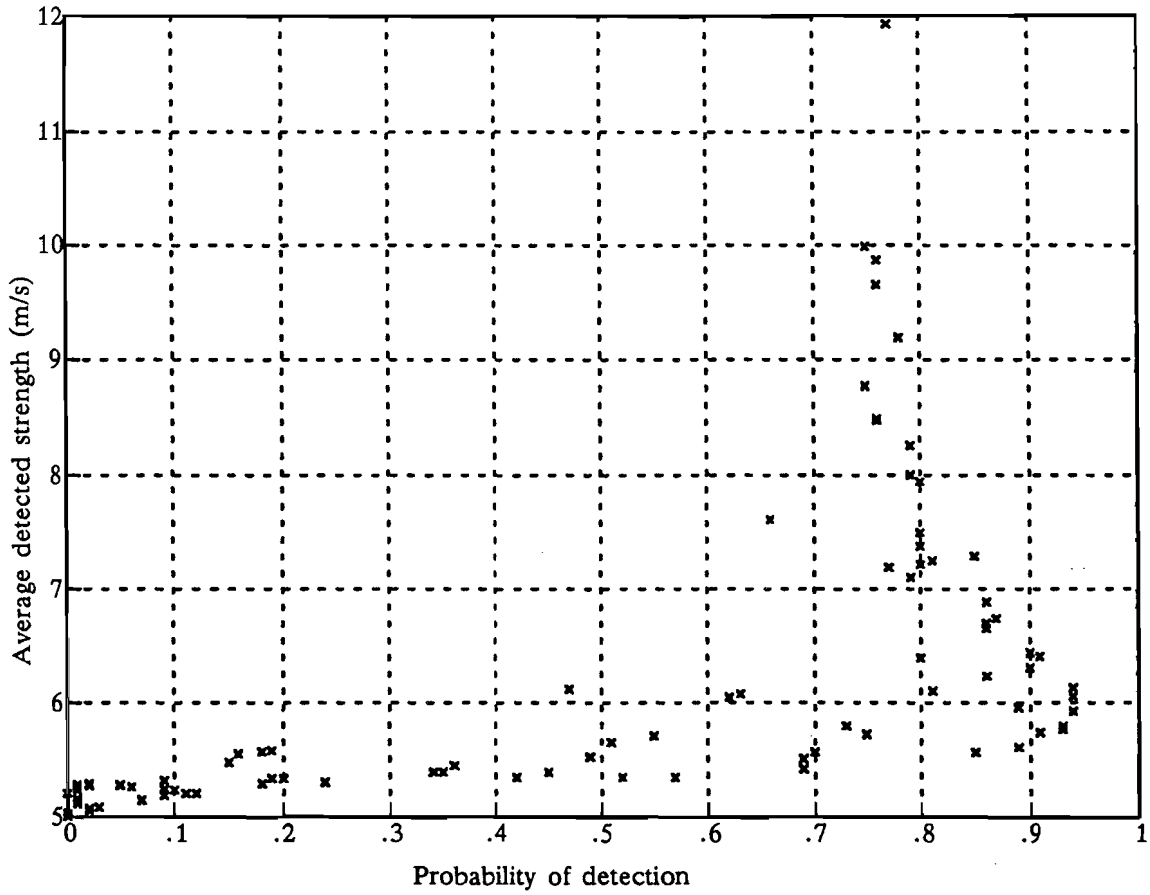

Figure 51: Detection performance scattergram for "small" shear profile.

other half of the combinations provide a wide range of detected strengths (from 5.5 to 12 m/s), with detection rates between .75 and .95. For these higher performance combinations, a tradeoff is required between average detected strength and probability of detection; the higher the detection rate, the lower the total detected strength.

While these results indicate that fairly good performance can be obtained by choosing the proper parameter combination for each of the three shear profiles, no single combination of window sizes and filter/shear algorithms provides optimal performance across all three profiles. Since the strong shear case is detected well by most combinations, the tradeoff between performance on the small profile and that on the weak profile is most pronounced. The scattergram in Figure 52 compares the average detected strengths for these two profiles for each algorithm variation. This comparison clearly indicates that detecting a substantial portion of the total strength for either signature profile requires that the detection of the other profile suffer a substantial loss in detected strength.

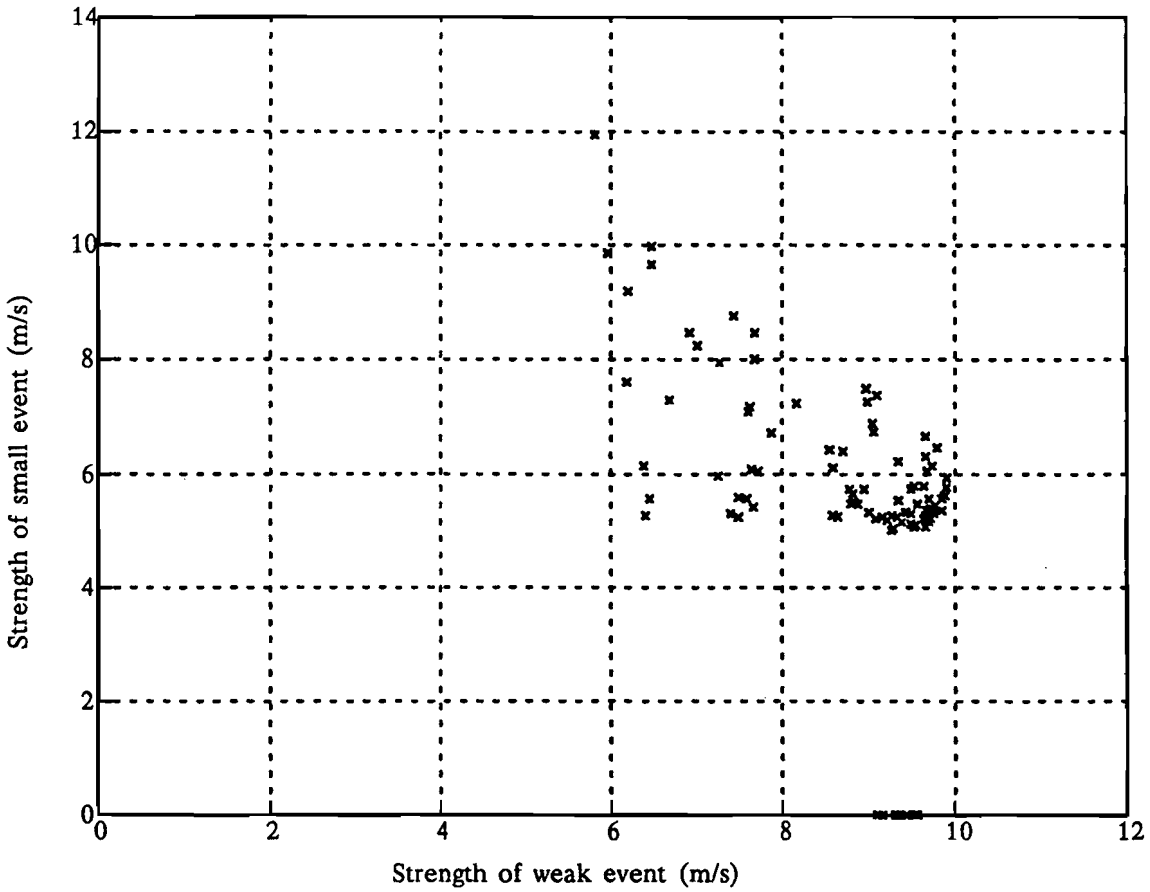

Figure 52: Comparison of average detected strength for the weak and the small shear profiles, for each of the filter/shear algorithm and window size variations.

7.4. Basic Computational Algorithm Applied to Measured Data

While the use of idealized microburst models can provide some insight into the performance of the algorithm, it is equally important to examine its operation on actual weather measurements. The basic computational algorithm described above was applied to a radar scan from the July 11, 1988 case described in Chapter 4. This version of the algorithm employed no smoothing filter and used a finite difference shear estimator over a five gate window. Segments were clustered using a 0.5 km range overlap requirement and clusters were required to have at least three segments and a total area of at least 1.0 square km. The shear segments and clusters obtained on the radar scan from time 22:09:00 are shown, along with the ground truth outline for this scan, in Figure 53. The shear segments shown in this figure cover the true shear region reasonably well but also display a number of false segments. Based on the model results from the previous section, only one to two percent of the gates outside of the shear region should be falsely alarmed, which is roughly equivalent to five to six range gates of false alarm per radial of measurements. The

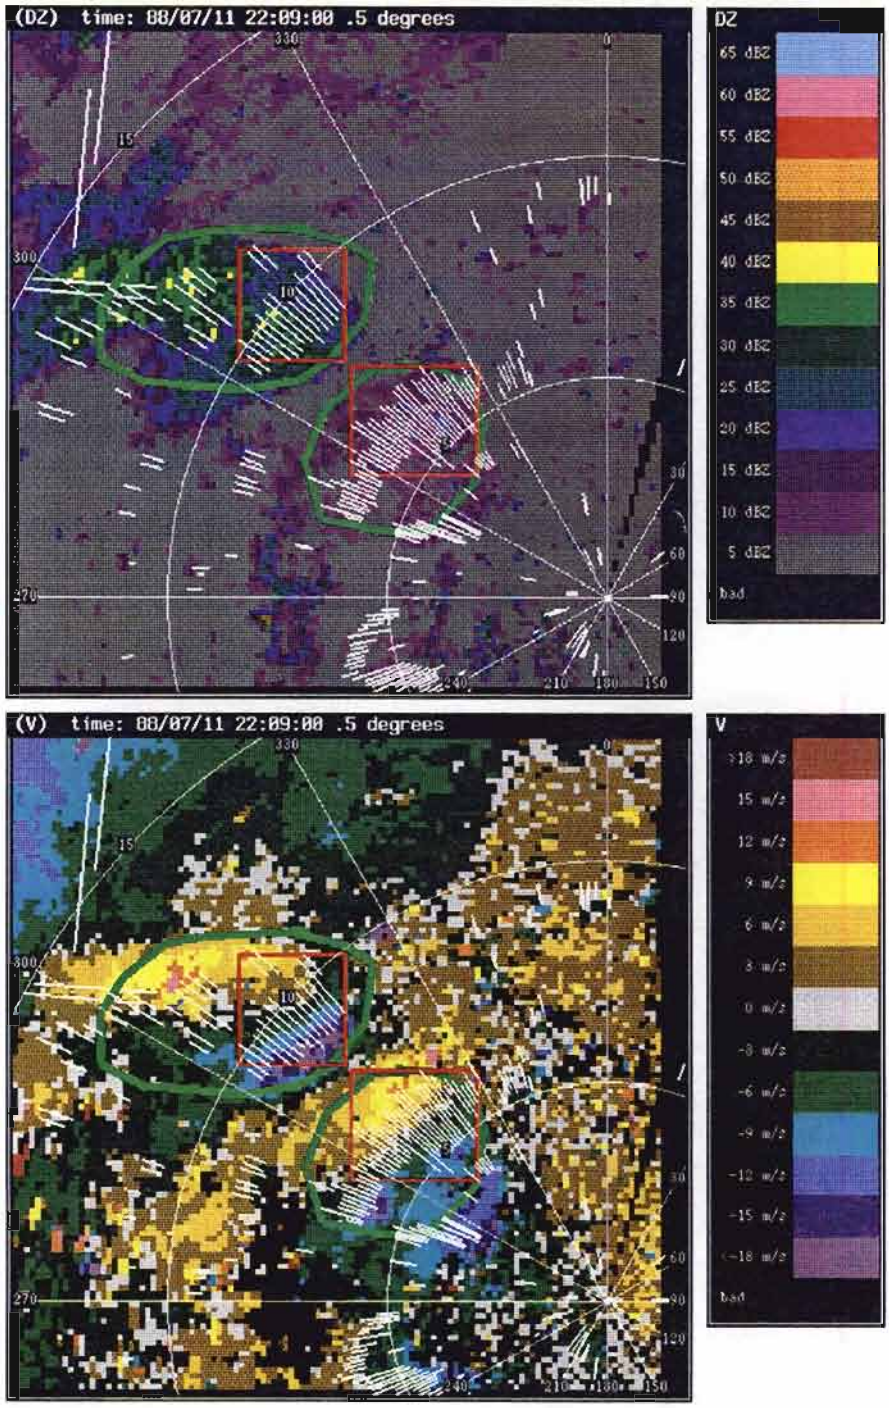

Figure 53: Shear segments and clusters from basic computational algorithm overlaid on base radar measurements from 22:09:00 on July 11, 1988. Ground truth outline is shown in green.

segments shown in Figure 53 appear to have a slightly higher false-alarm rate than the model results would predict. This observation suggests that a substantial fraction of the false alarms which occur with measured data fields are caused by natural variations in the true windfield and are not solely the result of spatially uncorrelated noise added during the sensing process.

When compared to the performance of the baseline TDWR algorithm on this same case (see Figure 23), it is evident that the basic computational algorithm has failed to detect a substantial portion of the outflow event near the runway (at the area centered at 11 km range and 306° azimuth). The shear signature in this region is rather noisy, as illustrated in the velocity profile of Figure 54. Between 10 and 12 km range, the velocity measurements in

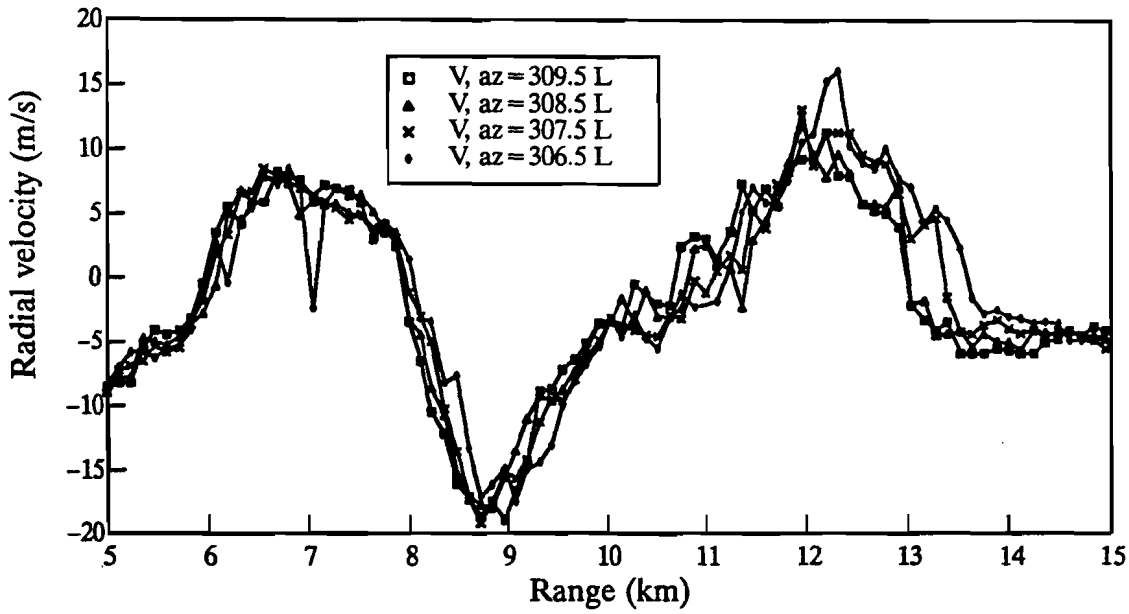

Figure 54: Velocity profiles for microburst region near runways on July 11, 1988 at 22:09:00. Variability in shear signature between 10 - 12 km range results in basic shear detection algorithm failure.

this area exhibit both positive and negative variations in range. While the overall trend is an increase in range (i.e., a divergence), the trend is not sufficiently consistent to allow the five-gate shear estimator to reliably detect shear above the threshold level. Since the basic algorithm requires consecutive range gates to have shear estimates above threshold (i.e., a segment is terminated at the first non-detected gate), these fluctuations result in gaps in the detected shear region.

The variability of this shear signature from range gate to range gate is not uncommon and is indicative of the complexity of realistic shear signatures. The baseline TDWR algorithm was able to overcome this variability as a result of the decision criteria used to identify the end of a segment. The failure of the basic computational algorithm in this case suggests the need for a more sophisticated shear estimation and/or thresholding algorithm.

7.5. Use of Temporal Feedback to Improve Detection

The basic computational algorithm thresholds each point shear estimate against a fixed threshold level to make the initial determination whether or not that point is in a shear region or a non-shear region. This simple thresholding approach is very sensitive to the noise content of the shear estimate and is prone to both false alarms and missed shear detections (as illustrated in the example above). Simple thresholding makes no use of the substantial spatial and temporal continuity present in the signature of an actual windshear event. To take better advantage of this continuity, a new technique has been examined for classifying sample points into shear or non-shear regions. This approach relies heavily on the temporal history of the local shear estimates to adapt the shear threshold level for the current range sample, introducing a hysteresis effect.

In the temporal feedback technique, a classification map is maintained for each surface radar scan, and the classification thresholds used on the current tilt are adaptively chosen based on the local classification results from the previous tilt. This approach requires the identification of a compact region of strong shear to initiate the detection of a region but then reduces the shear threshold and expands the area to ensure detection on subsequent scans.

7.5.1. Data Editing and Shear Estimation

A data flow diagram for the processing used by this technique is shown in Figure 55. To reduce the likelihood of missing small shear regions, no initial smoothing filter is used. A simple filter is employed, however, to remove isolated patches of valid velocity measurements in regions generally flagged as invalid. Such isolated patches may be caused by ground clutter breakthrough or point target interference and have been observed to cause false shear regions to be detected. The "isolated patches" filter simply examines the data values in a rectangular range-azimuth window about the current range gate and counts the number of valid velocity measurements in the window. If too few data points in the window are marked as valid (by the base data quality algorithms), then the current point is made invalid. This filter does not alter valid data values but simply deletes small isolated regions of valid measurements. The same filtering process is applied to the shear estimates immediately after they are computed from the edited velocity field. In both passes of the filter, a window of three radials by five gates is examined, and at least 10 of the 15 data points in the window must be valid to prevent the center point from being marked invalid. The shear estimation is performed using a finite-difference estimator over a five-gate window in range.

7.5.2. Shear Classification

The classification process takes as input two fields: the estimated shear values surrounding the current range sample and the the classification map from the previous surface radar scan. The output from the process is a new classification map, where each range gate has been classified into one of the four categories: "divergence," "convergence," "stable" or "unknown."

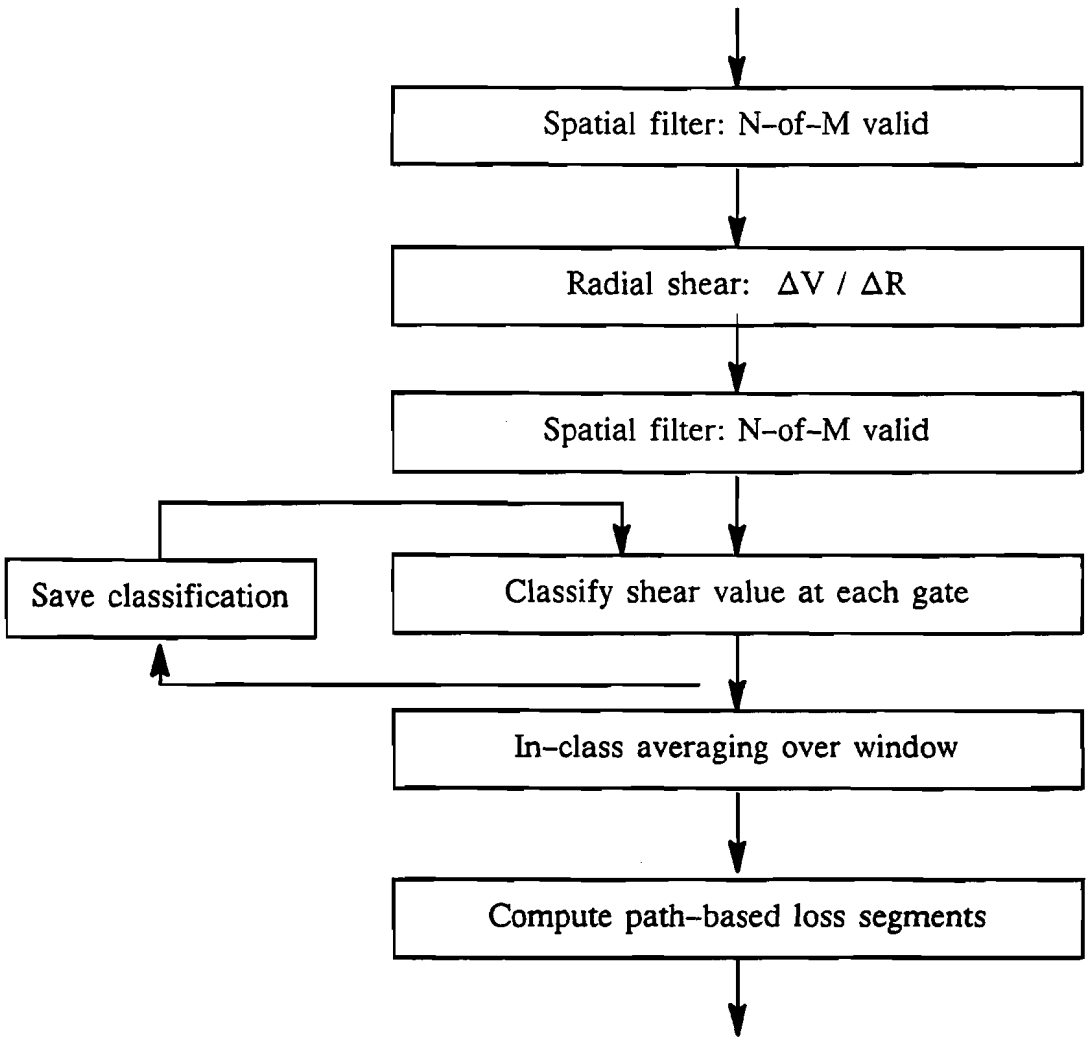

Figure 55: Overview of the stages in the computational divergence algorithm.

The decision procedure used to perform the classification at each point begins by examining the classification map from the previous scan. A window (of three radials by five gates) about the current point is scanned, and the number of samples in each of the four categories in computed. The category with the most number of samples is termed the "majority vote" for the previous scan and is used as an initial estimate of the classification at this point for the current scan. This initial estimate may be modified by the shear estimates computed for the current scan if there is sufficient evidence (using the rules to be described next) that the local "state" has changed.

The second step in the classification process is to examine the shear values in this same three radial by five gate window about the current point and to threshold the shear estimates into seven discrete levels as illustrated in Figure 56. The total number of sample points in the window whose shear values fall into each of these seven levels is then counted and used to decide if the initial classification estimate is to be modified. The rules which determine

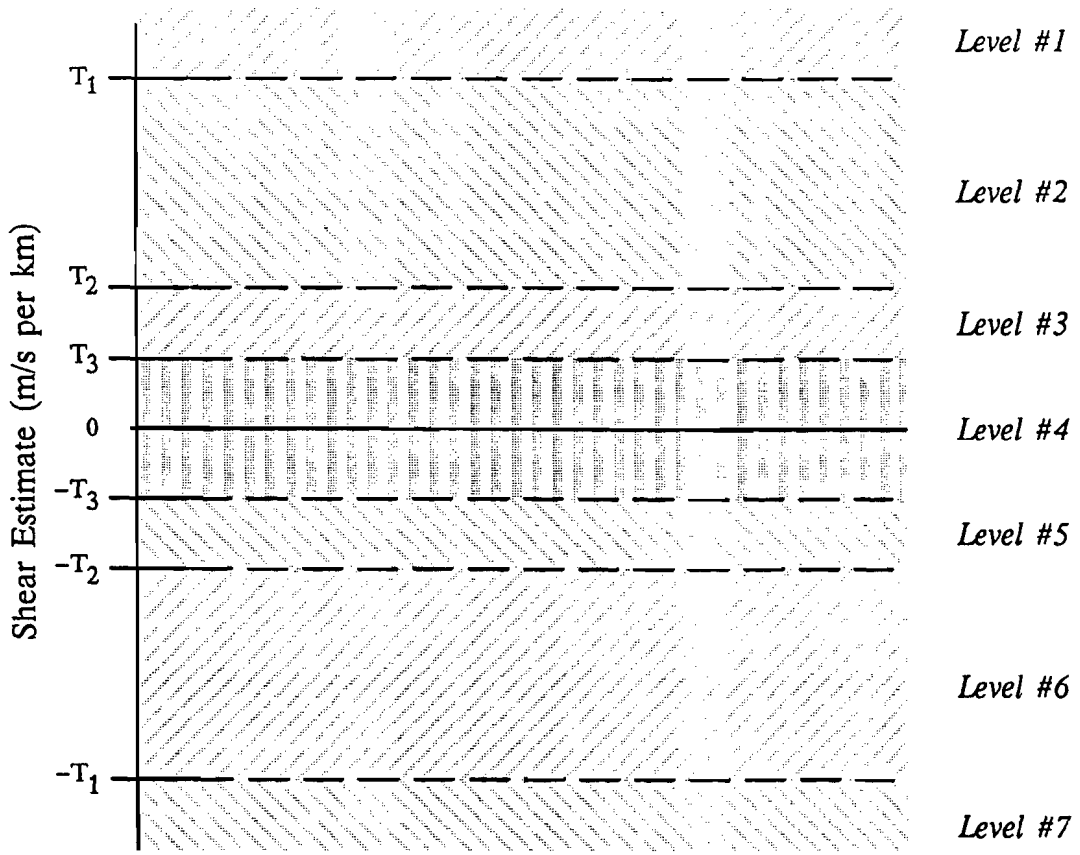

Figure 56: Quantization of shear values into seven levels. Three basic threshold parameters are applied to both positive and negative shear amplitudes, and the number of values for each region are tallied. The nominal values used for the three thresholds are 5, 3 and 2 m/s per km, respectively.

the state transitions are detailed in Table 9. This table indicates the sequence of tests applied, based on the initial classification estimate, to determine if a new classification should be assigned.

The initial state for each cell in the classification map (set at system initialization time) is UNKNOWN. From Table 9 it can be seen that the state will remain UNKNOWN until such time as at least 8 of the 15 values in the window become "valid," at which point the cell will be classified as STABLE. Transitions from STABLE to either DIVERGENCE or CONVERGENCE may then take place if a substantial number (at least seven) of strong shear values of the appropriate sign are found, with no more than two strong values of the opposite sign in the window. Using the decision strategy indicated in the table, a transition directly from UNKNOWN to DIVERGENCE or CONVERGENCE is also possible under these same conditions. These rules require substantial evidence of a shear region for a cell ever to be classified as DIVERGENCE (or CONVERGENCE), starting from the initial UNKNOWN state. Once in the DIVERGENCE state, however, the conditions for remaining in that state are much more liberal. As long as at least four of the 15 points in the window are

Table 9.
Rules used to update initial classification estimates based on current shear estimate statistics

Initial Estimate	Is changed to	When these conditions are met*
UNKNOWN	DIVERGENCE	Level #1 count $\geq$ 7; Level #7 count $\leq$ 2
	CONVERGENCE	Level #7 count $\geq$ 7; Level #1 count $\leq$ 2
	STABLE	Number valid $\geq$ 8
	UNKNOWN	otherwise
STABLE	DIVERGENCE	Level #1 count $\geq$ 5; Level #7 count $\leq$ 2
	CONVERGENCE	Level #7 count $\geq$ 5; Level #1 count $\leq$ 2
	STABLE	Number valid $\geq$ 8
	UNKNOWN	otherwise
DIVERGENCE	DIVERGENCE	Total level #1 + #2 + #3 counts $\geq$ 4
	CONVERGENCE	Total level #6 + #7 counts $\geq$ 5
	STABLE	Number valid $\geq$ 8
	UNKNOWN	otherwise
CONVERGENCE	CONVERGENCE	Total level #5 + #6 + #7 counts $\geq$ 4
	DIVERGENCE	Total level #1 + #2 counts $\geq$ 5
	STABLE	Number valid $\geq$ 8
	UNKNOWN	otherwise

*Rules are applied in the order listed for each initial estimate category, and the first transition with a satisfied condition is applied.

above the lowest magnitude threshold, then a cell previously classified as DIVERGENCE will retain that classification. This hysteresis effect attempts to prevent a region from being missed, once it has been detected originally, without suffering the false-alarm penalty associated with low initial detection thresholds.

7.5.3. In-Class Averaging

Once each range sample has been classified as UNKNOWN, STABLE, DIVERGENCE or CONVERGENCE, the shear field is spatially averaged. This averaging is intended to smooth the shear field so that the next stage, computing loss estimates and shear segment boundaries, will provide spatially consistent results. The averaging is also intended to remove some of the noise content of the shear estimates. To prevent attenuation of the shear magnitudes, however, this spatial averaging is done only among data values in the same classification category. At each range gate, all of the surrounding shear values (in the three radial by five gate window) which have the same classification as the center point are averaged. This "in-class" averaging technique prevents the shear values inside the shear

region from being biased by values outside the shear region, while still reducing the variance of the shear estimates and imparting spatial continuity to the shear field.

7.5.4. Shear Integration and Segment Formation

Once each point has been classified and the shear field has been spatially smoothed, the loss segment module is used to locate radial segments of loss. The basic loss segment process simply walks out along each radial in range, integrating each consecutive run of significant shear values. The shear values are multiplied by the range gate spacing at each point to convert to equivalent loss. Each run of consecutive points classified as significant shear result in a loss segment, subject to length and strength thresholding. Those loss segments which pass the nominal thresholds are then used in the same manner as in the current TDWR algorithm; they are azimuthally associated to form regions and passed on to the remainder of the microburst detection algorithm.

7.6. Temporal Classification Performance on July 11, 1988 Case

The computational algorithm, as described above, was applied to the July 11, 1988 microburst case discussed in Chapter 4. The plot shown in Figure 57 illustrates the divergence segments and resulting regions formed by the algorithm, overlaid with the microburst truth region (stippled area) for the first scan of the example, at 22:04:01. The shear classification map is initialized to UNKNOWN at all cells prior to this scan, so the algorithm will have reduced sensitivity for the first few scans. These results should be compared with the baseline algorithm outputs shown in Figure 18. The computational algorithm clearly generates far fewer segments on this scan, partly because of the greater smoothing performed by the computational algorithm and partly because the classification map initialization.

The computational algorithm results for the next scan (at 22:05:04) is shown in Figure 58, with two regions of shear detected near the airport. Compared to the detections from the baseline algorithm (Figure 19) we see the computational algorithm not only produces far fewer extraneous shear segments but also produces much more compact and consistent shear segments in these two regions. Figure 59 shows the results for the scan at 22:06:01, where the airport microburst is well detected by the computational algorithm. The baseline algorithm also detected this event but with fewer, more scattered shear segments (Figure 20). The next four scans, shown in Figure 60, Figure 61, Figure 62 and Figure 63, respectively, illustrate the ability of the computational algorithm to detect the shear region effectively. The shear segments produced by the computational algorithm cover the true shear region quite well, include very few missed segments in the interior of the region, and display considerable spatial continuity in the segment endpoint locations. In the final scan of the example (Figure 64) the computational algorithm splits the microburst into three detected regions based on the lack of overlap between the detected segments.

This example illustrates the ability of the computational algorithm to detect a particularly important microburst. The example also shows that the temporal filtering provided by the classification map may help smooth the shear detections generated by the algorithm and allow the shear region to be detected with a more complete set of segments than could be obtained from the baseline algorithm.

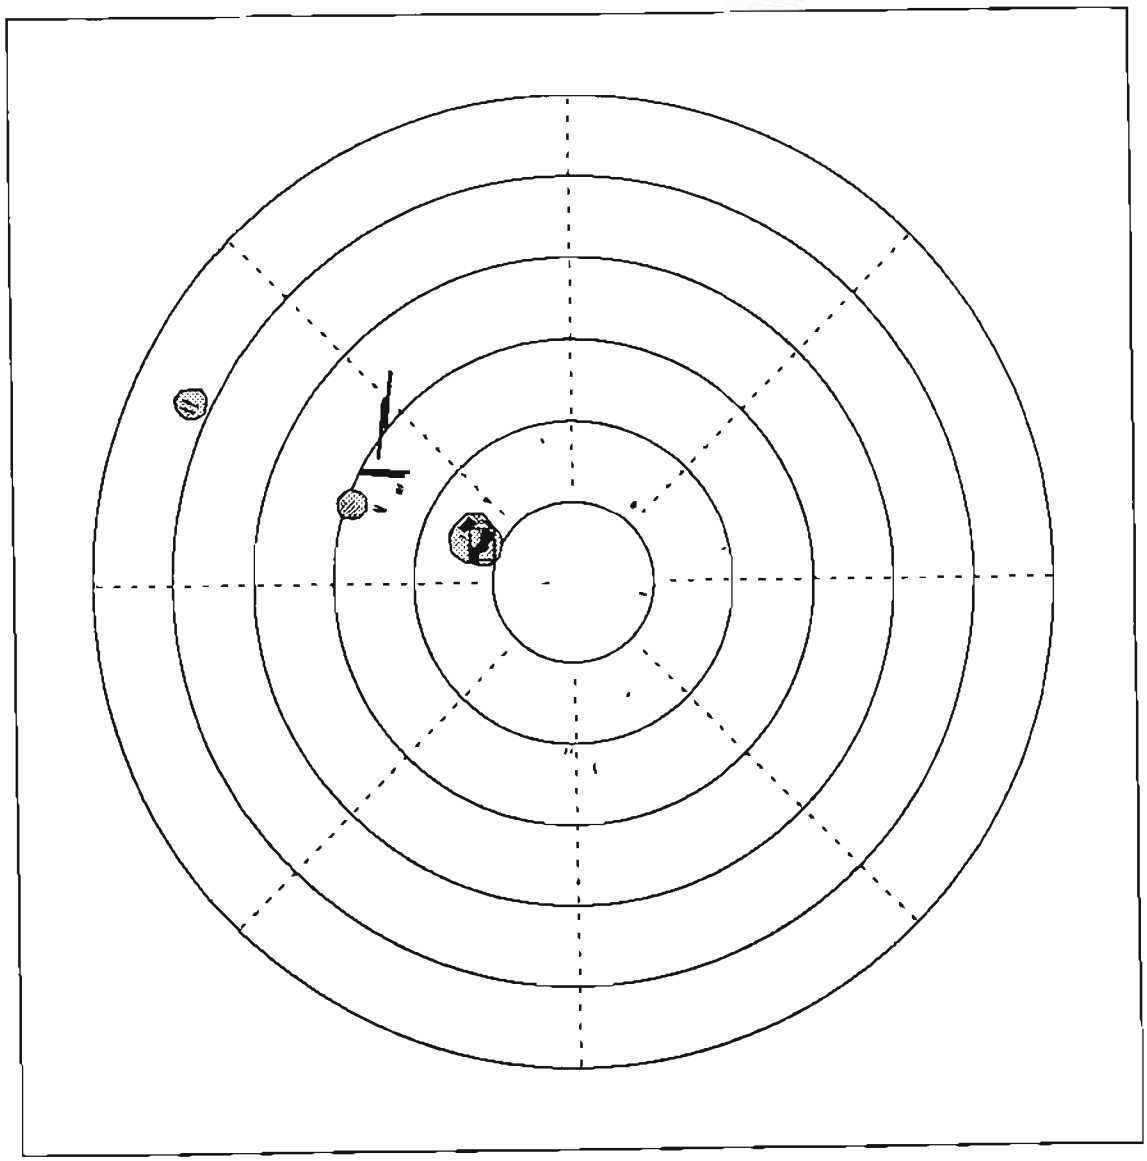

Figure 57: Computational algorithm results for 11 July 1988 case at radar scan time 22:04:01.

This one example does not provide much insight into the statistical POD and PFA performance of the algorithm; these performance metrics would require a much more comprehensive evaluation against a rather large set of cases. Using the simple hit/miss scoring criteria which has typically been applied to the microburst detection algorithm, it may be difficult to measure substantial differences between the baseline and computational algorithms. More detailed scoring methods, such as the path-based technique from

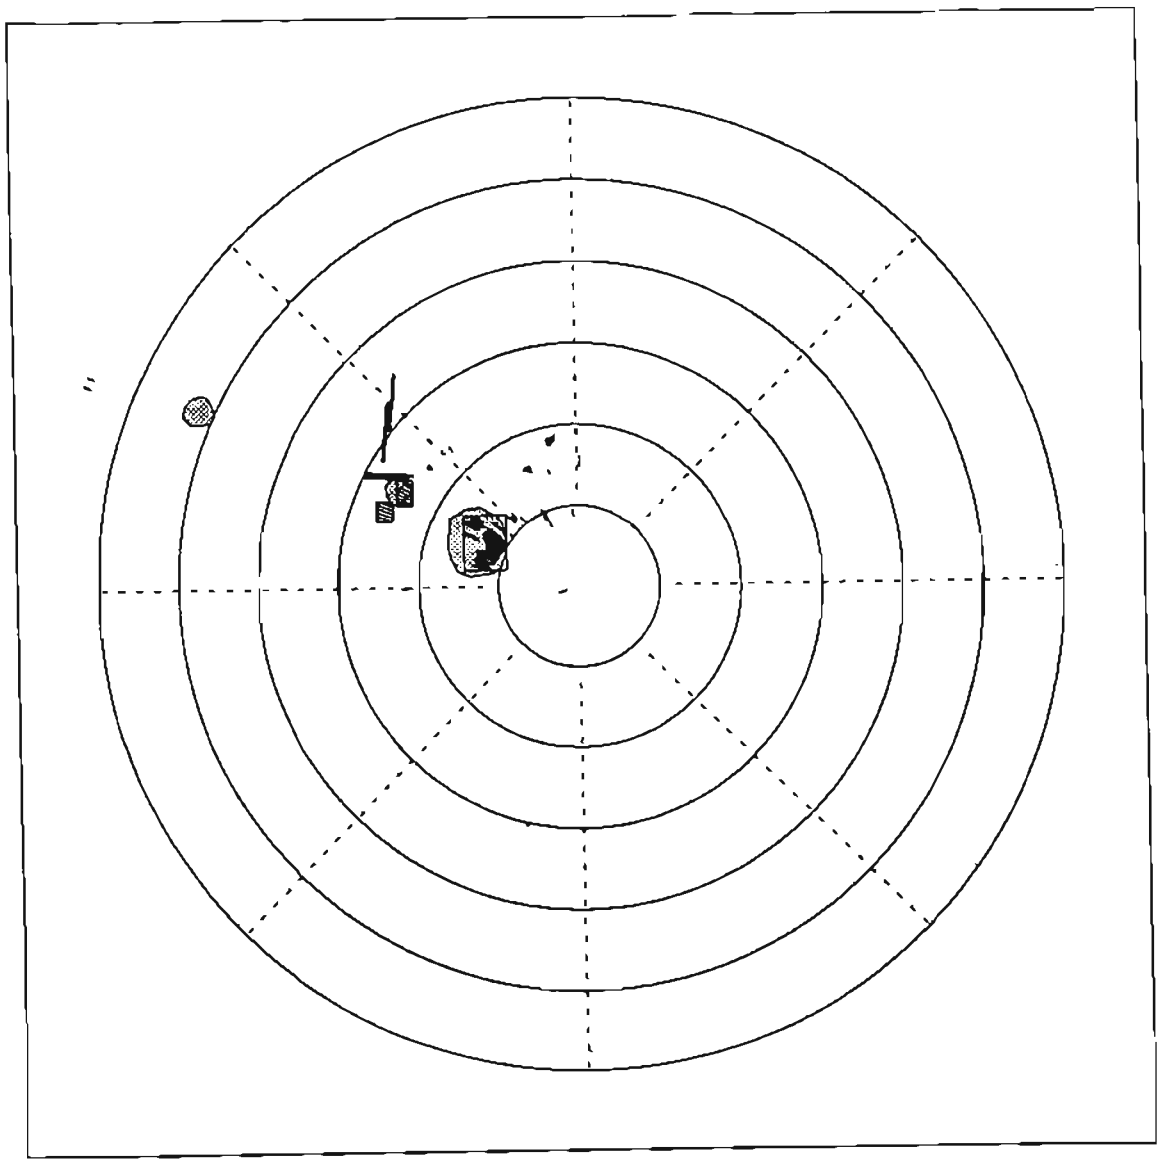

Figure 58: Computational algorithm results for 11 July 1988 case at radar scan time 22:05:04.

Chapter 5, may be necessary to properly evaluate the tradeoffs between these two classes of detection algorithm.

7.7. Additional Concepts Worth Exploring

The classification algorithm described above exploits the temporal continuity of outflow events to improve the detection process, but it makes rather weak use of the spatial continuity of the outflow signature. One possible approach for extracting more information

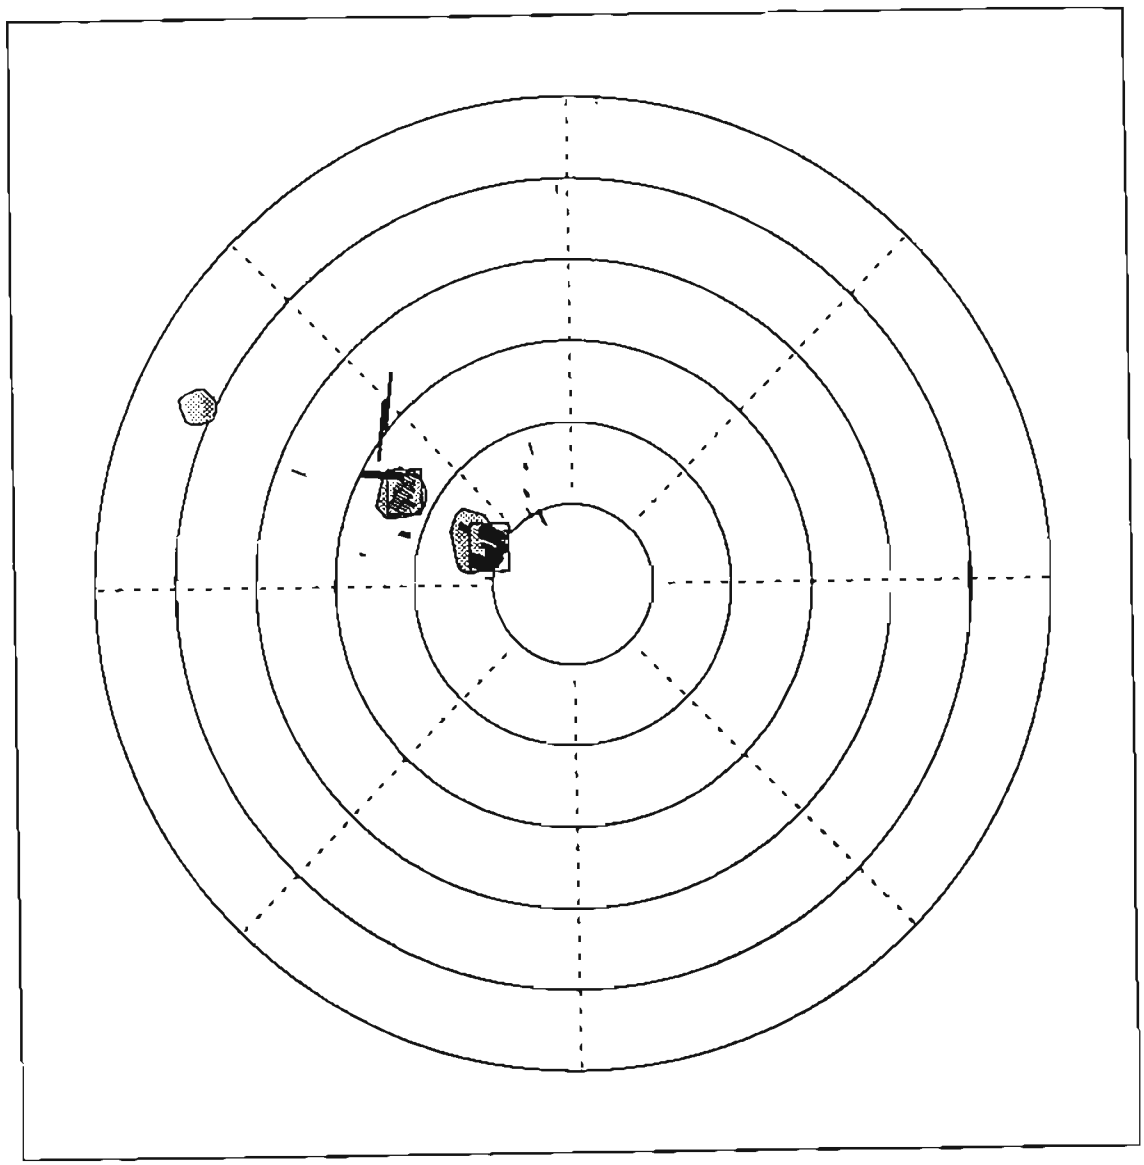

Figure 59: Computational algorithm results for 11 July 1988 case at radar scan time 22:06:01.

from the spatial distribution of shear would be to correlate the velocity field with a microburst outflow template to form a coefficient of match. This approach has been explored for the detection of rotation regions within the downdraft of microburst storm cells [Stillson, 1989]. A correlation coefficient might be very useful in the shear classification process, i.e., to determine if a point shear value should be included in a divergence region or not. This factor could be integrated with the temporal threshold adjustment to achieve an improved region definition.

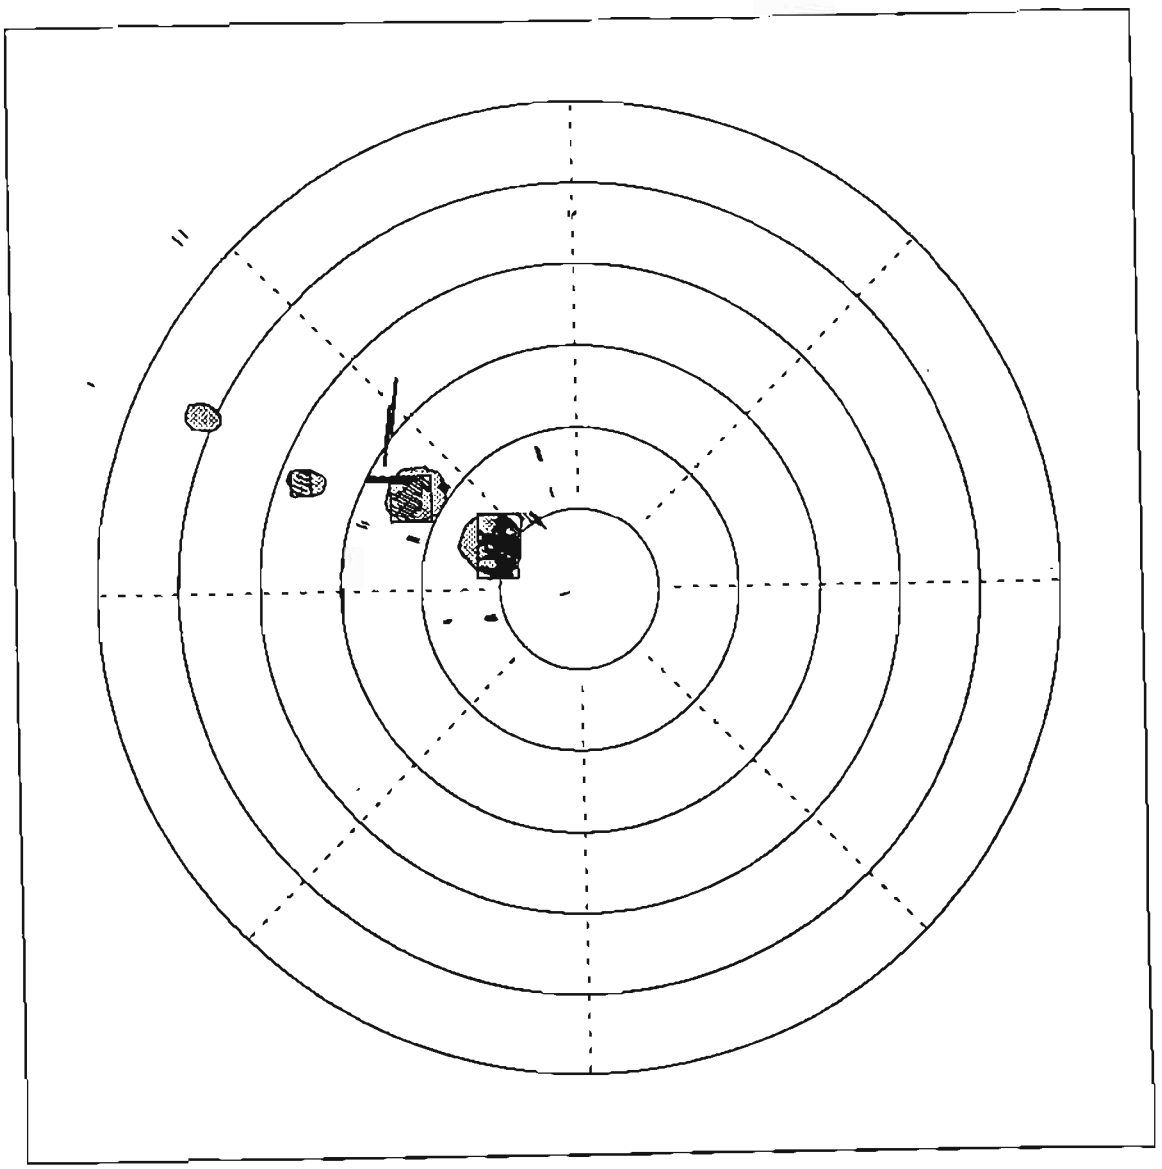

Figure 60: Computational algorithm results for 11 July 1988 case at radar scan time 22:06:58.

The analysis of filtering methods presented in this chapter concluded that smoothing filters can significantly attenuate the detected strength of small microbursts and hence are quite costly in terms of performance degradation. The benefits of the smoothing, however, include an improved ability to detect weaker, extended shear regions. The detection of these regions is quite important, given the possible effects of microburst asymmetry. Adaptive filtering methods may exist which could provide a better tradeoff between attenuation and smoothing than the simple filters studied here. An example of one candidate filter is that

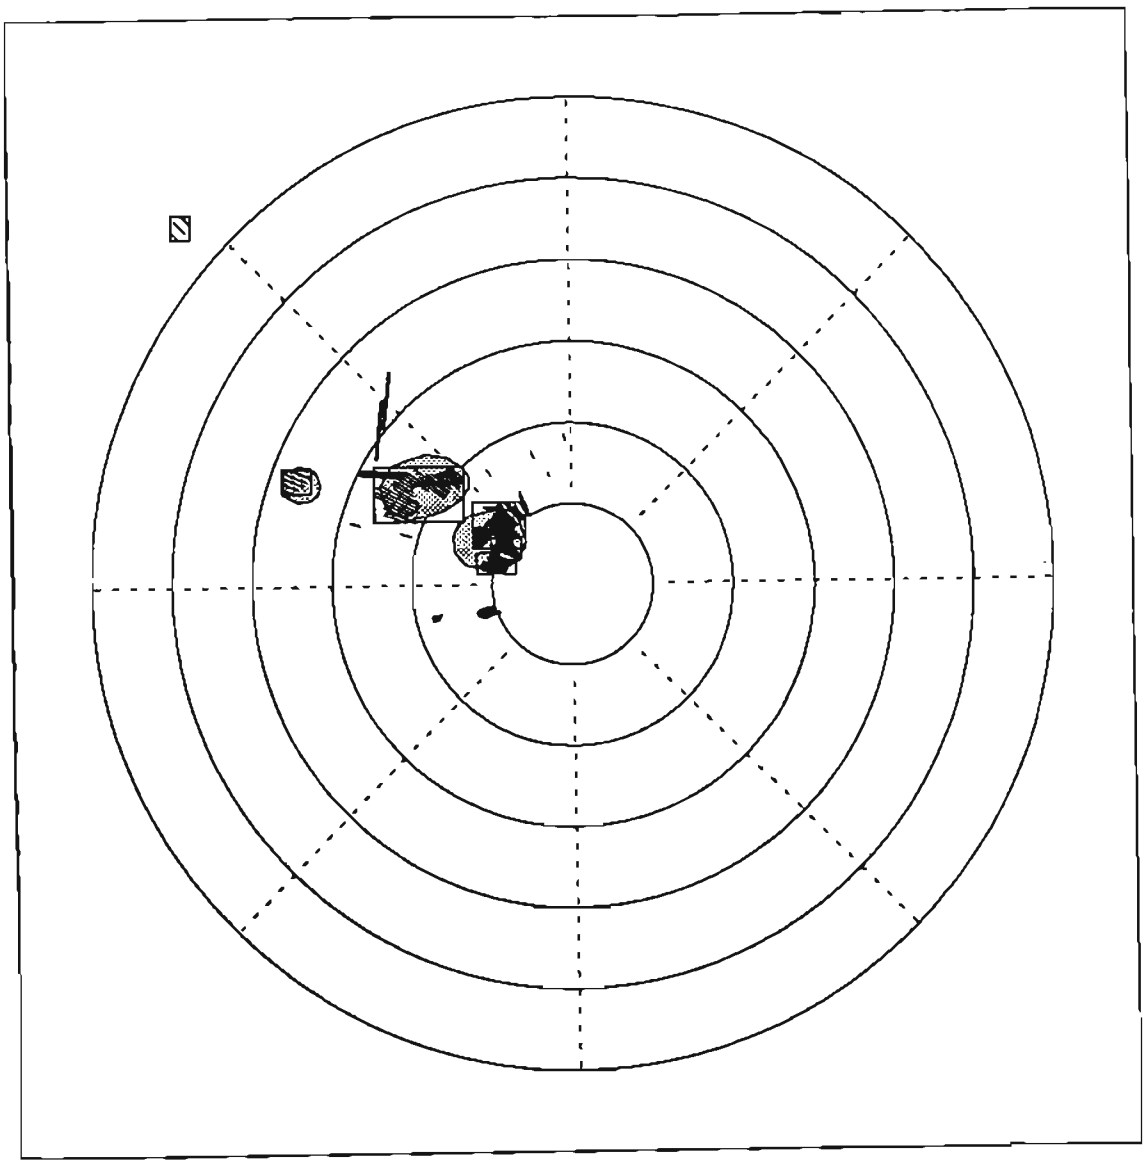

Figure 61: Computational algorithm results for 11 July 1988 case at radar scan time 22:08:03.

described by [Saint-Marc and Medioni, 1988]. This filter performs an iterative smoothing operation where the smoothing effect is regulated by the local derivative of the signal being smoothed. The result of this adaptation is the partitioning of the signal into piecewise constant sections, without smoothing out the sharp transitions between sections. If this algorithm were applied to the raw shear field, prior to shear classification, it may reduce unwanted variations without significantly disturbing the strong shear sections. Operation of the filter requires the choice of two parameters: a spatial scale parameter and the number of

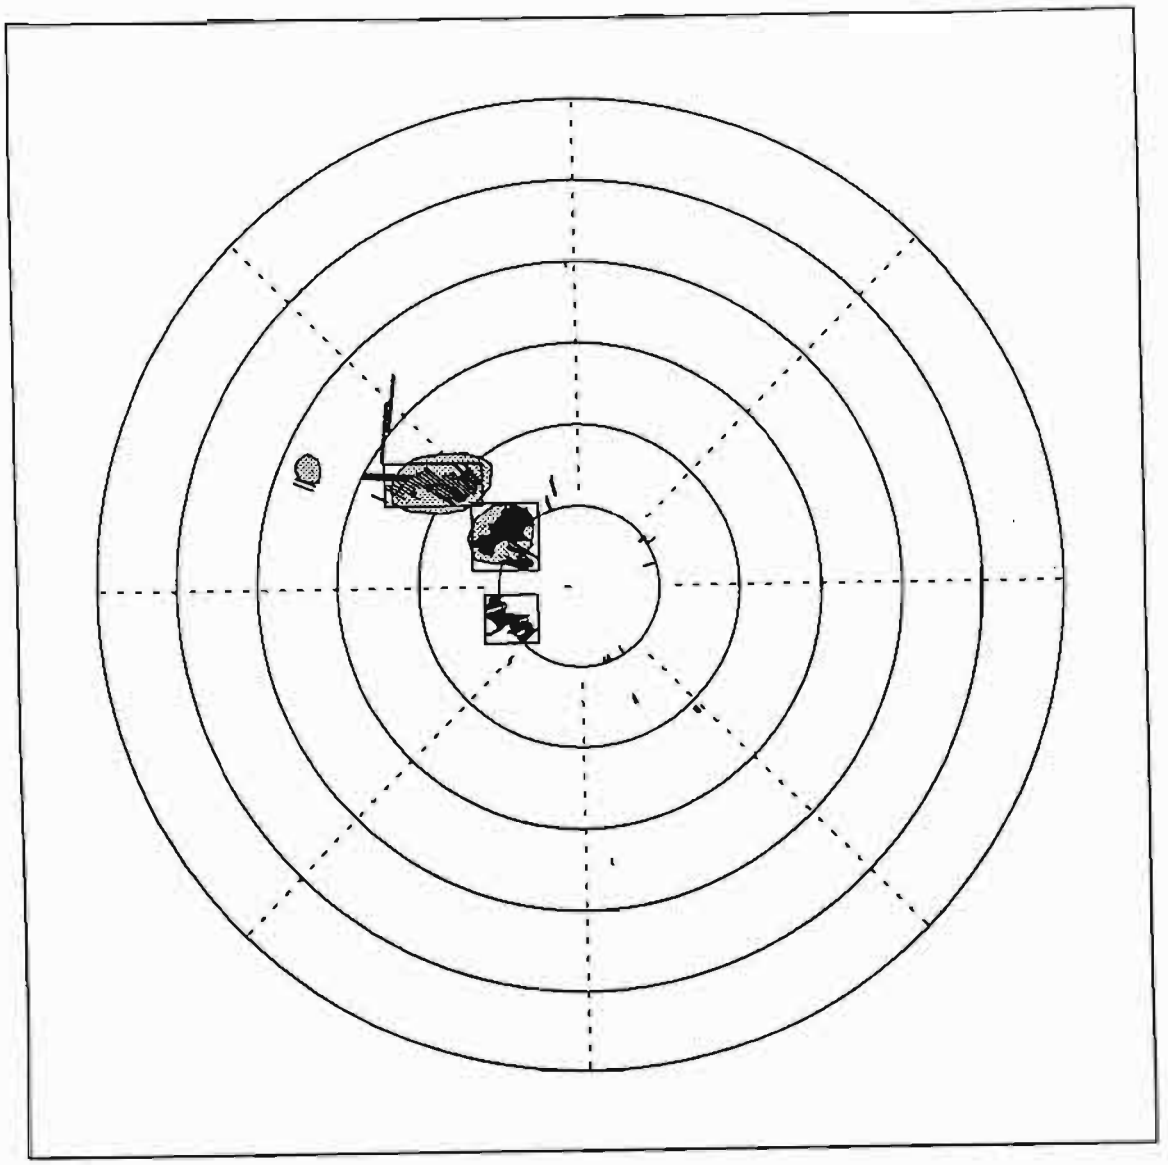

Figure 62: Computational algorithm results for 11 July 1988 case at radar scan time 22:09:00.

iterations to perform. These two parameters allow the spatial resolution and the degree of smoothing to be chosen independently; the mean and median filters do not have this desirable property. An experimental implementation of this approach showed mixed results: on some signatures the filter performed extremely well, but it occasionally produced very anomalous results. It appears that the choice of parameters has a substantial impact on the performance of the filter, and that the variance of the smoothed shear field (from one noise realization to the next) is quite high. Refined variations on this adaptive scheme may

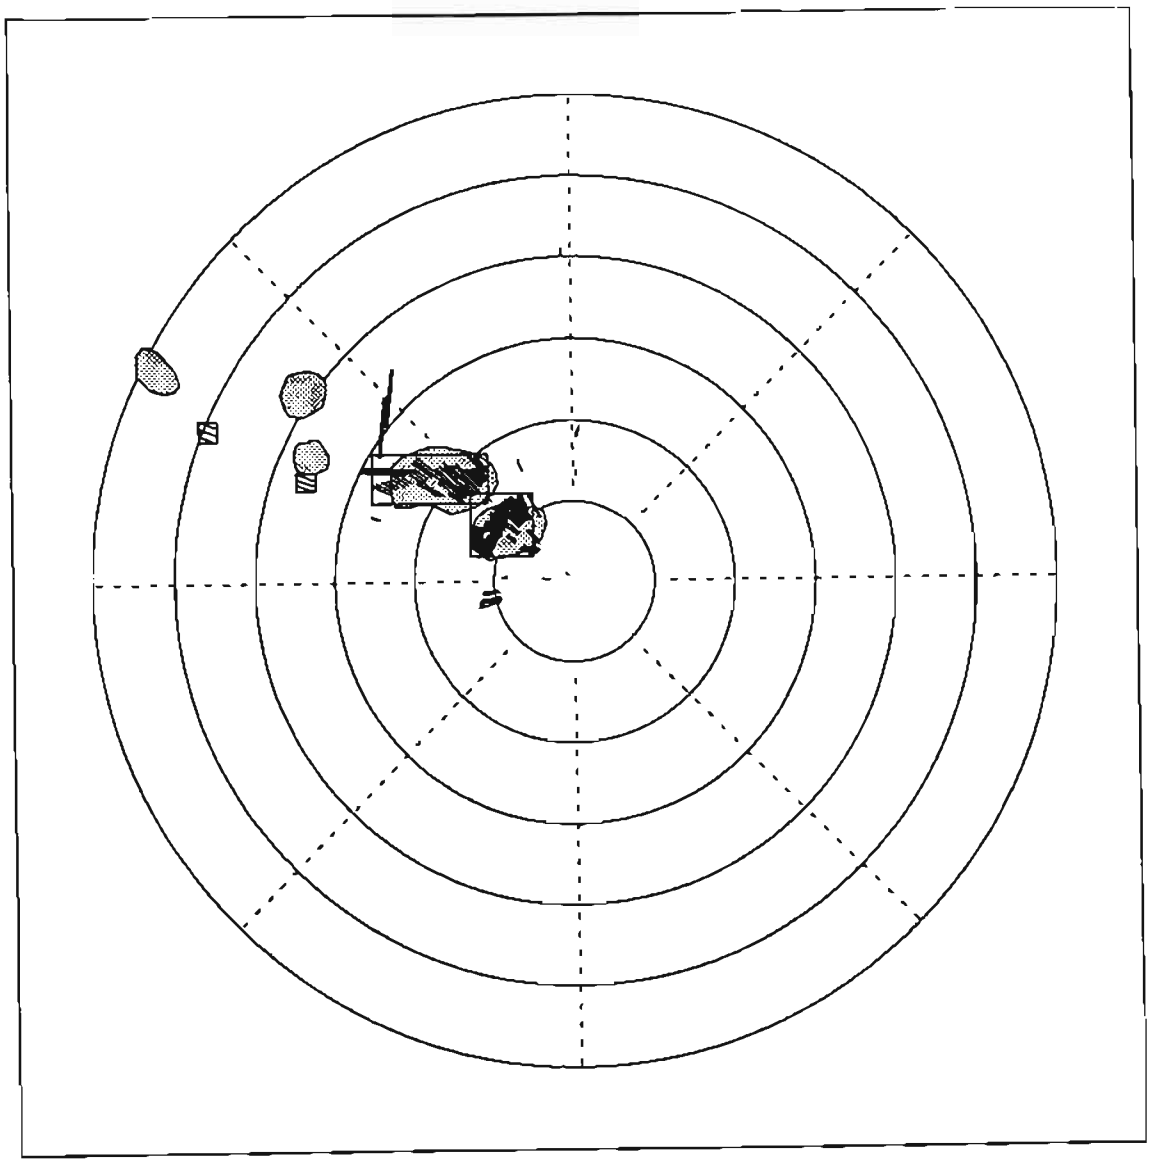

Figure 63: Computational algorithm results for 11 July 1988 case at radar scan time 22:10:03.

well be quite successful at improving the detectability of shear regions, but the computational cost of this technique would appear to be quite high.

A final technique which may warrant study is the "direction of gradient" technique, described in [Zhou et al., 1986]; In this approach, the radial velocity field is used to compute a vector gradient field, and the direction of the gradient vector is computed at each point. Experience in other image processing applications has indicated that the gradient direction field is often more consistent than the gradient magnitude field and that detection of a

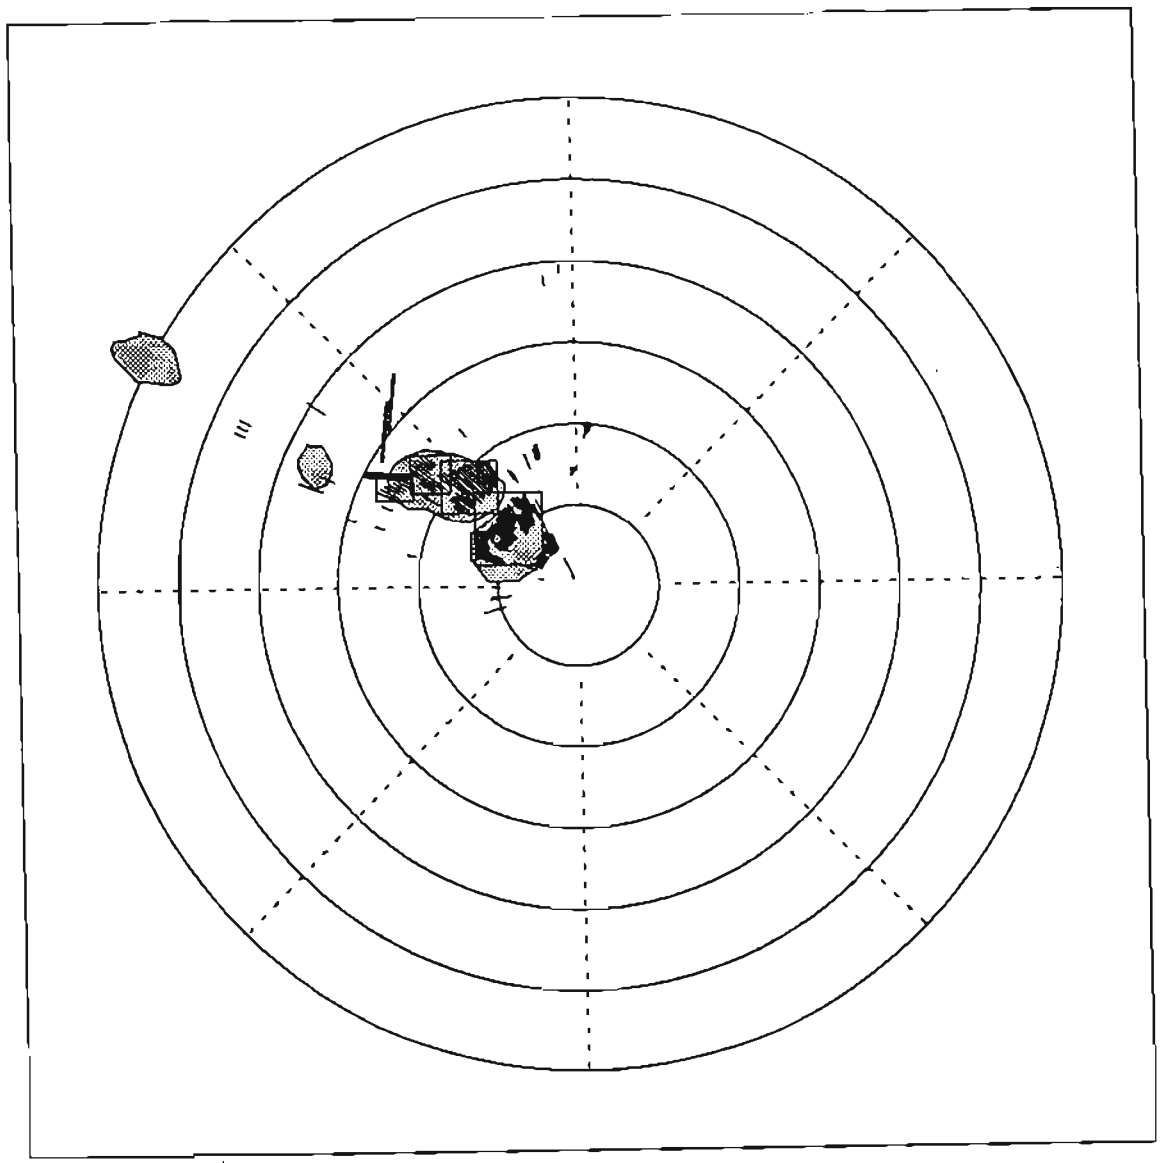

Figure 64: Computational algorithm results for 11 July 1988 case at radar scan time 22:11:00.

divergent region might be more reliable from this information. The fact that the radar measurement of the windfield only obtains the radial component could potentially defeat the benefit of the gradient direction approach, but the use of a spatially extended, strength-independent quantity could provide improved detection capability, particularly for weak or asymmetric events.

7.8. Summary of Computational Alternatives

This chapter has introduced a number of alternative concepts for the basic detection of shear segments using computational techniques tied to the estimation of a shear field. The use of a very simple shear estimation and thresholding approach was shown to be too sensitive to the variability present in realistic signatures, and more sophisticated shear thresholding techniques were discussed. The primary goal of any alternative algorithm approaches should focus on exploiting the spatial and temporal continuity present in the outflow signatures. The use of shear-based computational methods is most attractive because of the strong relationship between shear and the actual aircraft hazard level, and care should be taken in the basic shear segment detection stage to preserve this hazard-related information and allow for the accurate delineation of the hazard extent in subsequent algorithm processing stages.

The techniques discussed here are no more than exploratory algorithm concepts and have not yet been subjected to any form of quantitative performance evaluation. To adequately evaluate the performance differences between the baseline TDWR algorithm and the class of computational algorithms described in this chapter, the path-based scoring technique of Chapter 5 would be required. Simple hit/miss scoring methods, as used in previous algorithm performance analyses, would not be adequate to distinguish the operational differences between these algorithms. Such an evaluation requires a significant number of high-quality dual-doppler data cases and a major commitment of both computer and analyst resources.

8. SUMMARY

The divergence outflow algorithm is the primary component of the TDWR microburst detection system and is responsible for identifying surface divergence regions based on the radar velocity measurements. This algorithm is the result of several years of development and evaluation, based extensively on actual testbed radar system measurements.

The divergence algorithm operates as one component in a chain of processing stages, preceded by data quality pre-processors and followed by sophisticated feature integration and validation stages. Each stage in this chain has a significant effect on the performance of the overall system, but the ability to detect a microburst outflow is determined largely by the detection rate of the divergence algorithm studied here. Performance of this algorithm has been evaluated using a large sample of measurements from the Denver, 1988 field measurement program, where algorithm detections were compared to ground truth information to compute both probability of detection and false alarm. Ground truth information was obtained by a careful manual analysis of the radar velocity and reflectivity measurements to locate actual microburst regions. This performance analysis indicated a POD of roughly 90 percent for all microbursts with a radar-measured strength of 10 m/s or greater. The POD increases rapidly with increasing microburst strength with 98 percent of the microbursts with radar-measured strengths of 15 m/s or greater detected.

The PFA of the divergence algorithm was approximately 30 percent for this evaluation data set, but this alarm rate is reduced substantially by the subsequent processing stages of the overall microburst system. Evaluations of the overall system, using a similar methodology and data set, resulted in an overall system PFA of five to seven percent. These performance figures meet the minimum requirements of the TDWR System Requirements Statement and have been used to validate the TDWR system design for the purposes of system procurement and deployment.

While the demonstrated performance of the divergence detection algorithm is adequate for meeting the basic system requirements there are several considerations which indicate that improvements to the performance of this algorithm should be sought. Chief among these issues is the impact of microburst outflow asymmetry on detection performance.

The asymmetric nature of the microburst outflow can result in a radar-measured strength which is significantly different from that along the runway direction (i.e., that which affects the safety of aircraft flight). Since the detection rate of the divergence algorithm drops off rapidly as the radar-oriented strength falls below about 12 m/s, it is quite possible for a microburst which is strong in another direction to be poorly detected as observed by the radar. The performance impact of asymmetry as a function of radar viewing angle relative to runway direction has been examined using observed microburst asymmetry distributions and algorithm detection statistics. For microbursts with a runway-oriented strength above 15 m/s, the basic detection performance results (POD of 98 percent) apply if the radar is sited so as to view directly along the runway. As the radar viewing angle is rotated relative to the runway direction, the radar-oriented strengths no longer match the runway-oriented strengths because of outflow asymmetry. While the average difference between the radar and runway strengths is zero (because exactly half of the time the radar strength is greater

than the runway strength and half of the time the runway strength is greater than the radar strength), the highly nonlinear change in detection probability with strength results in a net decrease in POD. For a radar viewing angle rotated 45° from the runway orientation the POD (for runway strengths of 15 m/s or greater) drops to 83 percent. This decrease in performance is largely influenced by the divergence algorithm detection rate for weak radar-measured strengths. If the divergence algorithm were improved to provide near-100 percent probability of detection for outflows with radar-measured strengths above 7 m/s, then the system POD (over all events with runway-oriented strengths of 15 m/s or more) would increase to 96 percent.

These quantitative results are strongly dependent on both the asymmetry observations and microburst algorithm performance statistics derived from the Denver demonstrations. Qualitatively, however, these results do indicate that significant improvements in the ability to detect the weak (measured) outflows would result in a substantial increase in safety provided by the system.

Alternative detection algorithm concepts have been briefly explored, with the goals of simplifying the logical structure of the detection algorithm and improving the detection performance. The computational techniques described are based heavily on the computation of a radial shear estimate and make aggressive use of the temporal and spatial continuity of the outflow signature. The limited number of cases used for the exploration of these algorithm techniques are not adequate for quantifying the performance of the computational approach but do suggest that a more comprehensive performance evaluation for this class of algorithm is warranted. Such an evaluation would require the application of the technique to a larger set of test cases and would also require the use of detailed performance analysis techniques (such as the path-based scoring approach) to adequately measure the performance improvements obtained.

The studies presented in this report provide an important insight for the development of future, more advanced detection algorithms, namely: that the primary challenge in the detection process is the discrimination and classification of weak outflows in complex surface windfields, not the detection of simple idealized signatures in a uniform noise background. Simulation studies using an idealized conceptual model have demonstrated that even weak shear signatures can be reliably detected in the presence of realistic noise levels by even the most basic detection algorithms. Actual windfield measurements present a much more complex environment, where multiple interacting outflows of assorted strengths and spatial scales are a common occurrence. In these complex cases the performance of a detection algorithm is dominated by its ability to accurately identify the extent of the various shear regions and to appropriately characterize their strengths. Frequently, an objective definition of the "correct" result is hard to provide in these cases, even if the exact, true windfield is known. In these complex scenarios it may be possible to estimate the flight performance loss for a particular path, given the detailed aircraft state vector and pilot response to the shear. Attempts to define hazard regions without knowledge of these detailed aircraft parameters is often an ill-defined problem and the detection performance of a divergence algorithm may well be limited by fundamental ambiguities in the definition of what the output of a "perfect" detection algorithm should be. Any efforts to improve detection performance beyond that obtained by the current baseline TDWR divergence algorithm should confront the existing ambiguity in what constitutes the

"correct" answer in these complex windfield scenarios as a primary focus. Meaningful evaluations of performance improvements can only be obtained with far more sophisticated definitions of "hazard" than those used in the past. The path-based windfield scoring methods discussed in this report suggest a possible direction for an evaluation metric.

9. REFERENCES

Steven D. Campbell, "Use of features aloft in the TDWR microburst recognition algorithm," *Preprints, 24th conference on radar meteorology*, American Meteorological Society, Tallahassee, FL, (1989), pp. 167–170.

Steven D. Campbell and M.W. Merritt, "Advanced microburst recognition algorithm," MIT Lincoln Laboratory, Lexington, MA, Project Report ATC–145, (1988).

Steven D. Campbell, M.W. Merritt, and J.T. DiStefano, "Microburst recognition performance of TDWR operational testbed," *Third international conference on the aviation weather system*, American Meteorological Society, Anaheim, CA, (1988).

Steven D. Campbell, Personal communication, (1990).

J.T. DiStefano, "Observability of microbursts using doppler weather radar and surface anemometers during 1987 in Denver, CO," MIT Lincoln Laboratory, Lexington, MA, Project Report ATC–161, (1988).

J.T. DiStefano and D.A. Clark, "Microburst observability and frequency during 1988 in Denver, CO," MIT Lincoln Laboratory, Lexington, MA, Project Report ATC–170, (1990).

Michael D. Eilts, "Use of a single Doppler radar to estimate the runway wind shear component in microburst outflows," *AIAA 26th Aerospace Sciences Meeting*, Reno, Nevada, (January, 1988).

K.L. Elmore and W.R. Sand, "A cursory study of F–factor applied to Doppler radar," *3rd International Conference on the Aviation Weather System*, Anaheim, CA, American Meteorological Society, Boston, MA, 1989.

J.E. Evans and D. Turnbull, "Development of an automated windshear detection system using Doppler weather radar," *Proc. of the IEEE*, vol. 77, No. 11, pp. 1661–1673, (November, 1989)

J.E. Evans, ed., "Results of the Kansas City 1989 terminal doppler weather radar (TDWR) operational evaluation testing," MIT Lincoln Lincoln Laboratory, Lexington, MA, Project Report ATC–171, DOT/FAA/NR–90/1, (17 August 1990).

Robert G. Hallowell, "Aspect angle dependence of outflow strength in Denver microbursts: spatial and temporal variations," *16th Conference on Severe Local Storms*, Alberta, Canada, American Meteorological Society, (October, 1990).

Mark R. Hjelmfelt, "Structure and life cycle of microburst outflows observed in Colorado," *J. Appl. Meter.*, vol. 27, no. 8, pp. 900–927, (August 1988).

W. Mahoney, W. Wilson, K. Brislawn, and M.K. Politovich, "Microburst detection by TDWR: performance assessment," *Third international conference on the aviation weather system*, American Meteorological Society, Anaheim, CA, (1989).

Terri A. Noyes, "Divergence detection in wind fields estimated by an airport surveillance radar," MIT Lincoln Lincoln Laboratory, Lexington, MA, Project Report ATC–172, DOT/FAA/NR–90/2, (15 October 1990).

Philippe Saint-Marc and G. Medioni, "Adaptive smoothing for feature extraction," *Proceedings: Image Understanding Workshop*, DARPA, pp. 1100-1113, (April, 1988).

Herbert W. Schlickenmaier, "Windshear case study: Denver, Colorado, July 11, 1988," FAA Report DOT/FAA/DS-89/19, (November 1989).

J.H. Stillson (internal memorandum, 1989).

Russell Targ and R.L. Bowles, "Investigation of airborne lidar for avoidance of windshear hazards," *AIAA Conference on Sensor and Measurment Techniques for Aeronautical Applications*, Atlanta, GA, (September, 1988).

D.J. Turnbull, J. McCarthy, J. Evans, and D. Zrnic', "The FAA terminal Doppler weather radar (TDWR) program," *Third international conference on the aviation weather system*, American Meteorological Society, Anaheim, CA, (1989).

Mark E. Weber, "Dual-beam autocorrelation based wind estimates from airport surveillance radar signals," MIT Lincoln Lincoln Laboratory, Lexington, MA, Project Report ATC-167, DOT/FAA-PS-89/5, (21 June 1989).

F.W. Wilson, Jr., R.K. Goodrich, K. Brislawn, "Enclosing shapes for single Doppler radar features," accepted for publcation by the *Journal of Atmospheric and Oceanic Technology* (April 1991).

James W. Wilson, R. Roberts, C. Kessinger, and J. McCarthy, " Microburst wind structure and evaluation of Doppler radar for airport wind shear detection," *J. Climate and Applied Meteor.*, vol. 23, pp. 898-915, (June 1984).

Y.T. Zhou, R. Chellappa, V. Venkateswar, "Edge detection using the directional derivatives of a space varying correlated random field model," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Miami Beach, FL, pp. 115-121, (22-26 June 1986).

10. LIST OF ABBREVIATIONS

AEL	Algorithm Enunciation Language
AGC	Automatic Gain Control
ASR-9	Air Surveillance Radar
ATC	Air Traffic Control
FAA	Federal Aviation Administration
FL-2	FAA / Lincoln Laboratory testbed radar system
GSD	Geographical Situation Display
LLWAS	Low Level Windshear Alert System
MIT	Massachusetts Institute of Technology
NCAR	National Center for Atmospheric Research
OT&E	Operational Test and Evaluation
PFA	Probability of False Alarm
POD	Probability of Detection
PRF	Pulse Repetition Frequency
RDA	Radar Data Acquisition
RPG	Radar Products Generator
SNR	Signal to Noise Ratio
STC	Sensitivity Time Control
TDWR	Terminal Doppler Weather Radar
TRACON	Terminal Radar Control
UND	University of North Dakota

APPENDIX 1.
PERFORMANCE DATA FOR 1988 SCORING

This appendix illustrates the minute–by–minute results of the scoring analysis conducted using the three microburst days from the 1988 operational demonstration in Denver, CO. The figures shown below are indexed by time on the horizontal axis, labelled on the bottom in UTC. Each page contains three sections, corresponding to the three case days: June 10, June 21 and June 25, 1988 (top to bottom, respectively).

Each microburst event present in the ground truth database is depicted for each minute it was observed by the radar. The strength of the microburst is plotted at each minute interval, and the strength value is enclosed in a box if the event was considered eligible for scoring on that minute. A microburst may be ineligible if it is too close to the radar (within 6 km), too far from the radar (outside 30 km), too weak (below 10 m/s) or if it lies outside the azimuthal edge of the TDWR sector scanning region. Each microburst has a symbolic name (e.g., "MB5") which is shown to the left of the first observation of the event. If a microburst was detected by the divergence region algorithm, the box for the event is filled with a stippled pattern; otherwise, the box is left empty. No indication is made for false alarms in this appendix.

This graphical summary of the ground truth and detection performance concisely presents both the temporal distribution of the actual microburst events and the detection performance of the algorithm. A quick scan of the appendix will verify that indeed the algorithm detects a very large percentage of the events, and most of the misses are for weak events.

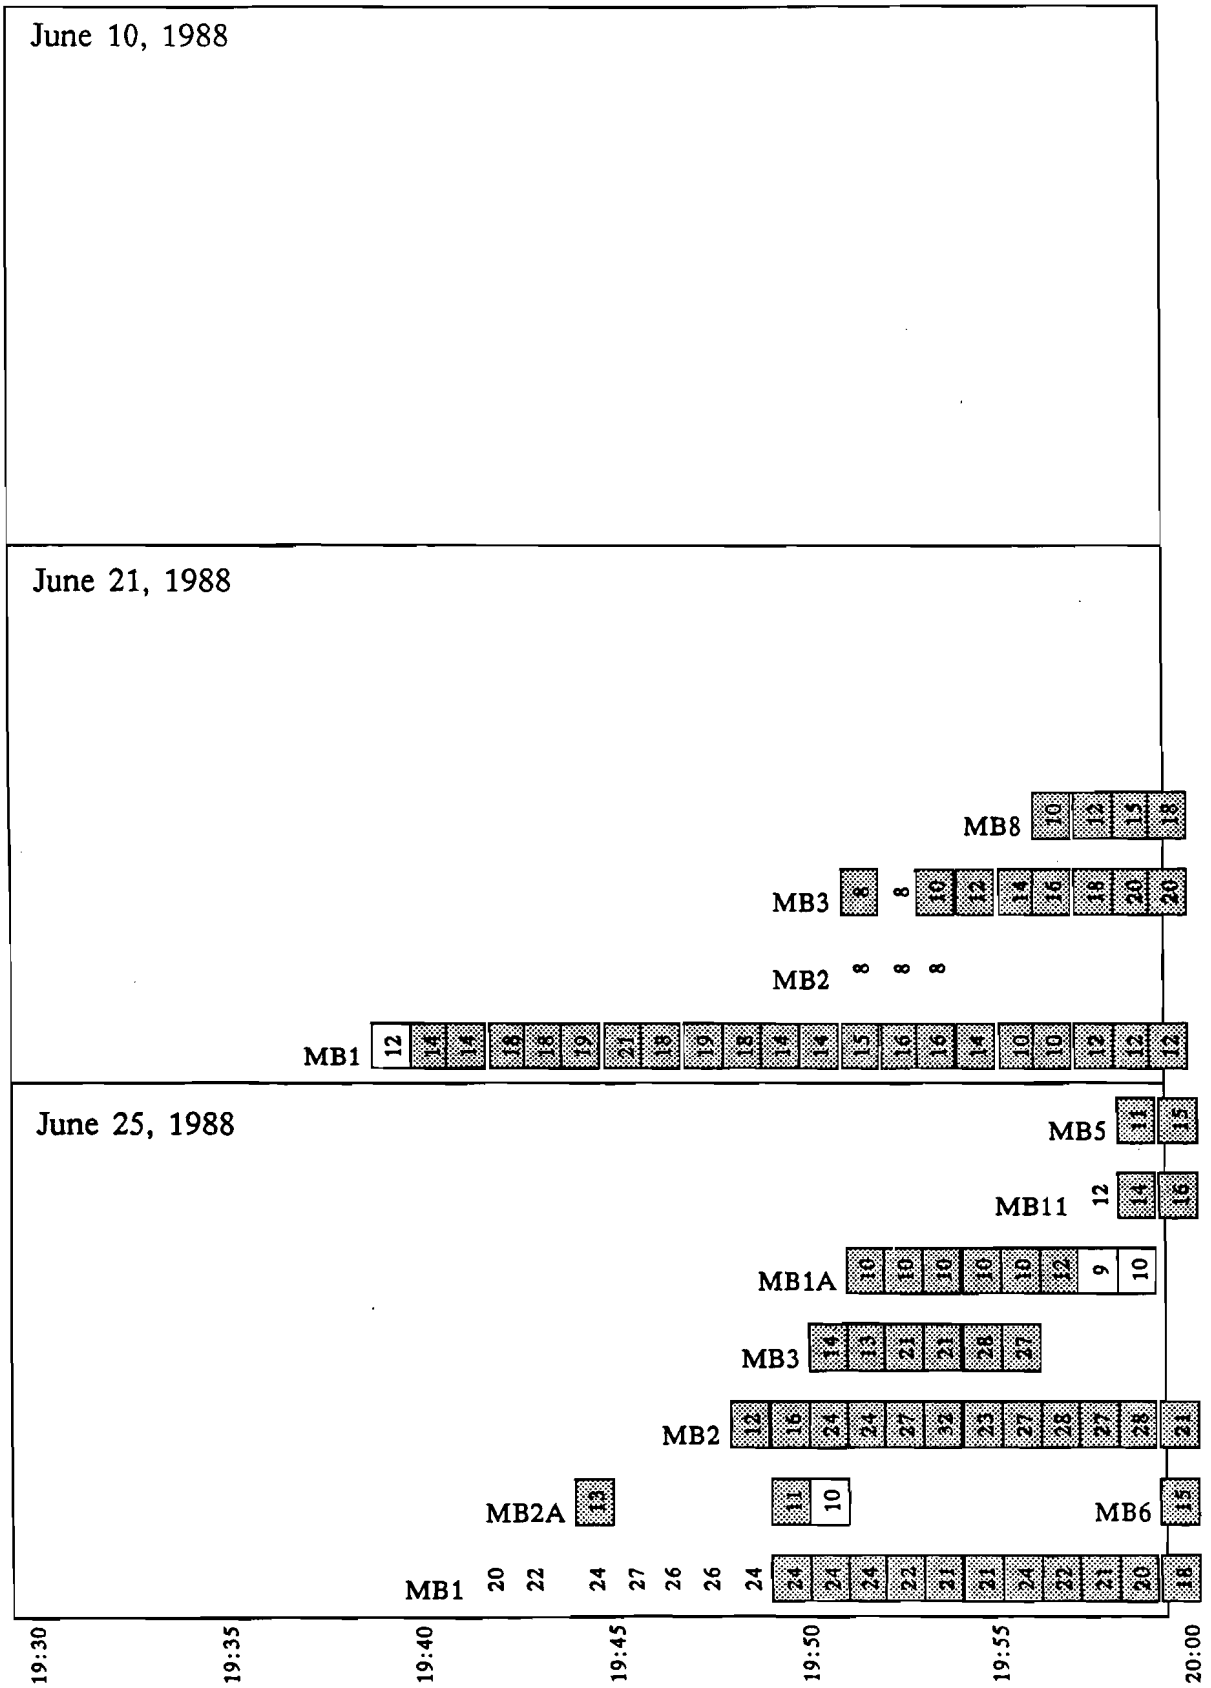
June 10, 1988
June 21, 1988
MB8
MB3
MB2
MB1
June 25, 1988
MB5
MB11
MB1A
MB3
MB2
MB2A
MB6
MB1
19:30
19:35
19:40
19:45
19:50
19:55
20:00

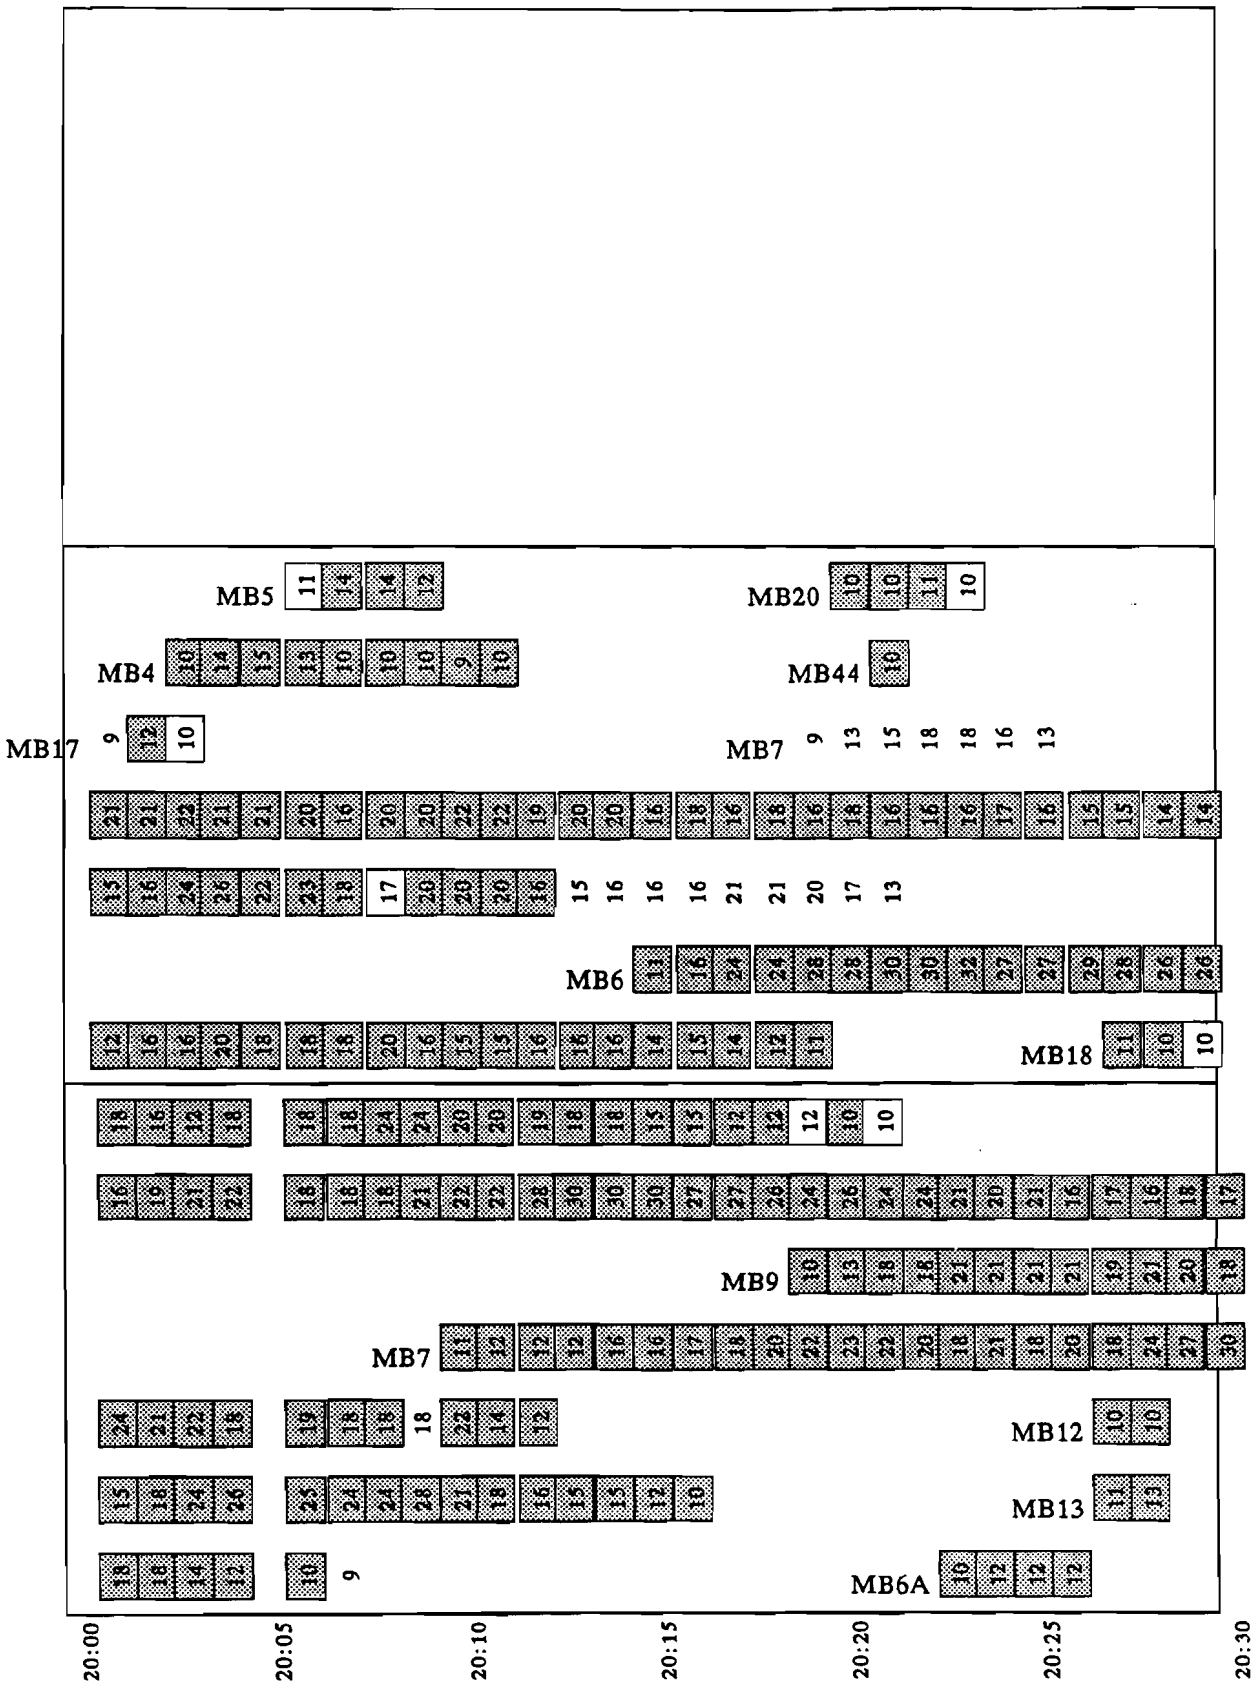
MB5
MB20
MB4
MB44
MB17
MB7
MB6
MB18
MB9
MB7
MB12
MB13
MB6A
20:00
20:05
20:10
20:15
20:20
20:25
20:30

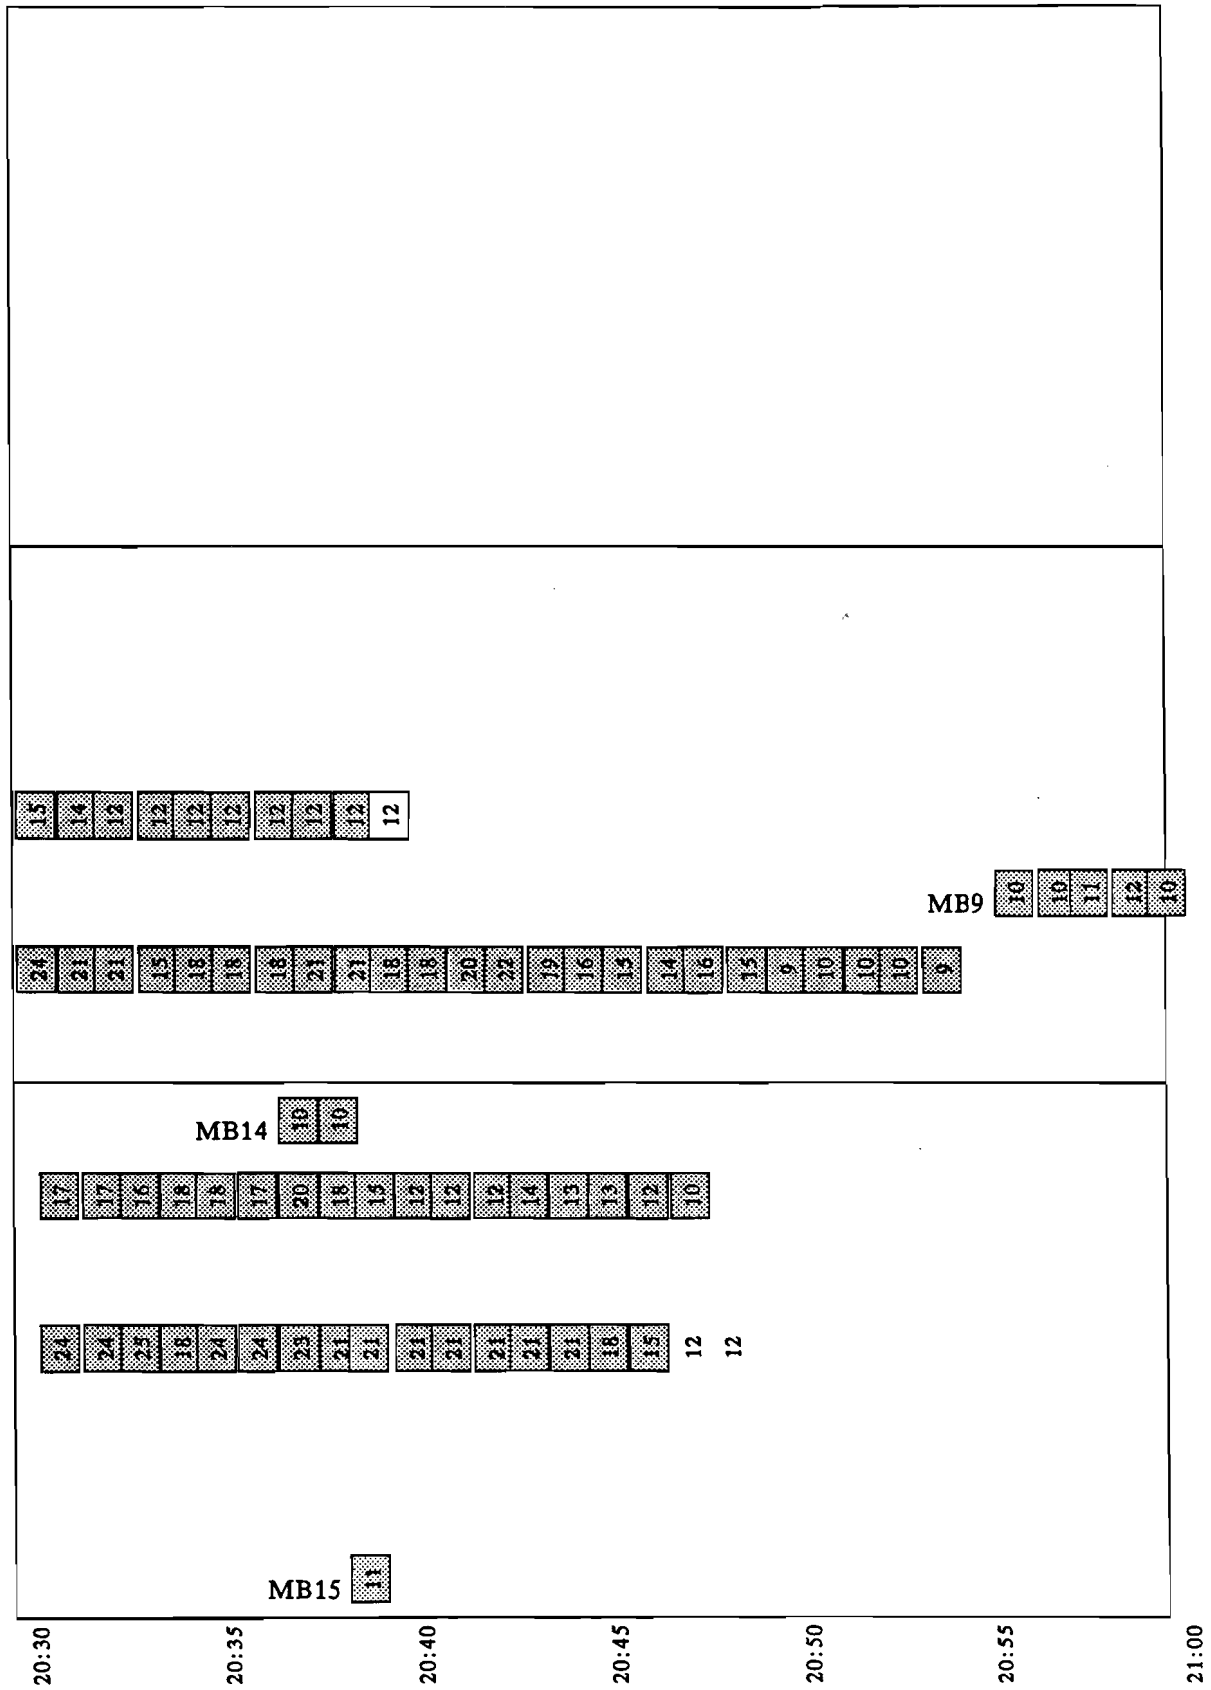
MB9
MB14
MB15
20:30
20:35
20:40
20:45
20:50
20:55
21:00

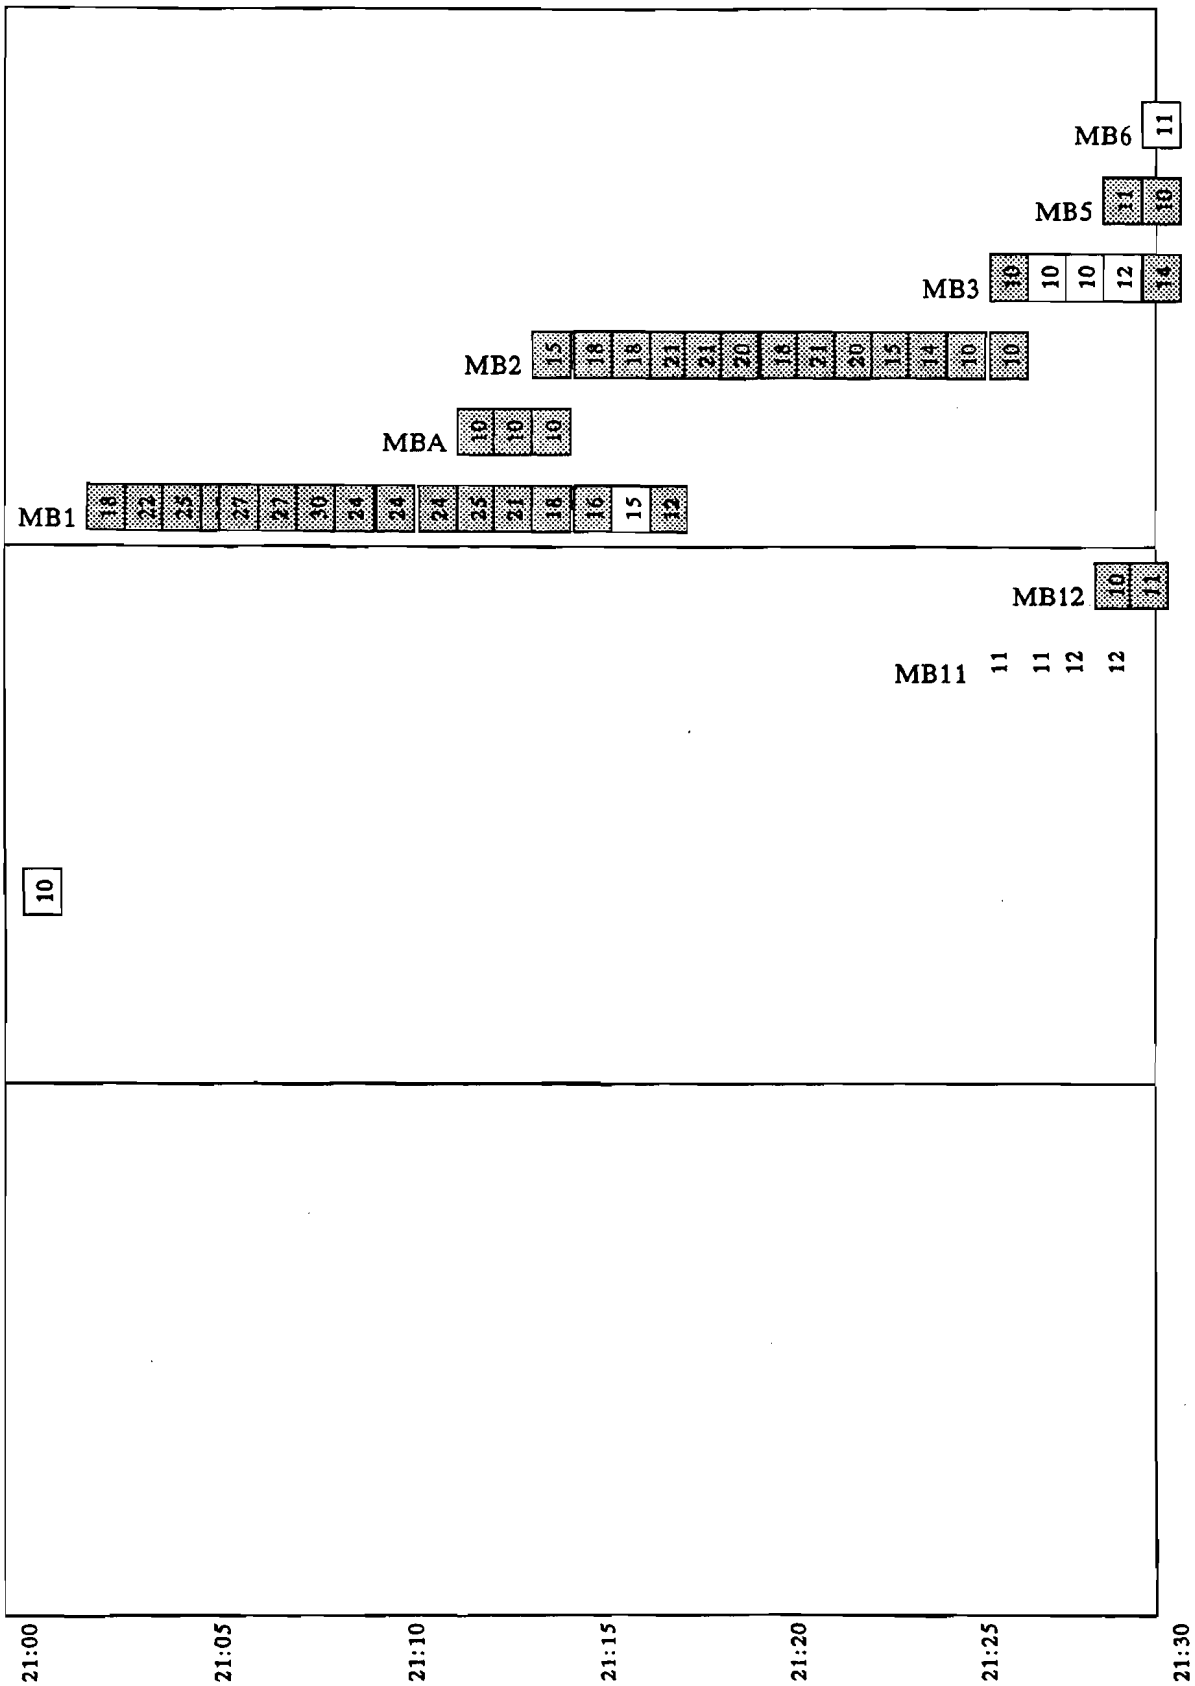
MB6
11
MB5
11
10
MB3
10
10
10
12
14
MB2
15
18
18
21
21
20
18
21
20
15
14
10
10
MBA
10
10
10
MB1
18
22
23
27
27
30
24
24
24
25
21
16
16
15
12
MB12
10
11
MB11
11
11
12
12
10
21:00
21:05
21:10
21:15
21:20
21:25
21:30

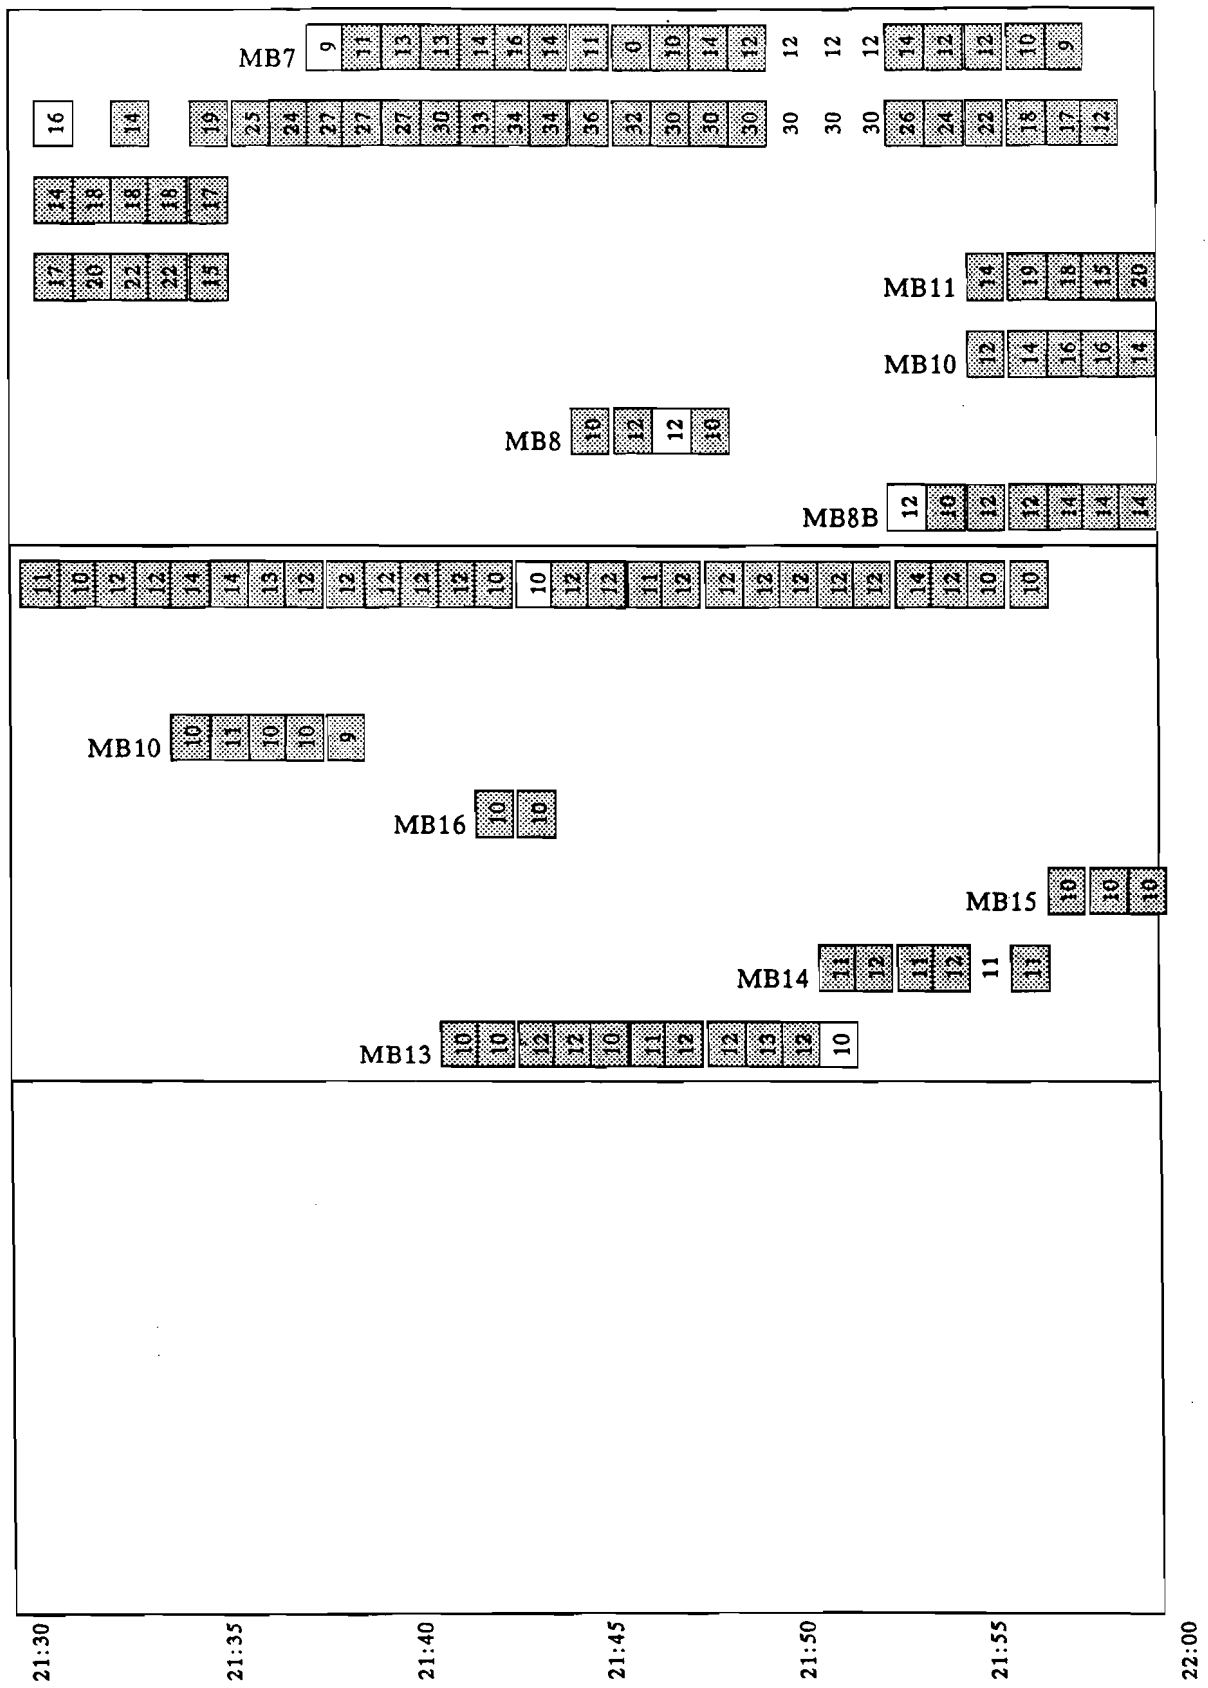
MB7
MB11
MB10
MB8
MB8B
MB10
MB16
MB15
MB14
MB13
21:30
21:35
21:40
21:45
21:50
21:55
22:00

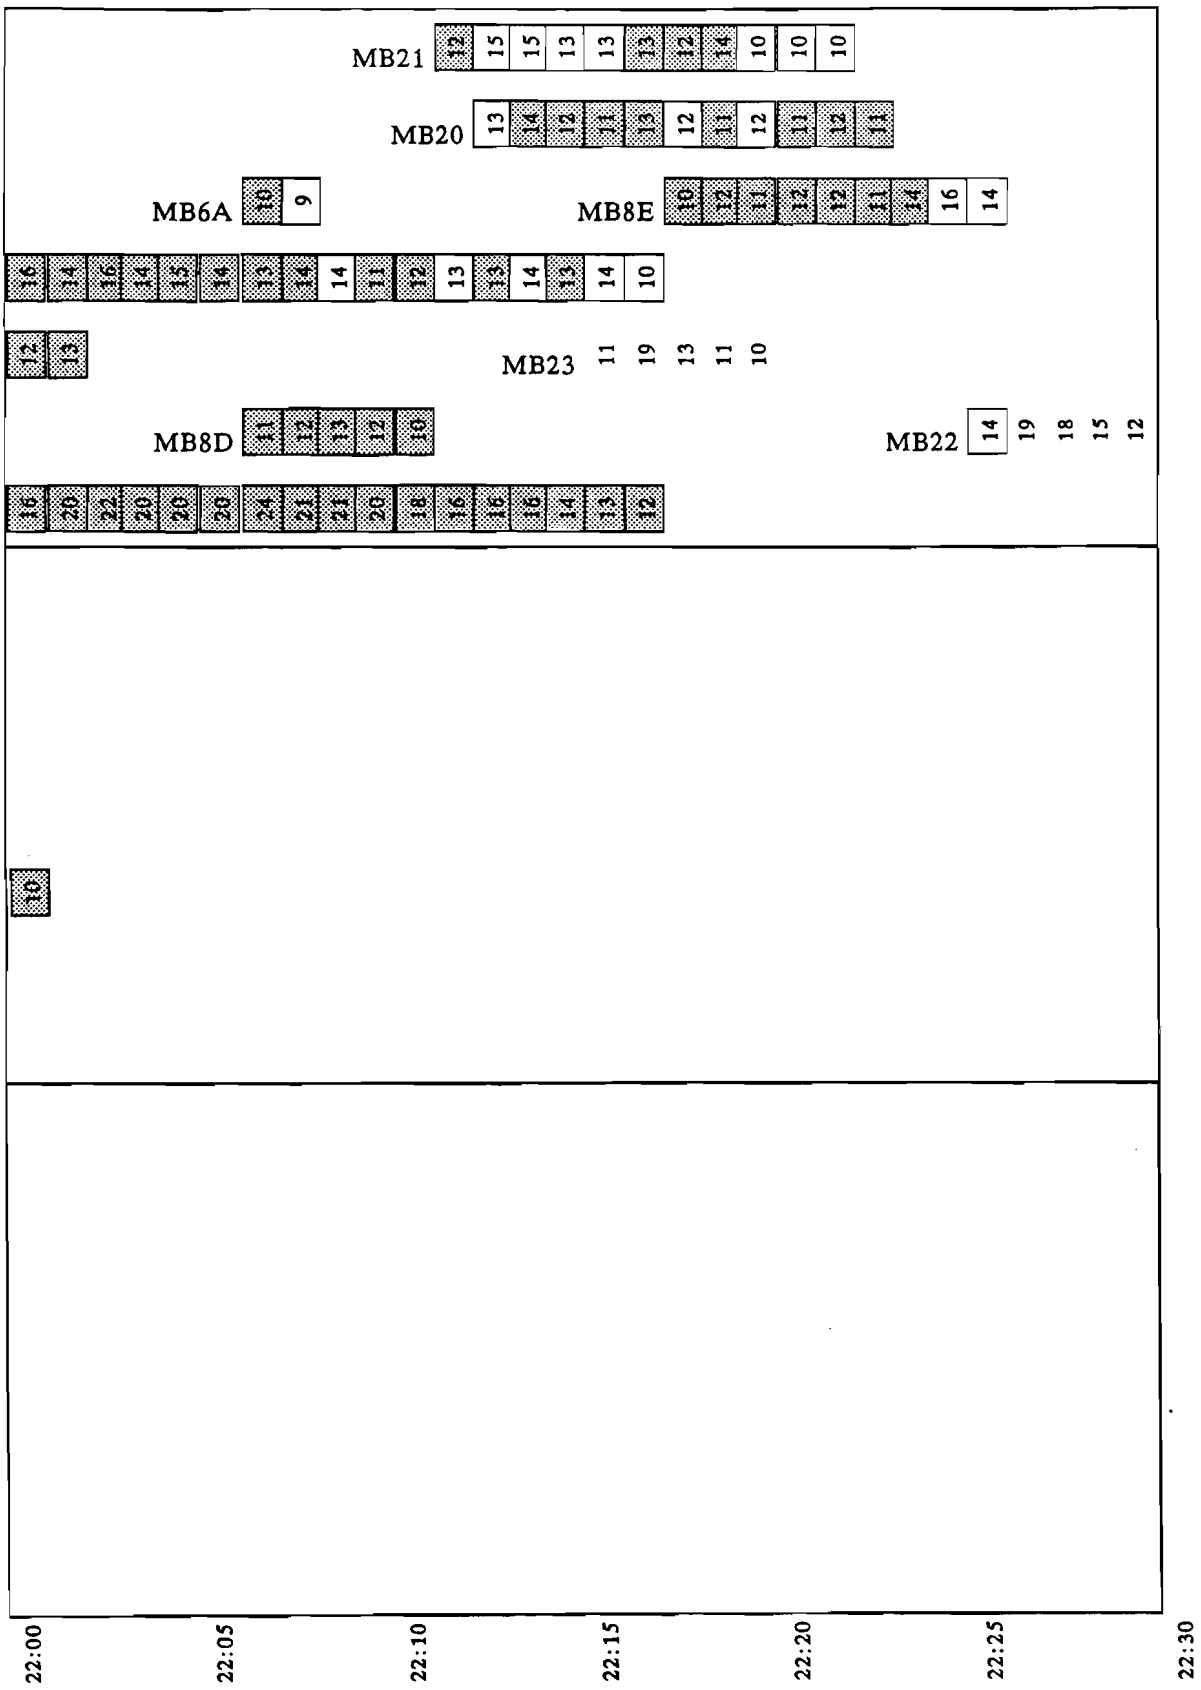
MB21
MB20
MB6A
MB8E
MB23
MB8D
MB22
22:00
22:05
22:10
22:15
22:20
22:25
22:30

APPENDIX 2.
NOISE SENSITIVITY OF SHEAR ESTIMATION ALGORITHMS

The goal of the shear estimation calculation is to compute an estimate of the rate of change of radial velocity with range at each sample point in range. The radar-measured radial velocity field has several types of noise contributions, all of which will be amplified by the differentiation process. To reduce the impact of this noise contribution, the shear estimate must be computed over some spatial window in a manner which provides a low-variance estimate without significantly attenuating the small-scale shears.

To put the shear estimation problem in perspective, consider that the minimum shear threshold of interest is roughly 2.5 m/s per km, which corresponds to a gate-to-gate velocity difference of 0.3 m/s (for TDWR range gate spacing of 125 m). The TDWR specification allows the velocity sample estimator variance for typical microburst measurements (8 dB SNR and 4 m/s spectrum width) to be as large as 1 m/s. Clearly the shears for weaker microbursts generate gate-to-gate differences which are small compared to the inherent signal noise level, and must therefore be observed over a larger spatial extent.

The two shear estimation techniques considered here are the finite difference approach and the least squares approach. Both techniques model the velocity measurements as a linearly-varying function, and estimate the slope of the line at each point. The finite difference approach simply computes the average rate of change over some fixed distance ($\Delta V/\Delta R$) while the least squares approach computes the slope of the best-fit line over some fixed distance. The least squares approach was evaluated for use in the ASR-9 windshear detection application by [Noyes, 1990].

The finite difference shear estimate, computed over a distance of 2k gates, is defined as:

$$S_{d,k}(n) \equiv \frac{V(n+k)-V(n-k)}{2k\,\Delta R} \tag{16}$$

where V(n) is the radial velocity value (m/s) at gate n and ΔR = is the spacing between range gates (km). The least-squares shear estimate, applied over a radius of k gates, is defined as:

$$S_{ls,k}(n) \equiv \frac{\sum_{i=-k}^{k} iV(n+i)}{\sum_{i=-k}^{k} i^2 \Delta R} \tag{17}$$

An important observation regarding these apparently unrelated approaches is that the finite difference estimator can be used as a basis for the least squares operator, as follows:

$$S_{ls,k}(n) = \sum_{i=-k}^{k} iV(n+i)\left[\frac{1}{\sum_{j=-k}^{k} j^2\Delta R}\right] = \sum_{i=1}^{k} i\left[V(n+i)-V(n-i)\right]\left[\frac{1}{2\sum_{j=1}^{k} j^2\Delta R}\right] \quad (18)$$

$$= \sum_{i=1}^{k}\left[\frac{V(n+i)-V(n-i)}{2i\,\Delta R}\right]\left[\frac{i^2}{\kappa}\right] \quad (19)$$

$$= \sum_{i=1}^{k} S_{d,i}(n)\left[\frac{i^2}{\kappa}\right] \quad (20)$$

This relationship gives some insight into the performance of the two estimators: the least squares estimate is largely the same as the finite difference estimator, but has the finite difference shear estimates at smaller scales averaged in as well.

If the velocity signature being processed is a perfect linear ramp (no noise) then the above two estimators will produce the same result, being the true slope of the ramp. The response of the two estimators to various forms of noise is different, and the responses of the two estimators to various forms of noise are very relevant to the optimum choice of estimator. Since both estimators are linear shift-invariant operators, the noise and true signal inputs may be considered separately, and the composite response will be the sum of the responses to the signal and noise inputs.

The two noise types considered here are white noise (uncorrelated from gate to gate) and impulse noise (a single large spike at one gate). The white noise input is representative of velocity errors resulting from estimator uncertainty, and would be most significant at low signal-to-noise ratios. The impulse noise type is typical of data contamination effects from point target or clutter interference.

For a zero-mean white noise input (with variance $\sigma 2$) both estimators will produce zero-mean outputs with variances:

$$var(S_{d,k}) = var\left(\frac{V(n+k)-V(n-k)}{2k\,\Delta R}\right) = \frac{2\sigma^2}{4k^2\Delta R^2} = \frac{\sigma^2}{2k^2\Delta R^2} \quad (21)$$

$$var(S_{ls,k}) = \sum_{i=1}^{k} var(S_{d,k})\left[\frac{i^4}{\kappa^2}\right] \quad (22)$$

$$= \frac{\sigma^2}{2\Delta R^2\kappa} \quad (23)$$

The least-squares estimator therefore has a lower variance than the finite difference estimator, with the ratio:

$$R_{\sigma^2} \equiv \frac{var(S_{d,k})}{var(S_{ls,k})} = \frac{(k+1)(2k+1)}{6k} \tag{24}$$

When considering the response to impulse noise, the total time-integrated output power is a reasonable metric for comparing performance. For a unit impulse input, the integrated output powers for the two estimators are:

$$U_{ls,k} = \sum_{i=1}^{k} S_{ls,k}^2(i) = \sum_{i=1}^{k} (\frac{1}{2i})^2 (\frac{i^2}{\kappa})^2 = \sum_{i=1}^{k} \frac{i^2}{4\kappa^2} = \frac{1}{4\kappa} \tag{25}$$

$$U_{d,K} = 2(\frac{1}{2k})^2 = \frac{1}{2k^2} \tag{26}$$

hence

$$R_U \equiv \frac{U_{d,K}}{U_{ls,k}} = \frac{(k+1)(2k+1)}{3k} \tag{27}$$

In this case, the least-squares estimator is again seen to provide better rejection of the impulse noise compared to the finite difference estimator. The relative performance against these two noise sources for the estimators at various window half-sizes is shown in Table 10. For any given window size, the least-squares shear estimator provides significant noise reduction compared to the finite difference estimator for both noise input types.

Table 10: Comparison of Shear Estimator Responses to White and Impulse Noise Inputs

k	R_{σ^2}	R_U
2	1.25	2.50
3	1.56	3.11
4	1.88	3.75
5	2.20	4.40
6	2.17	4.33

APPENDIX 3.
SIMULATION RESULT TABLES FOR FILTER AND SHEAR ALGORITHM TRADEOFF STUDY

This appendix lists the numerical results of several simulation runs of the basic computational shear detection algorithm against each of three idealized microburst shear signatures, as described in Chapter 2. A list is provided for each signature type, indicating the following quantities:

1) Filter algorithm name
2) Filter window size (in range gates)
3) Shear algorithm name
4) Shear window width (in range gates)
5) Test signal standard deviation (always 1.00 m/s)
6) Total POD for test
7) Segment POD
8) Total PFA
9) Raw PFA
10) Average loss for segments overlapping true region

"Easy" sinusoid signature

Total velocity difference = 30 m/s

1	2	3	4	5	6	7	8	9	10
nofilt	0	fd	1	1.00	0.81	1.00	0.00	0.27	27.72
mean	1	fd	1	1.00	0.96	1.00	0.01	0.15	29.49
median	1	fd	1	1.00	0.88	1.00	0.00	0.19	28.49
mean	2	fd	1	1.00	0.98	1.00	0.01	0.07	29.17
median	2	fd	1	1.00	0.89	1.00	0.00	0.12	28.30
mean	3	fd	1	1.00	0.99	1.00	0.01	0.03	28.61
median	3	fd	1	1.00	0.89	1.00	0.00	0.08	27.81
mean	4	fd	1	1.00	0.99	1.00	0.01	0.03	27.92
median	4	fd	1	1.00	0.87	1.00	0.00	0.07	27.30
mean	5	fd	1	1.00	1.00	1.00	0.00	0.03	27.07
median	5	fd	1	1.00	0.85	1.00	0.00	0.05	26.80
nofilt	0	fd	2	1.00	0.95	1.00	0.01	0.17	29.21
mean	1	fd	2	1.00	0.98	1.00	0.01	0.06	29.29
median	1	fd	2	1.00	0.97	1.00	0.01	0.11	29.30
mean	2	fd	2	1.00	0.99	1.00	0.01	0.03	28.87
median	2	fd	2	1.00	0.97	1.00	0.01	0.07	28.89
mean	3	fd	2	1.00	1.00	1.00	0.01	0.01	28.32
median	3	fd	2	1.00	0.96	1.00	0.00	0.03	28.39
mean	4	fd	2	1.00	1.00	1.00	0.00	0.01	27.63
median	4	fd	2	1.00	0.94	1.00	0.00	0.02	27.79
mean	5	fd	2	1.00	1.00	1.00	0.00	0.02	26.79
median	5	fd	2	1.00	0.92	1.00	0.00	0.02	27.16
nofilt	0	fd	3	1.00	0.97	1.00	0.01	0.09	28.91
mean	1	fd	3	1.00	0.99	1.00	0.01	0.02	28.79
median	1	fd	3	1.00	0.98	1.00	0.01	0.04	28.81
mean	2	fd	3	1.00	1.00	1.00	0.01	0.01	28.41
median	2	fd	3	1.00	0.98	1.00	0.01	0.02	28.48
mean	3	fd	3	1.00	1.00	1.00	0.00	0.01	27.89
median	3	fd	3	1.00	0.98	1.00	0.00	0.01	28.09
mean	4	fd	3	1.00	1.00	1.00	0.00	0.01	27.20
median	4	fd	3	1.00	0.97	1.00	0.00	0.01	27.61
mean	5	fd	3	1.00	1.00	1.00	0.00	0.00	26.37
median	5	fd	3	1.00	0.95	1.00	0.00	0.01	27.05
nofilt	0	fd	4	1.00	0.98	1.00	0.01	0.03	28.33
mean	1	fd	4	1.00	1.00	1.00	0.01	0.01	28.16
median	1	fd	4	1.00	0.99	1.00	0.01	0.01	28.20
mean	2	fd	4	1.00	1.00	1.00	0.00	0.00	27.79
median	2	fd	4	1.00	0.99	1.00	0.01	0.01	27.92
mean	3	fd	4	1.00	1.00	1.00	0.00	0.00	27.28
median	3	fd	4	1.00	0.99	1.00	0.00	0.01	27.60
mean	4	fd	4	1.00	1.00	1.00	0.00	0.00	26.61
median	4	fd	4	1.00	0.98	1.00	0.00	0.00	27.23
mean	5	fd	4	1.00	1.00	1.00	0.00	0.00	25.81
median	5	fd	4	1.00	0.98	1.00	0.00	0.00	26.80
nofilt	0	fd	5	1.00	0.99	1.00	0.01	0.01	27.57

mean 1	fd 5	1.00	1.00	1.00	0.00	0.00	27.38
median 1	fd 5	1.00	0.99	1.00	0.01	0.01	27.44
mean 2	fd 5	1.00	1.00	1.00	0.00	0.00	27.03
median 2	fd 5	1.00	0.99	1.00	0.00	0.00	27.20
mean 3	fd 5	1.00	1.00	1.00	0.00	0.00	26.53
median 3	fd 5	1.00	0.99	1.00	0.00	0.00	26.95
mean 4	fd 5	1.00	1.00	1.00	0.00	0.00	25.89
median 4	fd 5	1.00	0.99	1.00	0.00	0.00	26.67
mean 5	fd 5	1.00	1.00	1.00	0.00	0.00	25.12
median 5	fd 5	1.00	0.99	1.00	0.00	0.00	26.34
nofilt 0	lsq 1	1.00	0.81	1.00	0.00	0.27	27.72
mean 1	lsq 1	1.00	0.96	1.00	0.01	0.15	29.49
median 1	lsq 1	1.00	0.88	1.00	0.00	0.19	28.49
mean 2	lsq 1	1.00	0.98	1.00	0.01	0.07	29.17
median 2	lsq 1	1.00	0.89	1.00	0.00	0.12	28.30
mean 3	lsq 1	1.00	0.99	1.00	0.01	0.03	28.61
median 3	lsq 1	1.00	0.89	1.00	0.00	0.08	27.81
mean 4	lsq 1	1.00	0.99	1.00	0.01	0.03	27.92
median 4	lsq 1	1.00	0.87	1.00	0.00	0.07	27.30
mean 5	lsq 1	1.00	1.00	1.00	0.00	0.03	27.07
median 5	lsq 1	1.00	0.85	1.00	0.00	0.05	26.80
nofilt 0	lsq 2	1.00	0.96	1.00	0.01	0.14	29.45
mean 1	lsq 2	1.00	0.98	1.00	0.01	0.07	29.36
median 1	lsq 2	1.00	0.97	1.00	0.01	0.11	29.36
mean 2	lsq 2	1.00	0.99	1.00	0.01	0.03	28.92
median 2	lsq 2	1.00	0.97	1.00	0.01	0.06	28.92
mean 3	lsq 2	1.00	1.00	1.00	0.01	0.01	28.38
median 3	lsq 2	1.00	0.96	1.00	0.00	0.03	28.40
mean 4	lsq 2	1.00	1.00	1.00	0.00	0.01	27.68
median 4	lsq 2	1.00	0.94	1.00	0.00	0.02	27.78
mean 5	lsq 2	1.00	1.00	1.00	0.00	0.01	26.84
median 5	lsq 2	1.00	0.91	1.00	0.00	0.01	27.16
nofilt 0	lsq 3	1.00	0.98	1.00	0.01	0.05	29.18
mean 1	lsq 3	1.00	0.99	1.00	0.01	0.02	28.96
median 1	lsq 3	1.00	0.99	1.00	0.01	0.05	29.01
mean 2	lsq 3	1.00	1.00	1.00	0.01	0.01	28.59
median 2	lsq 3	1.00	0.98	1.00	0.01	0.03	28.64
mean 3	lsq 3	1.00	1.00	1.00	0.00	0.01	28.05
median 3	lsq 3	1.00	0.98	1.00	0.00	0.01	28.19
mean 4	lsq 3	1.00	1.00	1.00	0.00	0.00	27.36
median 4	lsq 3	1.00	0.96	1.00	0.00	0.01	27.65
mean 5	lsq 3	1.00	1.00	1.00	0.00	0.00	26.53
median 5	lsq 3	1.00	0.94	1.00	0.00	0.00	27.06
nofilt 0	lsq 4	1.00	0.99	1.00	0.01	0.01	28.71
mean 1	lsq 4	1.00	1.00	1.00	0.01	0.01	28.53
median 1	lsq 4	1.00	0.99	1.00	0.01	0.01	28.55
mean 2	lsq 4	1.00	1.00	1.00	0.00	0.00	28.16
median 2	lsq 4	1.00	0.99	1.00	0.01	0.01	28.23
mean 3	lsq 4	1.00	1.00	1.00	0.00	0.00	27.64
median 3	lsq 4	1.00	0.99	1.00	0.00	0.01	27.86

mean 4	lsq 4	1.00	1.00	1.00	0.00	0.00	26.96
median 4	lsq 4	1.00	0.98	1.00	0.00	0.00	27.40
mean 5	lsq 4	1.00	1.00	1.00	0.00	0.00	26.14
median 5	lsq 4	1.00	0.97	1.00	0.00	0.00	26.88
nofilt 0	lsq 5	1.00	1.00	1.00	0.00	0.00	28.17
mean 1	lsq 5	1.00	1.00	1.00	0.00	0.00	27.98
median 1	lsq 5	1.00	1.00	1.00	0.00	0.01	28.02
mean 2	lsq 5	1.00	1.00	1.00	0.00	0.00	27.64
median 2	lsq 5	1.00	1.00	1.00	0.00	0.00	27.74
mean 3	lsq 5	1.00	1.00	1.00	0.00	0.00	27.13
median 3	lsq 5	1.00	1.00	1.00	0.00	0.00	27.43
mean 4	lsq 5	1.00	1.00	1.00	0.00	0.00	26.47
median 4	lsq 5	1.00	0.99	1.00	0.00	0.00	27.06
mean 5	lsq 5	1.00	1.00	1.00	0.00	0.00	25.67
median 5	lsq 5	1.00	0.99	1.00	0.00	0.00	26.63

"Weak" sinusoid signature

Total velocity difference = 10 m/s

1	2	3	4	5	6	7	8	9	10
nofilt	0	fd	1	1.00	0.05	0.18	0.00	0.29	5.82
mean	1	fd	1	1.00	0.21	0.50	0.00	0.15	6.47
median	1	fd	1	1.00	0.09	0.26	0.00	0.20	5.97
mean	2	fd	1	1.00	0.41	0.75	0.00	0.07	7.26
median	2	fd	1	1.00	0.12	0.32	0.00	0.12	6.19
mean	3	fd	1	1.00	0.65	0.95	0.00	0.04	8.55
median	3	fd	1	1.00	0.13	0.32	0.00	0.09	6.38
mean	4	fd	1	1.00	0.74	0.99	0.00	0.03	8.84
median	4	fd	1	1.00	0.15	0.33	0.00	0.07	6.45
mean	5	fd	1	1.00	0.80	1.00	0.00	0.03	9.15
median	5	fd	1	1.00	0.14	0.33	0.00	0.06	6.39
nofilt	0	fd	2	1.00	0.13	0.34	0.00	0.17	6.20
mean	1	fd	2	1.00	0.46	0.81	0.00	0.06	7.69
median	1	fd	2	1.00	0.30	0.62	0.00	0.11	7.01
mean	2	fd	2	1.00	0.70	0.96	0.00	0.03	8.98
median	2	fd	2	1.00	0.42	0.75	0.00	0.07	7.60
mean	3	fd	2	1.00	0.83	1.00	0.00	0.02	9.69
median	3	fd	2	1.00	0.48	0.82	0.00	0.04	7.64
mean	4	fd	2	1.00	0.85	1.00	0.00	0.01	9.68
median	4	fd	2	1.00	0.48	0.83	0.00	0.03	7.50
mean	5	fd	2	1.00	0.86	1.00	0.00	0.02	9.54
median	5	fd	2	1.00	0.47	0.82	0.00	0.03	7.39
nofilt	0	fd	3	1.00	0.25	0.52	0.00	0.09	6.68
mean	1	fd	3	1.00	0.73	0.98	0.00	0.02	9.04
median	1	fd	3	1.00	0.53	0.88	0.00	0.05	7.87
mean	2	fd	3	1.00	0.82	1.00	0.00	0.01	9.67
median	2	fd	3	1.00	0.67	0.97	0.00	0.03	8.58
mean	3	fd	3	1.00	0.86	1.00	0.00	0.01	9.85
median	3	fd	3	1.00	0.72	0.98	0.00	0.02	8.81
mean	4	fd	3	1.00	0.87	1.00	0.00	0.01	9.69
median	4	fd	3	1.00	0.73	0.99	0.00	0.01	8.80
mean	5	fd	3	1.00	0.87	1.00	0.00	0.00	9.50
median	5	fd	3	1.00	0.72	0.99	0.00	0.01	8.59
nofilt	0	fd	4	1.00	0.44	0.78	0.00	0.04	7.24
mean	1	fd	4	1.00	0.81	1.00	0.00	0.01	9.49
median	1	fd	4	1.00	0.70	0.97	0.00	0.02	8.78
mean	2	fd	4	1.00	0.86	1.00	0.00	0.00	9.78
median	2	fd	4	1.00	0.79	0.99	0.00	0.01	9.34
mean	3	fd	4	1.00	0.87	1.00	0.00	0.00	9.73
median	3	fd	4	1.00	0.81	1.00	0.00	0.01	9.43
mean	4	fd	4	1.00	0.88	1.00	0.00	0.00	9.58
median	4	fd	4	1.00	0.81	1.00	0.00	0.00	9.27
mean	5	fd	4	1.00	0.87	1.00	0.00	0.00	9.34
median	5	fd	4	1.00	0.80	1.00	0.00	0.00	9.08
nofilt	0	fd	5	1.00	0.57	0.91	0.00	0.01	7.66

mean 1	fd 5	1.00	0.84	1.00	0.00	0.00	9.50
median 1	fd 5	1.00	0.77	1.00	0.00	0.01	9.01
mean 2	fd 5	1.00	0.87	1.00	0.00	0.00	9.65
median 2	fd 5	1.00	0.82	1.00	0.00	0.00	9.33
mean 3	fd 5	1.00	0.88	1.00	0.00	0.00	9.54
median 3	fd 5	1.00	0.84	1.00	0.00	0.00	9.39
mean 4	fd 5	1.00	0.87	1.00	0.00	0.00	9.36
median 4	fd 5	1.00	0.84	1.00	0.00	0.00	9.27
mean 5	fd 5	1.00	0.87	1.00	0.00	0.00	9.15
median 5	fd 5	1.00	0.83	1.00	0.00	0.00	9.09
nofilt 0	lsq 1	1.00	0.05	0.18	0.00	0.29	5.82
mean 1	lsq 1	1.00	0.21	0.50	0.00	0.15	6.47
median 1	lsq 1	1.00	0.09	0.26	0.00	0.20	5.97
mean 2	lsq 1	1.00	0.41	0.75	0.00	0.07	7.26
median 2	lsq 1	1.00	0.12	0.32	0.00	0.12	6.19
mean 3	lsq 1	1.00	0.65	0.95	0.00	0.04	8.55
median 3	lsq 1	1.00	0.13	0.32	0.00	0.09	6.38
mean 4	lsq 1	1.00	0.74	0.99	0.00	0.03	8.84
median 4	lsq 1	1.00	0.15	0.33	0.00	0.07	6.45
mean 5	lsq 1	1.00	0.80	1.00	0.00	0.03	9.15
median 5	lsq 1	1.00	0.14	0.33	0.00	0.06	6.39
nofilt 0	lsq 2	1.00	0.21	0.51	0.00	0.14	6.47
mean 1	lsq 2	1.00	0.42	0.76	0.00	0.07	7.44
median 1	lsq 2	1.00	0.29	0.62	0.00	0.12	6.90
mean 2	lsq 2	1.00	0.72	0.97	0.00	0.03	9.10
median 2	lsq 2	1.00	0.43	0.76	0.00	0.07	7.63
mean 3	lsq 2	1.00	0.83	1.00	0.00	0.01	9.75
median 3	lsq 2	1.00	0.49	0.82	0.00	0.04	7.72
mean 4	lsq 2	1.00	0.86	1.00	0.00	0.01	9.72
median 4	lsq 2	1.00	0.49	0.83	0.00	0.03	7.59
mean 5	lsq 2	1.00	0.86	1.00	0.00	0.01	9.56
median 5	lsq 2	1.00	0.48	0.82	0.00	0.02	7.49
nofilt 0	lsq 3	1.00	0.48	0.82	0.00	0.05	7.68
mean 1	lsq 3	1.00	0.69	0.96	0.00	0.03	8.97
median 1	lsq 3	1.00	0.57	0.90	0.00	0.05	8.17
mean 2	lsq 3	1.00	0.81	1.00	0.00	0.01	9.67
median 2	lsq 3	1.00	0.68	0.97	0.00	0.03	8.71
mean 3	lsq 3	1.00	0.87	1.00	0.00	0.01	9.92
median 3	lsq 3	1.00	0.73	0.99	0.00	0.02	8.94
mean 4	lsq 3	1.00	0.88	1.00	0.00	0.00	9.77
median 4	lsq 3	1.00	0.74	0.99	0.00	0.01	8.87
mean 5	lsq 3	1.00	0.88	1.00	0.00	0.00	9.55
median 5	lsq 3	1.00	0.73	0.99	0.00	0.01	8.65
nofilt 0	lsq 4	1.00	0.73	0.98	0.00	0.02	9.07
mean 1	lsq 4	1.00	0.83	1.00	0.00	0.01	9.79
median 1	lsq 4	1.00	0.76	0.99	0.00	0.02	9.34
mean 2	lsq 4	1.00	0.86	1.00	0.00	0.01	9.92
median 2	lsq 4	1.00	0.80	1.00	0.00	0.01	9.54
mean 3	lsq 4	1.00	0.88	1.00	0.00	0.00	9.86
median 3	lsq 4	1.00	0.82	1.00	0.00	0.01	9.57

mean 4	lsq 4	1.00	0.88	1.00	0.00	0.00	9.67
median 4	lsq 4	1.00	0.82	1.00	0.00	0.01	9.42
mean 5	lsq 4	1.00	0.87	1.00	0.00	0.00	9.43
median 5	lsq 4	1.00	0.81	1.00	0.00	0.00	9.21
nofilt 0	lsq 5	1.00	0.83	1.00	0.00	0.01	9.65
mean 1	lsq 5	1.00	0.87	1.00	0.00	0.00	9.89
median 1	lsq 5	1.00	0.83	1.00	0.00	0.01	9.71
mean 2	lsq 5	1.00	0.87	1.00	0.00	0.00	9.85
median 2	lsq 5	1.00	0.85	1.00	0.00	0.01	9.74
mean 3	lsq 5	1.00	0.88	1.00	0.00	0.00	9.71
median 3	lsq 5	1.00	0.85	1.00	0.00	0.00	9.67
mean 4	lsq 5	1.00	0.88	1.00	0.00	0.00	9.53
median 4	lsq 5	1.00	0.85	1.00	0.00	0.00	9.49
mean 5	lsq 5	1.00	0.87	1.00	0.00	0.00	9.31
median 5	lsq 5	1.00	0.83	1.00	0.00	0.00	9.28

"Strong" model signature

Total velocity difference = 14 m/s

1	2	3	4	5	6	7	8	9	10
nofilt	0	fd	1	1.00	0.77	1.00	0.00	0.31	11.93
mean	1	fd	1	1.00	0.75	1.00	0.00	0.17	9.98
median	1	fd	1	1.00	0.76	1.00	0.00	0.21	9.87
mean	2	fd	1	1.00	0.80	1.00	0.00	0.06	7.94
median	2	fd	1	1.00	0.66	1.00	0.00	0.13	7.60
mean	3	fd	1	1.00	0.91	1.00	0.00	0.03	6.41
median	3	fd	1	1.00	0.47	0.80	0.00	0.09	6.12
mean	4	fd	1	1.00	0.69	0.69	0.00	0.02	5.51
median	4	fd	1	1.00	0.16	0.28	0.00	0.07	5.56
mean	5	fd	1	1.00	0.10	0.10	0.00	0.04	5.24
median	5	fd	1	1.00	0.02	0.03	0.00	0.06	5.28
nofilt	0	fd	2	1.00	0.78	1.00	0.00	0.19	9.20
mean	1	fd	2	1.00	0.76	1.00	0.00	0.06	8.48
median	1	fd	2	1.00	0.79	1.00	0.00	0.12	8.25
mean	2	fd	2	1.00	0.81	1.00	0.00	0.02	7.25
median	2	fd	2	1.00	0.79	1.00	0.00	0.07	7.10
mean	3	fd	2	1.00	0.94	0.99	0.00	0.01	6.06
median	3	fd	2	1.00	0.63	0.79	0.00	0.03	6.07
mean	4	fd	2	1.00	0.42	0.42	0.00	0.01	5.36
median	4	fd	2	1.00	0.19	0.25	0.00	0.02	5.59
mean	5	fd	2	1.00	0.02	0.02	0.00	0.03	5.06
median	5	fd	2	1.00	0.02	0.02	0.00	0.02	5.29
nofilt	0	fd	3	1.00	0.85	1.00	0.00	0.09	7.29
mean	1	fd	3	1.00	0.86	1.00	0.00	0.01	6.89
median	1	fd	3	1.00	0.86	1.00	0.00	0.04	6.70
mean	2	fd	3	1.00	0.90	1.00	0.00	0.00	6.30
median	2	fd	3	1.00	0.81	0.93	0.00	0.02	6.10
mean	3	fd	3	1.00	0.85	0.86	0.00	0.00	5.57
median	3	fd	3	1.00	0.51	0.57	0.00	0.01	5.65
mean	4	fd	3	1.00	0.11	0.11	0.00	0.01	5.21
median	4	fd	3	1.00	0.15	0.16	0.00	0.01	5.48
mean	5	fd	3	1.00	0.00	0.00	0.00	0.01	0.00
median	5	fd	3	1.00	0.01	0.01	0.00	0.01	5.28
nofilt	0	fd	4	1.00	0.89	0.94	0.00	0.03	5.96
mean	1	fd	4	1.00	0.91	0.91	0.00	0.00	5.74
median	1	fd	4	1.00	0.75	0.78	0.00	0.01	5.73
mean	2	fd	4	1.00	0.69	0.69	0.00	0.00	5.43
median	2	fd	4	1.00	0.49	0.50	0.00	0.00	5.53
mean	3	fd	4	1.00	0.12	0.12	0.00	0.00	5.21
median	3	fd	4	1.00	0.19	0.19	0.00	0.00	5.34
mean	4	fd	4	1.00	0.00	0.00	0.00	0.00	0.00
median	4	fd	4	1.00	0.05	0.05	0.00	0.00	5.28
mean	5	fd	4	1.00	0.00	0.00	0.00	0.00	0.00
median	5	fd	4	1.00	0.00	0.00	0.00	0.00	5.21
nofilt	0	fd	5	1.00	0.34	0.34	0.00	0.01	5.40

mean 1	fd 5	1.00	0.24	0.24	0.00	0.01	5.31
median 1	fd 5	1.00	0.20	0.20	0.00	0.01	5.34
mean 2	fd 5	1.00	0.09	0.09	0.00	0.01	5.20
median 2	fd 5	1.00	0.09	0.09	0.00	0.01	5.25
mean 3	fd 5	1.00	0.00	0.00	0.00	0.00	0.00
median 3	fd 5	1.00	0.01	0.01	0.00	0.01	5.14
mean 4	fd 5	1.00	0.00	0.00	0.00	0.00	0.00
median 4	fd 5	1.00	0.00	0.00	0.00	0.01	5.01
mean 5	fd 5	1.00	0.00	0.00	0.00	0.00	0.00
median 5	fd 5	1.00	0.00	0.00	0.00	0.00	0.00
nofilt 0	lsq 1	1.00	0.77	1.00	0.00	0.31	11.93
mean 1	lsq 1	1.00	0.75	1.00	0.00	0.17	9.98
median 1	lsq 1	1.00	0.76	1.00	0.00	0.21	9.87
mean 2	lsq 1	1.00	0.80	1.00	0.00	0.06	7.94
median 2	lsq 1	1.00	0.66	1.00	0.00	0.13	7.60
mean 3	lsq 1	1.00	0.91	1.00	0.00	0.03	6.41
median 3	lsq 1	1.00	0.47	0.80	0.00	0.09	6.12
mean 4	lsq 1	1.00	0.69	0.69	0.00	0.02	5.51
median 4	lsq 1	1.00	0.16	0.28	0.00	0.07	5.56
mean 5	lsq 1	1.00	0.10	0.10	0.00	0.04	5.24
median 5	lsq 1	1.00	0.02	0.03	0.00	0.06	5.28
nofilt 0	lsq 2	1.00	0.76	1.00	0.00	0.15	9.65
mean 1	lsq 2	1.00	0.75	1.00	0.00	0.07	8.77
median 1	lsq 2	1.00	0.76	1.00	0.00	0.13	8.47
mean 2	lsq 2	1.00	0.80	1.00	0.00	0.02	7.38
median 2	lsq 2	1.00	0.77	1.00	0.00	0.07	7.18
mean 3	lsq 2	1.00	0.94	0.99	0.00	0.00	6.13
median 3	lsq 2	1.00	0.62	0.80	0.00	0.03	6.05
mean 4	lsq 2	1.00	0.45	0.45	0.00	0.01	5.39
median 4	lsq 2	1.00	0.18	0.24	0.00	0.02	5.57
mean 5	lsq 2	1.00	0.03	0.03	0.00	0.02	5.10
median 5	lsq 2	1.00	0.01	0.02	0.00	0.01	5.25
nofilt 0	lsq 3	1.00	0.79	1.00	0.00	0.05	8.00
mean 1	lsq 3	1.00	0.80	1.00	0.00	0.02	7.49
median 1	lsq 3	1.00	0.80	1.00	0.00	0.04	7.22
mean 2	lsq 3	1.00	0.86	1.00	0.00	0.01	6.65
median 2	lsq 3	1.00	0.80	0.97	0.00	0.03	6.39
mean 3	lsq 3	1.00	0.93	0.94	0.00	0.00	5.77
median 3	lsq 3	1.00	0.55	0.65	0.00	0.01	5.72
mean 4	lsq 3	1.00	0.18	0.18	0.00	0.00	5.30
median 4	lsq 3	1.00	0.15	0.18	0.00	0.00	5.48
mean 5	lsq 3	1.00	0.00	0.00	0.00	0.01	0.00
median 5	lsq 3	1.00	0.01	0.01	0.00	0.00	5.24
nofilt 0	lsq 4	1.00	0.87	1.00	0.00	0.01	6.74
mean 1	lsq 4	1.00	0.90	1.00	0.00	0.00	6.44
median 1	lsq 4	1.00	0.86	0.97	0.00	0.01	6.23
mean 2	lsq 4	1.00	0.94	0.98	0.00	0.00	5.93
median 2	lsq 4	1.00	0.73	0.79	0.00	0.00	5.80
mean 3	lsq 4	1.00	0.57	0.57	0.00	0.00	5.36
median 3	lsq 4	1.00	0.36	0.37	0.00	0.00	5.46

mean 4	lsq 4	1.00	0.02	0.02	0.00	0.00	5.08
median 4	lsq 4	1.00	0.09	0.09	0.00	0.00	5.33
mean 5	lsq 4	1.00	0.00	0.00	0.00	0.01	0.00
median 5	lsq 4	1.00	0.01	0.01	0.00	0.00	5.17
nofilt 0	lsq 5	1.00	0.93	0.94	0.00	0.00	5.80
mean 1	lsq 5	1.00	0.89	0.89	0.00	0.00	5.61
median 1	lsq 5	1.00	0.70	0.71	0.00	0.00	5.57
mean 2	lsq 5	1.00	0.52	0.52	0.00	0.00	5.35
median 2	lsq 5	1.00	0.35	0.35	0.00	0.00	5.39
mean 3	lsq 5	1.00	0.07	0.07	0.00	0.00	5.15
median 3	lsq 5	1.00	0.06	0.06	0.00	0.00	5.26
mean 4	lsq 5	1.00	0.00	0.00	0.00	0.00	0.00
median 4	lsq 5	1.00	0.01	0.01	0.00	0.00	5.12
mean 5	lsq 5	1.00	0.00	0.00	0.00	0.00	0.00
median 5	lsq 5	1.00	0.00	0.00	0.00	0.00	5.04