

**Project Report
ATC-196**

Machine Intelligent Gust Front Algorithm

**R. L. Delanoy
S. W. Troxel**

4 November 1993

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

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.

1. Report No. ATC-196		2. Government Accession No. DOT/FAA/RD-93/1		3. Recipient's Catalog No.	
4. Title and Subtitle Machine Intelligent Gust Front Algorithm				5. Report Date 4 November 1993	
				6. Performing Organization Code	
7. Author(s) Richard L. Delanoy, Seth W. Troxel				8. Performing Organization Report No. ATC-196	
9. Performing Organization Name and Address Lincoln Laboratory, MIT P.O. Box 73 Lexington, MA 02173-9108				10. Work Unit No. (TRAIS)	
				11. Contract or Grant No. DTFA-01-93-Z-02012	
12. Sponsoring Agency Name and Address Department of Transportation Federal Aviation Administration Systems Research and Development Service Washington, DC 20591				13. Type of Report and Period Covered 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 Air Force under Contract F19628-90-C-0002.					
16. Abstract The Federal Aviation Administration has sponsored research and development of algorithms for automatic gust front detection as part of a suite of hazardous weather detection capabilities for airports. These algorithms are intended for use with Doppler radar systems, specifically the Terminal Doppler Weather Radar (TDWR) and the Airport Surveillance Radar enhanced with a Wind Shear Processor (ASR-9 WSP). Although gust fronts are observable with fairly reliable signatures in TDWR data, existing gust front detection algorithms have achieved only modest levels of detection performance. For smaller airports not slated to receive a dedicated TDWR, the ASR-9 WSP will provide a less expensive wind shear detection capability. Gust front detection in ASR-9 SP data is an even more difficult problem, given the reduced sensitivity and less reliable Doppler measurements of this radar. A Machine Intelligent Gust Front Algorithm (MIGFA) has been constructed at Lincoln Laboratory that is a radical departure from previous design strategies. Incorporating knowledge-based, signal-processing techniques initially developed at Lincoln Laboratory for automatic target recognition, MIGFA uses meteorological knowledge, spatial and temporal context, conditional data fusion, delayed thresholding, and pixel-level fusion of evidence to improve gust front detection performance significantly. In tests comparing MIGFA with an existing state-of-the-art algorithm applied to ASR-9 WSP data, MIGFA has substantially outperformed the older algorithm. In fact, by some measures, MIGFA has done as well or better than human interpreters of the same data. Operational testing of this version was done during 1992 in Orlando, Florida. The design, test results, and performance evaluation of the ASR-9 WSP version of MIGFA are presented in this report, which was prepared as part of the documentation package for the ASR-9 WSP gust front algorithm.					
17. Key Words gust front detection MIGFA wind shear detection ASR-9 functional template correlation			18. Distribution Statement This document is available to the public through the National Technical Information Service, Springfield, VA 22161.		
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of Pages 129	22. Price

ABSTRACT

The Federal Aviation Administration (FAA) has sponsored research and development of algorithms for automatic gust front detection as part of a suite of hazardous weather detection capabilities for airports. These algorithms are intended for use with Doppler radar systems, specifically the Terminal Doppler Weather Radar (TDWR) and the Airport Surveillance Radar, enhanced with a Wind Shear Processor (ASR-9 WSP). Although gust fronts are observable with fairly reliable signatures in TDWR data, existing gust front detection algorithms have achieved only modest levels of detection performance. For smaller airports not slated to receive a dedicated TDWR, the ASR-9 WSP will provide a less expensive wind shear detection capability. Gust front detection in ASR-9 WSP data is an even more difficult problem, given the reduced sensitivity and less reliable Doppler measurements of this radar.

A Machine Intelligent Gust Front Algorithm (MIGFA) has been constructed at Lincoln Laboratory that is a radical departure from previous design strategies. Incorporating knowledge-based signal processing techniques initially developed at Lincoln Laboratory for automatic target recognition, MIGFA uses meteorological knowledge, spatial and temporal context, conditional data fusion, delayed thresholding, and pixel-level fusion of evidence to significantly improve gust front detection performance.

In tests comparing MIGFA with an existing state of the art algorithm applied to ASR-9 WSP data, MIGFA has substantially outperformed the older algorithm. In fact, by some measures MIGFA has done as well or better than human interpreters of the same data. Operational testing of this version was done during 1992 in Orlando, Florida.

The design, test results, and performance evaluation of the ASR-9 WSP version of MIGFA are presented in this report, which was prepared as part of the documentation package for the ASR-9 WSP gust front algorithm.

ACKNOWLEDGMENTS

The authors would like to thank the staff of the FL-3 ASR-9 radar site in Orlando, FL, who were responsible for running MIGFA, recording data, and logging results during the operational test period in 1992. These people include Wes Johnston, Craig McFarland, Jeff Boisseau, and Cindy Meuse. The authors would also like to thank Joe Cullen for generating “truth” from ASR-9 WSP data, against which all algorithm results have been scored. Thanks to Mike Donovan for generating Kansas City statistics on gust front propagation speeds. Thanks to Leslie Mahn, who assisted in editing and proof-reading this report.

The development environment SKETCH was first created by Robert Walton while on the staff of MIT Lincoln Laboratory. Other contributors include Jacques Verly, Patrick van Hove, and Carol Lazott. XTRS, the prototype object recognition system upon which the Machine Intelligent Gust Front Algorithm is based, has evolved over several years of collaboration with Jacques Verly and Dan Dudgeon.

TABLE OF CONTENTS

Abstract	iii
Acknowledgments	v
List of Illustrations	ix
List of Tables	xi
1. INTRODUCTION	1
2. PROBLEM DESCRIPTION	3
2.1 Overview of the ASR-9 WSP Testbed System	3
2.2 ASR-9 WSP Base Data	7
2.3 Gust Front Signatures	7
2.4 Required Algorithm Products	16
3. PRIOR RESEARCH AND DEVELOPMENT	21
4. ALGORITHM DESCRIPTION	25
4.1 Design Philosophies	25
4.2 Knowledge-based Signal Processing	26
4.3 Overview	28
4.4 Image Preparation	33
4.5 Feature Detection	33
4.6 Combining Evidence	52
4.7 Extraction	57
4.8 Tracking/Heuristics	63
4.9 Prediction	65
4.10 Computing Wind Shift and Wind Shear Hazard	66
4.11 Gust Front Update Task	67
5. RESULTS	69
5.1 Truthing and Scoring	69
5.2 Comparison of AGFA and MIGFA	70
5.3 Results of 1992 Field Testing	75
5.4 Accuracy of Wind Estimates	78
5.5 FAA Survey of Air Traffic Controllers	81

TABLE OF CONTENTS
(Continued)

6. EVALUATION	85
6.1 MIGFA Versus Human Performance	85
6.2 Failure Modes	86
6.3 Site Dependent Knowledge	109
6.4 Resource Requirements	110
7. CONCLUSIONS	111
GLOSSARY	113
REFERENCES	115

LIST OF ILLUSTRATIONS

Figure No.		Page
1	Block diagram of the ASR-9 WSP testbed processing and recording system	5
2	High-level block diagram of the ASR-9 WSP gust front detection system	6
3	Comparison of reflectivity and Doppler velocity gust front signatures seen by TDWR and ASR-9 WSP	9
4	Example of sharp decrease in radial velocity indicating convergence associated with a gust front	11
5	Minimum detectable signal expressed in weather reflectivity units (dBZ) as a function of range for the low beam of the ASR-9 WSP testbed.	13
6	Distribution of gust front propagation speeds for Kansas City and Orlando	15
7	Example GSD. The current position of a gust front (solid purple curve) is shown along with 10 and 20 minute forecast gust front positions (dashed purple curves)	17
8	Principal steps of the GFDA algorithm	22
9	Thin line thresholding procedure applied along a radial of reflectivity data	23
10	Example functional template for feature detector TL-DZ	27
11	MIGFA block diagram	29
12	Processed scan summary	31
13	Standard deviation texture map	35
14	DZ thin line detector	35
15	Functional template for feature detector TL-SD	37
16	SD thin line detector	39
17	Functional template for feature detector DZ-MOTION	41
18	Functional template for detecting patches of precipitation in reflectivity images	42
19	Feature detector DZ-MOTION	45
20	Feature detector SD-MOTION	45
21	Dual functional template for out-of-trip weather	47
22	Feature detector for out-of-trip weather	49

LIST OF ILLUSTRATIONS (Continued)

Figure No.		Page
23	Combining interest: strong evidence	53
24	Combining interest: weak evidence	55
25	Extraction steps	59
26	Bow-tie filter for thin line smoothing	62
27	Example extraction of most interesting chains from a simple graph structure	64
28	ASR-9 WSP gust front wind shear hazard (ΔV_h) estimation	68
29	Example of MIGFA gust front detection extending beyond region identified by a human interpreter	73
30	Example radial velocity profiles across three fronts having different cross-front gradients but the same end-to-end velocity change	82
31	Comparison of gust front visibility in matching TDWR and ASR-9 WSP images	87
32	ASR-9 WSP reflectivity image sequence illustrating contrast between nearly invisible inbound gust front signatures and more distinct outbound gust front signatures	89
33	Example of a thin line echo associated with stratiform rain	93
34	Example of a high altitude thin line echo not associated with a gust front	95
35	Example of sidelobe contamination in ASR-9 WSP reflectivity and velocity images	99
36	Example of gust front obscuration due to clutter breakthrough from anomalous propagation (AP)	103
37	Obscuration of a gust front by range ambiguous (out-of-trip) weather echoes	105
38	ASR-9 WSP reflectivity images illustrating obscuration of a gust front by storm cells	107

LIST OF TABLES

Table No.		Page
1	Parameters for the ASR-9	4
2	AGFA and MIGFA performance on ASR-9 WSP data as scored against human interpretations	71
3	AGFA and MIGFA performance on ASR-9 WSP data as scored against human interpretations of matching TDWR data	75
4	MIGFA results on ASR-9 WSP data for the period 1 August to 20 September, 1992 in Orlando, Florida	77
5	Observed vs. Forecast Gust Front Wind Shift	79
6	Observed (LLWAS) vs. Reported (ASR-9 WSP) Gust Front ΔV_h	81

1. INTRODUCTION

A gust front is the leading edge of a cold air outflow from a thunderstorm. The outflow, which is deflected at the ground, may propagate many miles ahead of the generating thunderstorm and may persist as an outflow boundary long after the original storm has dissipated. Although gust fronts are generally deemed less dangerous than microbursts, the associated turbulence, crosswinds, and convergent wind shear can produce conditions that are hazardous to landing aircraft. In addition, the change in wind speed and direction behind a gust front may require alteration of runway usage; advance notice of gust front arrivals is desirable since it allows Air Traffic Control (ATC) to more efficiently organize traffic flow in anticipation of a runway reconfiguration. Reliable detection and forecasting of gust fronts would improve both air safety and reduce costly delays. The Federal Aviation Administration (FAA) has sponsored research and development of an automated gust front detection algorithm to be included as a critical component in a suite of hazardous weather detection capabilities for the Airport Surveillance Radar with Wind Shear Processor (ASR-9 WSP) system and the Terminal Doppler Weather Radar (TDWR).

Gust fronts are observable in Doppler radar images as boundaries between regions of converging Doppler velocity values and as thin lines of increased reflectivity caused by the concentration of particulate scatterers in the convergent boundary. Existing algorithms work reasonably well in TDWR imagery using these signatures. However, many fronts or parts of fronts are missed due to the ambiguous or conditional nature of the observable signatures. For example, convergence signatures disappear as gust fronts become radially aligned; reflectivity thin lines can be obscured by low altitude precipitation.

Detecting gust fronts in ASR-9 WSP data is an even harder problem. With the reduced sensitivity inherent in the fan-beam design of the ASR-9, clear air velocity estimates are unreliable, thereby eliminating most convergence signatures. The reduced sensitivity also makes the reflectivity thin lines more fragmented and harder to resolve from background. Nevertheless, human observers can reliably detect gust fronts in ASR-9 WSP data, arguing that with sufficient perceptual skills, automated gust front detection should also be possible.

This report describes a Machine Intelligent Gust Front Algorithm (MIGFA), which uses knowledge-based signal processing and object recognition techniques originally developed at Lincoln Laboratory for the automatic detection and recognition of military vehicles in radar imagery. This approach uses contextual knowledge, spatial and temporal context, conditional data fusion, delayed thresholding, and pixel-level maps of evidence/belief to improve gust front detection performance. Although developed for both ASR-9 WSP and TDWR radar data, this report focuses on the ASR-9 WSP version of MIGFA.

The next chapter is a detailed description of the problem domain, including a description of the radar, the radar products serving as algorithm inputs, and the required algorithm outputs. The third chapter is a review of the existing gust front detection algorithms and a discussion of their limitations. This is followed by a chapter describing details of the MIGFA design and

implementation. The fifth and sixth chapters present test results and an evaluation of MIGFA performance, respectively.

2. PROBLEM DESCRIPTION

2.1 Overview of the ASR-9 WSP Testbed System

The ASR-9 uses a broad elevation fan-shaped beam to detect aircraft over altitudes up to 35,000 feet within a 60 nmi radius of the airport. Although primarily designed for detecting and tracking aircraft, many characteristics (e.g., transmit frequency, pulse repetition frequency (PRF), pulse width) of the ASR-9 make it suitable for weather detection. The ASR-9 possesses a separate dedicated weather processing channel that provides air traffic controllers with quantitative reports of precipitation intensity. Radar parameters for the ASR-9 are given in Table 1. Since 1986, Lincoln Laboratory has been developing a low-cost Wind Shear Processor (WSP) add-on to the ASR-9 that will provide wind shear detection capabilities at smaller airports not slated to receive a dedicated TDWR. Descriptions of the Lincoln Laboratory ASR-9 WSP testbed can be found in [1, 2, 3].

The Lincoln Laboratory ASR-9 WSP testbed is a production ASR-9 that has been modified to allow simultaneous reception of signals from the low and high receiving beams. Figure 1 is a high level block diagram of the current Lincoln Laboratory ASR-9 WSP testbed (also known as FL-3). Signals from the high and low receiving beams of the ASR-9 are shunted to receivers and digitizers built at Lincoln Laboratory. The digitized I/Q outputs of the A/D converters are distributed simultaneously to a high density 28-track data recorder and to the signal processor. The signal processor uses a series of array processing cards interconnected via a VME-bus backplane to perform adaptive clutter filtering, reflectivity measurement, radial velocity estimation, and data quality assessment. The reflectivity and velocity estimates are sent to separate processors which host the various algorithms used to detect microbursts and gust fronts as well as provide estimates of storm motion. Anemometer data, measuring winds near the airport, are provided by the Low-Level Wind Shear Alert System (LLWAS). Outputs from these algorithms are sent to a local Geographical Situation Display (GSD) and passed along to remote GSDs located in air traffic control facilities. Figure 2 illustrates the end-to-end data flow through the real-time gust front detection system, of which MIGFA is a part.

The ASR-9 WSP signal processor provides reflectivity and velocity data to the gust front algorithm at prescribed intervals (currently 2 minutes). LLWAS anemometer data are received at 10 second intervals. Detection and forecast data produced by the gust front algorithm are passed to a separate GSD update timer task, which sends results to the GSD device at the desired display update rate (nominally 1 minute). The GSD update task uses MIGFA-generated forecasts to update the current location of the gust front on the GSD between actual algorithm updates.

The MIGFA algorithm (including wind shift and wind shear estimations) and the GSD update task are all currently installed on a single SUN Microsystems Sparc 2 workstation that operates at approximately 28 MIPS and is equipped with 32 Mbytes of memory.

TABLE 1**Parameters for the ASR-9. (From Weber et al., 1989 [1]).**

<i>Transmitter</i>	
Frequency	2.7-2.9 GHz
Polarization	Linear or Circular
Peak Power	1.1 MW
Pulse Width	1.0 μ s
Block-Staggered CPI Lengths	8 pulses/10 pulses
PRFs (Example)	972 s ⁻¹ /1250 s ⁻¹
<i>Receiver</i>	
Noise Figure	4.1 dB (max)
Sensitivity	-108 dBm
A/D Word Size	12 bit
<i>Antenna</i>	
Elevation Beamwidth	4.8° (min)
Azimuth Beamwidth	1.4°
Power Gain	34 dB
Rotation Rate	12.5 RPM

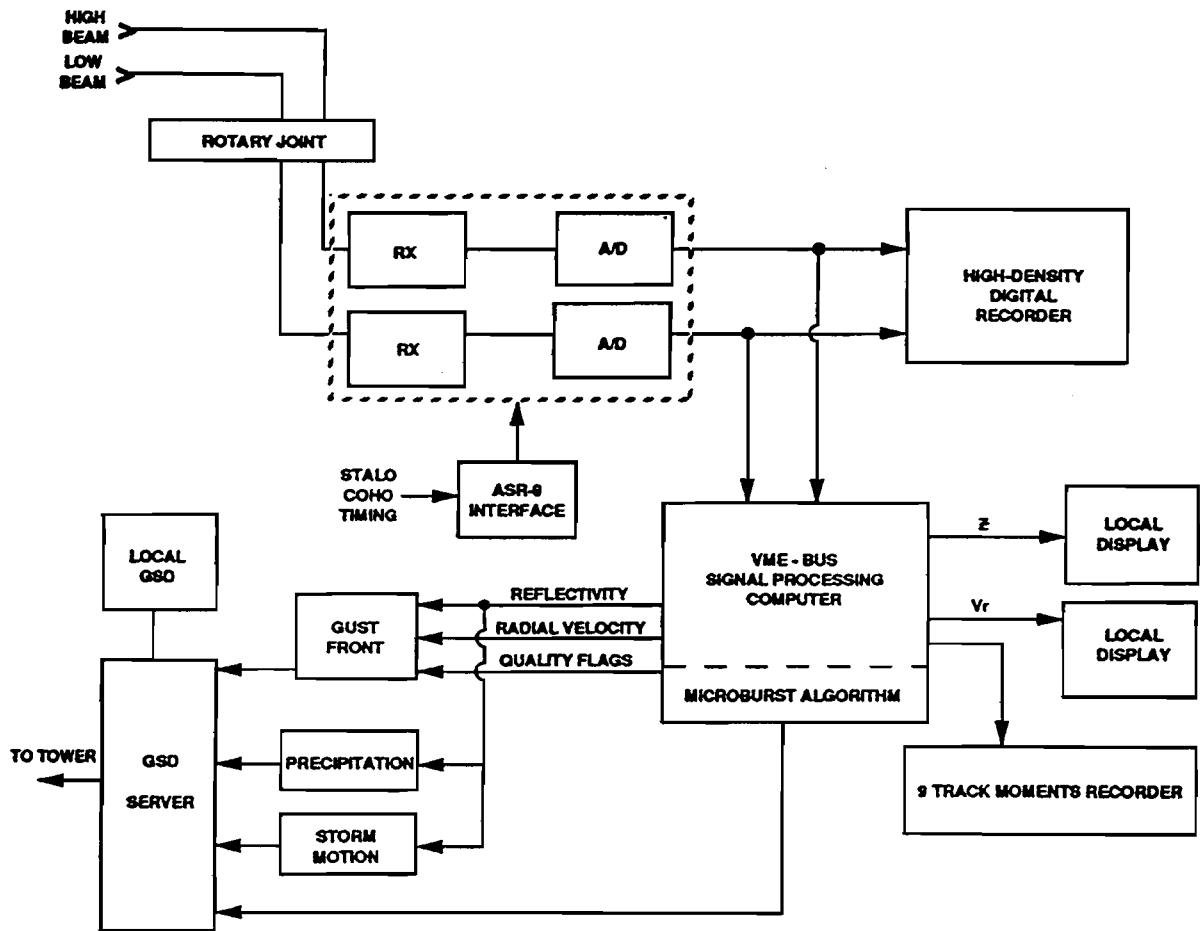


Figure 1. Block diagram of the ASR-9 WSP testbed processing and recording system.

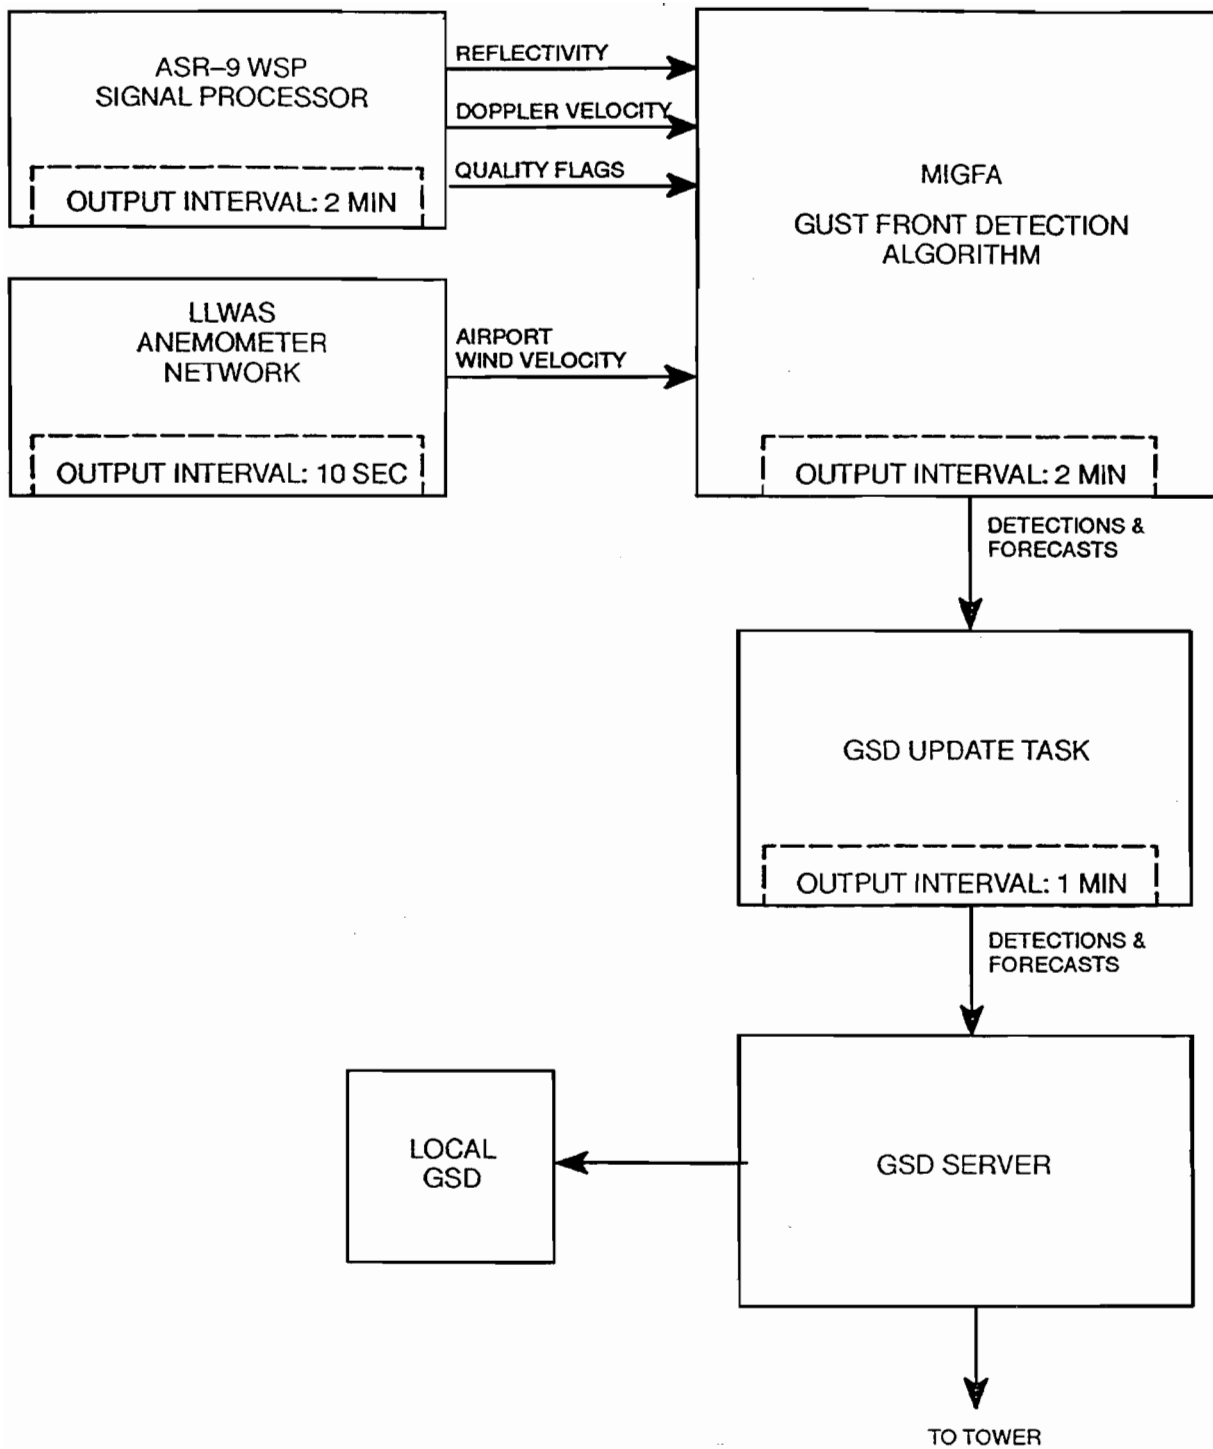


Figure 2. High-level block diagram of the ASR-9 WSP gust front detection system.

2.2 ASR-9 WSP Base Data

MIGFA receives reflectivity and velocity data in polar coordinate form. The data have resolution of 1.4 degrees in azimuth and 115.75 meters in range, extending out to 27.8 km range. Thus, each type of input data is comprised of 256 radials of 240 gates each.

The reflectivity base data ingested by the gust front algorithm are a composite of high and low beam data. Reflectivity estimates from the higher elevation beam are used at close range from 0 to 6 km to minimize ground clutter obscuration of gust front signatures, while reflectivity estimates derived from the low receiving beam are used from 6 to 27.8 km. The reflectivity estimates are smoothed temporally over a 30 second interval by averaging clutter-filtered estimates from the current and previous five scans. Averaging significantly reduces the noise variance, thereby eliminating the need to apply SNR thresholding when extracting the weather signal and increasing the probability of detection for relatively weak gust front signals.

Velocity estimates from the low receiving beam are the source of near-surface velocity data for the gust front algorithm. The velocity estimates are unthresholded with respect to signal-to-noise ratio and are currently calculated using a standard pulse-pair estimation technique applied to clutter-filtered samples obtained on a single scan. Alternative methods for producing more reliable velocity estimates in regions of low SNR are currently being investigated [4].

2.3 Gust Front Signatures

Gust fronts are distinguished in Doppler radar imagery on the basis of three physical properties:

1. Velocity convergence,
2. Thin lines, and
3. Motion.

Although many gust fronts are observable in ASR-9 WSP data, the relevant gust front signatures are often less distinguishable than those observed in TDWR data. Figure 3 contrasts the appearance of a typical, unobscured gust front as viewed by the TDWR and ASR-9 WSP systems and will be used to illustrate the following discussions of gust front signatures. The front had passed over the Orlando airport on October 4, 1992, approximately 10 minutes prior to the time shown in the figure and can be seen extending from northeast to southwest, passing through a point just south of the 3-mile extension of runway 17/35. An older and weaker surge can be seen in the TDWR data (upper two panes of the figure) located approximately 14 km southeast of the TDWR.

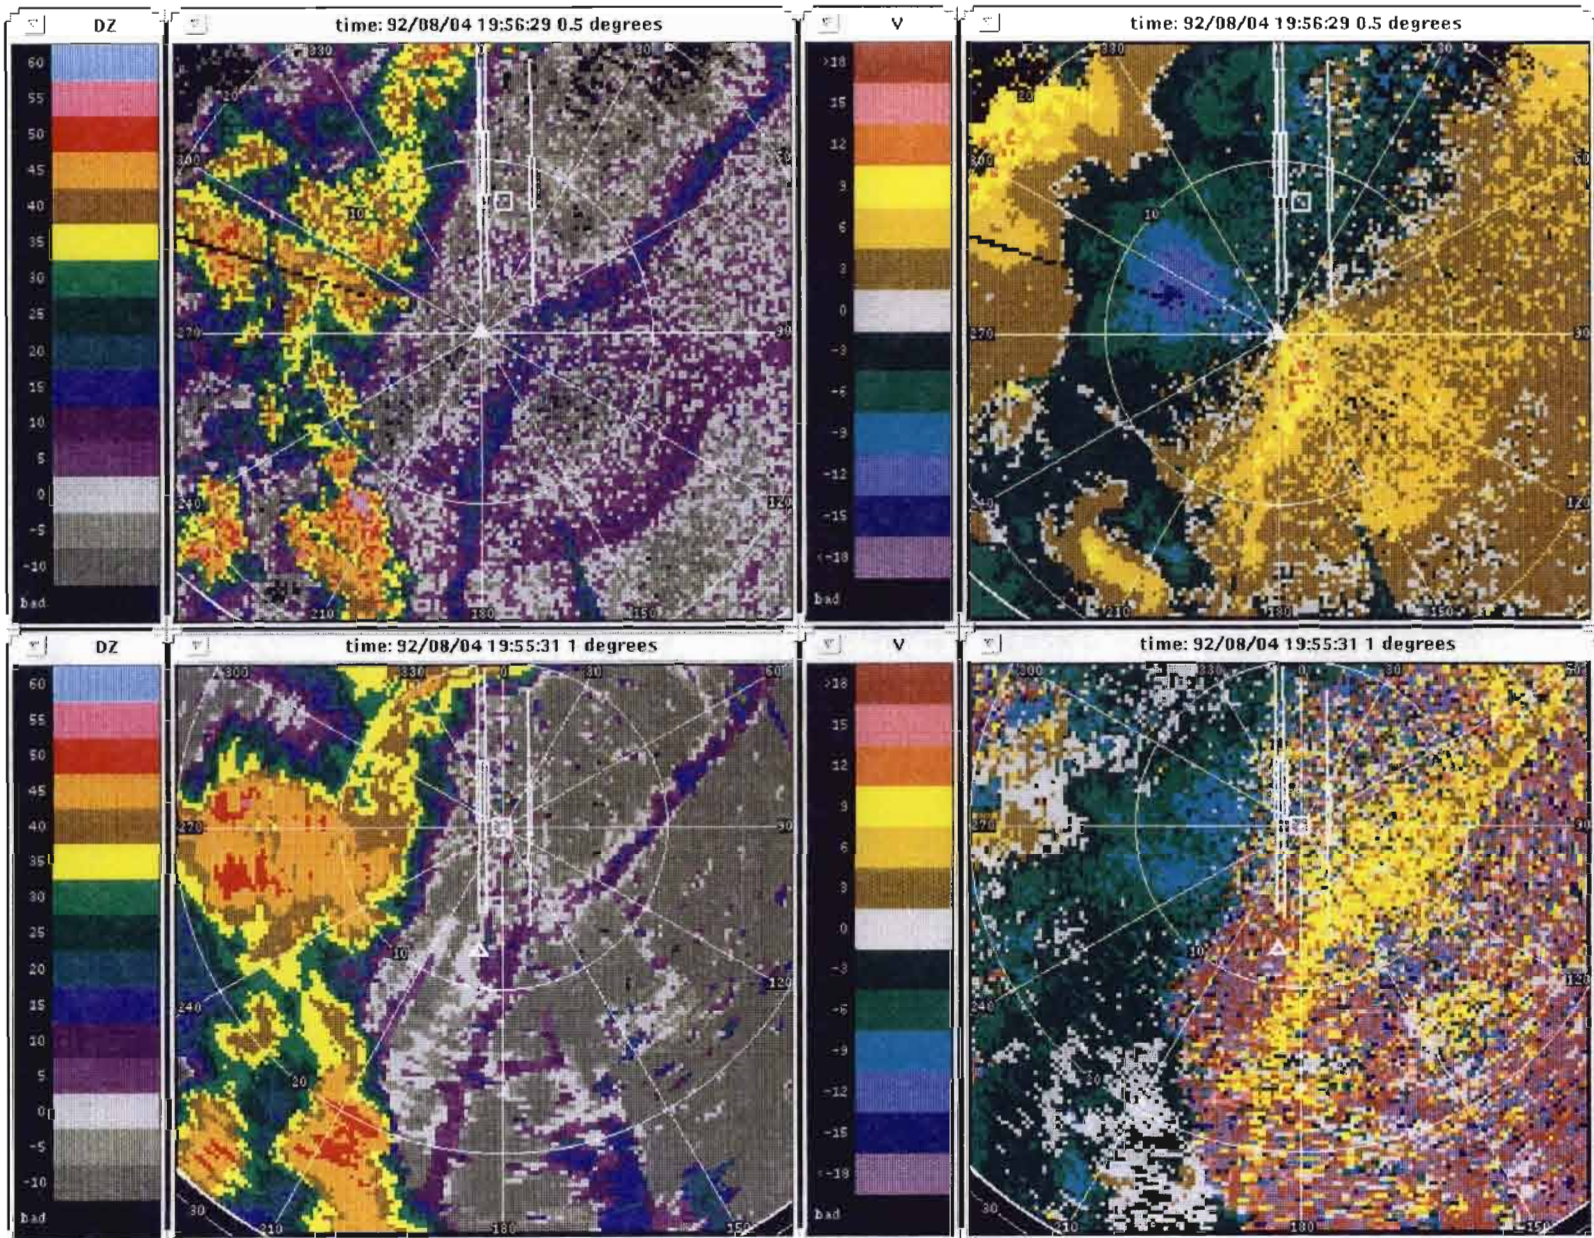


Figure 9. Comparison of reflectivity (left) and Doppler velocity (right) gust front signatures seen by TDWR (top) and ASR-9 WSP (bottom). Range rings are in km, reflectivity (DZ) in dBZ, and radial velocity (V) in m/s.

2.3.1 Convergence

The velocity of the wind within and behind the gust front differs from that of the air ahead of the gust front, creating a convergent velocity shear zone across the front. The TDWR velocity image shown in the upper right of Figure 3 demonstrates this problem. A convergent shear signature (seen as an enhancement in the receding velocities 3 km southeast of the TDWR) exists along the middle portion of the front where orientation is normal to radar beam. However, the convergence signature diminishes considerably as the front becomes more radially oriented along its southern end, and vanishes altogether along the northern end. When viewed along a single radial, the convergence signature is characterized by a relatively sharp decrease in Doppler radial velocity with distance (Figure 4). Since Doppler radars can only measure the component of the wind directed along the beam, velocity estimates based on Doppler measurements can often underestimate the true wind speed. In the extreme, the convergence signature of a gust front disappears completely when the direction of motion is across the radar beam.

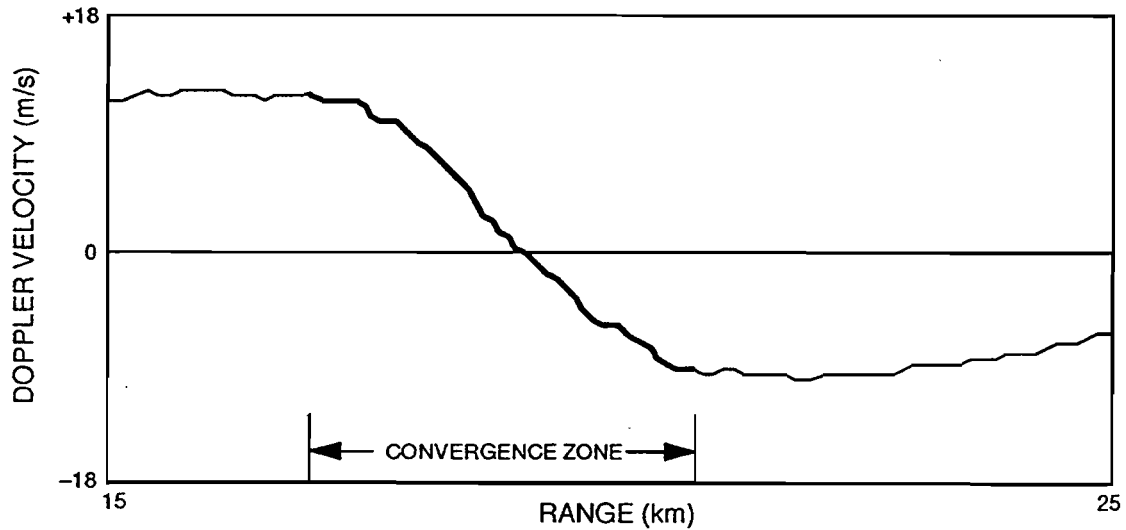


Figure 4. Example of sharp decrease in radial velocity indicating convergence associated with a gust front.

The TDWR velocity image shown in Figure 3 demonstrates this problem. The portion of the front closest to the radar site, having a direction of motion that is nearly radially aligned, has a pronounced convergent boundary. However, at the ends of the gust front, where the front is

radially aligned and the direction of motion is more azimuthal, the boundary becomes more difficult to detect.

Convergence is a viable signature for gust front detection in TDWR imagery because this radar system has sufficient sensitivity to reliably estimate Doppler values in the clear air ahead and behind gust fronts. In contrast, the same regions of clear air are usually below the sensitivity limits of the ASR-9. Doppler estimates in such regions tend to have values near the positive or negative Nyquist value as a consequence of using the pulse-pair velocity estimator on signals that have been passed through a high-pass clutter filter. Even in cases where gust fronts pass through regions of high reflectivity, convergence cannot be used reliably for gust front detection. The signal contribution from overhanging precipitation near the edges of storms can bias the low-level wind velocity estimate when there is vertical wind shear. This phenomenon has been observed to produce false convergence signatures, especially in data obtained in Kansas City, where vertical shear is more prevalent.

Note that in the case shown there is an apparent convergence signature in the ASR-9 WSP velocity image. This convergence is only an artifact of the boundary between reliable measurements within and the unusable high variance measurements ahead of the gust front. The accuracy of velocity estimations degrades markedly over the range of signal-to-noise values associated with clear air. For this reason, gust fronts typically appear in ASR-9 WSP velocity image as a thin line consisting of a band of low variance Doppler values, with high variance in the low signal-to-noise regions ahead and behind the gust front. This Doppler variance thin line is used in MIGFA as an alternative signature for gust front detection. Even though convergence is typically not visible in ASR-9 WSP radar data, implicit zones of convergence can be identified. If the environmental low-level wind velocity ahead of the storm is known, for example from LLWAS anemometers at the airport, then the velocity convergence can be inferred.

Finally, not all zones of convergence represent gust fronts in Doppler radar images. Convective circulations within storm cells can produce short lived periods of convergence. Only when a zone of convergence persists and moves in a coherent fashion is it likely to be a gust front.

2.3.2 Thin lines

The thin line echo is generally thought to be produced by the concentration of scatterers (dust, insects, rain droplets) along the leading edge of the horizontally spreading thunderstorm outflow. Some gust fronts produce a distinctive arcus cloud that is visible in satellite imagery and may also be detected by radar. The line may extend for tens of kilometers and varies in width but seldom exceeds 3 km [5].

Lacking convergence, thin line signatures become the primary means of detecting gust fronts in ASR-9 WSP imagery. In the example TDWR and ASR-9 WSP reflectivity images, the gust front is visible. Often, however, the ASR-9 WSP thin line shows only minimal contrast, is somewhat more fragmented, and does not extend as far as is apparent in TDWR data (for an example, see Figure 31).

Typical maximum reflectivities reported by the ASR-9 along gust front thin lines are in the range of 10-20 dBZ. However, much of the thin line region is comprised of echoes having reflectivities significantly lower; substantial portions of thin lines observed by the ASR-9 WSP can have reflectivities as low as -5 dBZ. Figure 5 plots the minimum detectable signal (assuming a filled beam) in equivalent weather reflectivity units (dBZ) as a function of range for the low beam of the ASR-9 WSP. For gust fronts, the filled beam assumption is unlikely to be met owing to the broad elevation fan beam, especially at ranges greater than about 10 km. Beam filling losses can reduce actual thin line reflectivity by as much as 5-10 dBZ. As evident from the figure, the resulting range of thin line reflectivities is often near the threshold of detectability for the ASR.

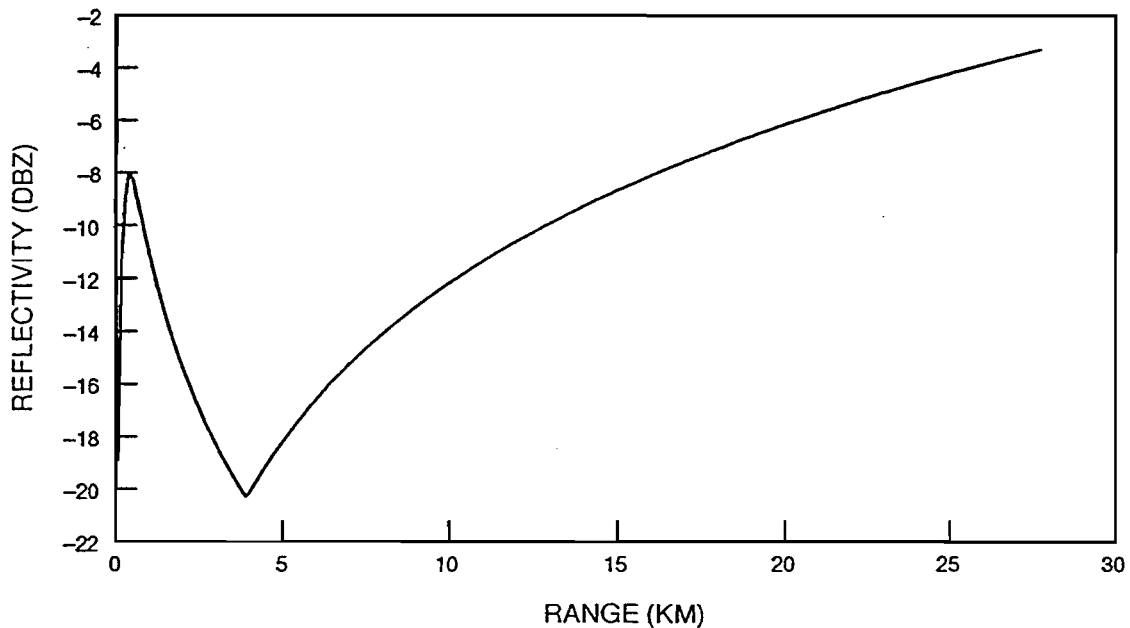


Figure 5. Minimum detectable signal expressed in weather reflectivity units (dBZ) as a function of range for the low beam of the ASR-9 WSP testbed. STC function ending at approximately 4 km is that used during 1992 operations.

Thin line signature quality often degrades at close range due to clutter obscuration and may even vanish as the gust front passes over the radar. The signature often re-establishes itself as the front moves out of the cluttered region. This problem is especially noteworthy for the ASR-9 since its on-airport location makes it more prone to detection loss when the gust front is impacting the airport.

Several weather-related phenomena can complicate thin line detection by obscuring the signatures or can cause false detections by having attributes similar to gust front thin lines. Large scale storm cells or regions of precipitation can obscure thin line signatures as gust fronts propagate through them, making detection more difficult. The integration of high and low altitude weather by the fan beam of the ASR-9 further increases the likelihood of obscuration by precipitation with this radar.

Range ambiguous echoes, also known as out-of-trip weather, can obscure thin line signatures as well as be a potential source of false alarms. These echoes occur when signals are reflected by weather more distant than the maximum unambiguous range. Because the signal has traveled farther, it arrives back at the radar receiver at the same time as signals being reflected from nearer weather. Because the apparent range extent of the storm echoes is maintained while the azimuthal extent is reduced proportional to the range, these signals have a distinctive appearance as narrow wedges of reflectivity that are radially oriented. Out-of-trip signals may have characteristics (e.g., intensity, length, width) similar to those of gust front thin lines (for example, see Figure 37 of Section 6.2). This phenomenon is well-known and very prevalent during periods of widespread thunderstorm activity. Fortunately, the pulse-to-pulse microstagger used by the ASR-9 decoheres echoes from beyond the unambiguous range of the radar; out-of-trip regions are therefore associated with large local velocity estimate variances. This characteristic can be exploited to discriminate out-of-trip weather from radially aligned gust fronts.

Other sources of false thin line signatures include: isolated and elongated low-intensity precipitation echoes, residual ground clutter, dust, and flocks of birds.

2.3.3 Motion

A final key gust front signature is motion. If sequential radar scans are compared, convergence and thin line signatures of a gust front move conspicuously in a direction approximately perpendicular to the orientation of the convergence boundary and reflectivity thin line. Signatures that do not move are either not gust fronts (e.g., false alarms from range ambiguous echoes, storm edges, or ground clutter) or are gust fronts that are not operationally significant. Within limits, gust fronts tend to move as a straight line or as an outwardly expanding curve. The propagation speed and direction tends to be consistent along its length and across time. While these tendencies apply to gust fronts in isolation, gust fronts may move more erratically as they collide with other outflow boundaries.

Figure 6 plots gust front propagation speeds for Kansas City and Orlando gust fronts. Mean propagation speeds were calculated by dividing the total gust front track distance by the total

time of the event. While mean propagation speeds for the Kansas City gust fronts varied widely, ranging from nearly stationary to over 20 m/s, a significant number of them (46%) moved with speeds greater than 10 m/s. This contrasts with the Orlando gust fronts which were nearly always slow-moving, ranging from 3 m/s to 14 m/s with an average of 8 m/s. Only 12% of the Orlando gust fronts examined moved faster than 10 m/s (some of the slower Orlando “gust fronts” were most likely sea breeze fronts).

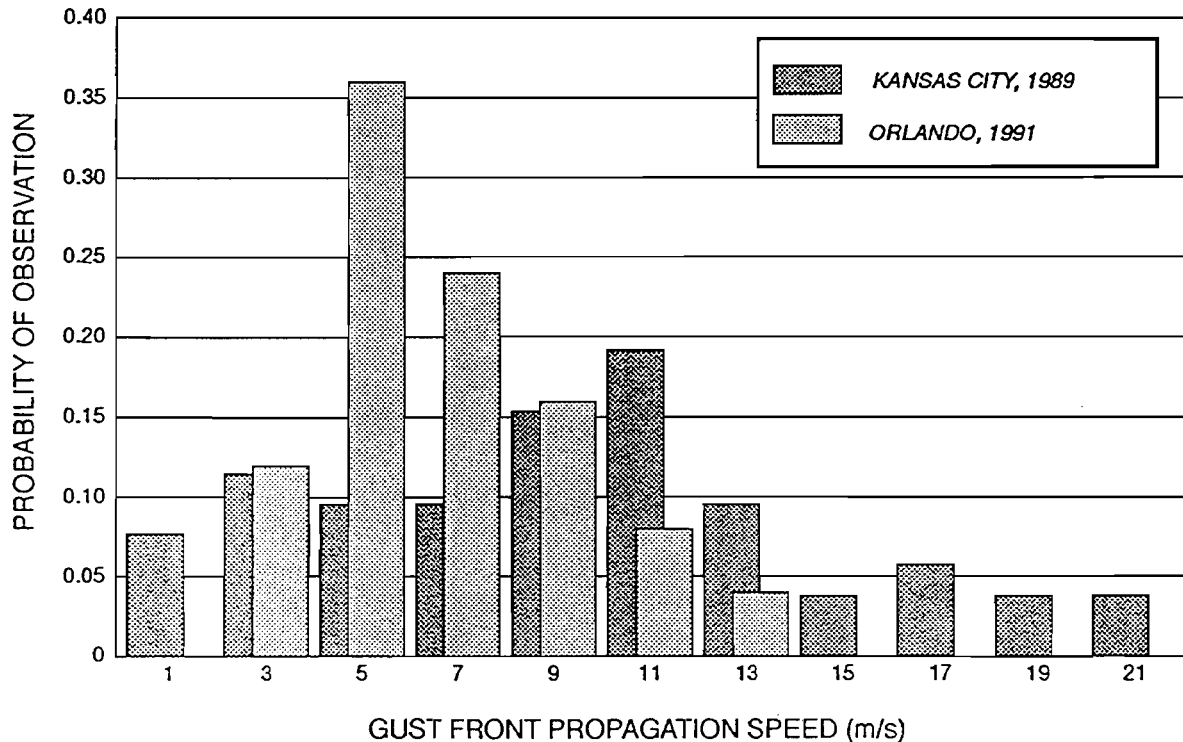


Figure 6. Distribution of gust front propagation speeds for Kansas City and Orlando.

Of course, not everything that moves is a gust front. Ambient winds can carry thin wisps of low-intensity precipitation that move in a manner indistinguishable from gust fronts in the absence of (explicit or implicit) convergence data.

2.3.4 Signature ambiguity

For both radar systems, gust front detection is not a trivial problem. Although gust fronts display multiple signatures, no signature is 100% reliable. For each signature discussed above, not all gust fronts will have the signature and not all visible signatures will be associated with gust fronts. For gust fronts that do display all the expected signatures, detection is easy. But oftentimes, gust fronts will display only a subset of the signatures and can display them in an ambiguous manner. And, none of these signatures are unique only to gust fronts. Many phenomena unrelated to gust fronts can display one or more of the signatures.

Only by weighing the quality of several signatures simultaneously can gust front detection performance reach levels comparable to human performance. The task is hard enough for the TDWR. With the ASR-9, the task is much more intimidating. Without the corroborative evidence of velocity convergence, discriminating between moving thin line features associated with a gust front and ones associated with other phenomena becomes more difficult.

2.4 Required Algorithm Products

The gust front algorithm provides air traffic controllers with products that enable them to perform two vital tasks: warning and planning. The warning function is fulfilled by the automatic generation of wind shear alerts (WSA) when a gust front of sufficient strength impacts a runway approach or departure corridor. The planning product alerts the ATC supervisor when a change in wind speed or direction associated with an approaching gust front may require a change in the active runway. Information generated by the algorithm is relayed to controllers and supervisors using two types of displays:

1. Geographical Situation Display (GSD).
2. Ribbon Display Terminal (RDT).

Figure 7 shows the GSD display provided to air traffic control supervisors at the Orlando FAA facility. The GSD overlays weather data in graphic format onto a map of the terminal airspace. Its design allows the tower or Terminal Radar Approach Control (TRACON) supervisor to quickly assess the current weather conditions in the terminal area. This information allows controllers to more effectively vector aircraft by avoiding weather impacted areas or anticipating changes in active runway due to changing winds.

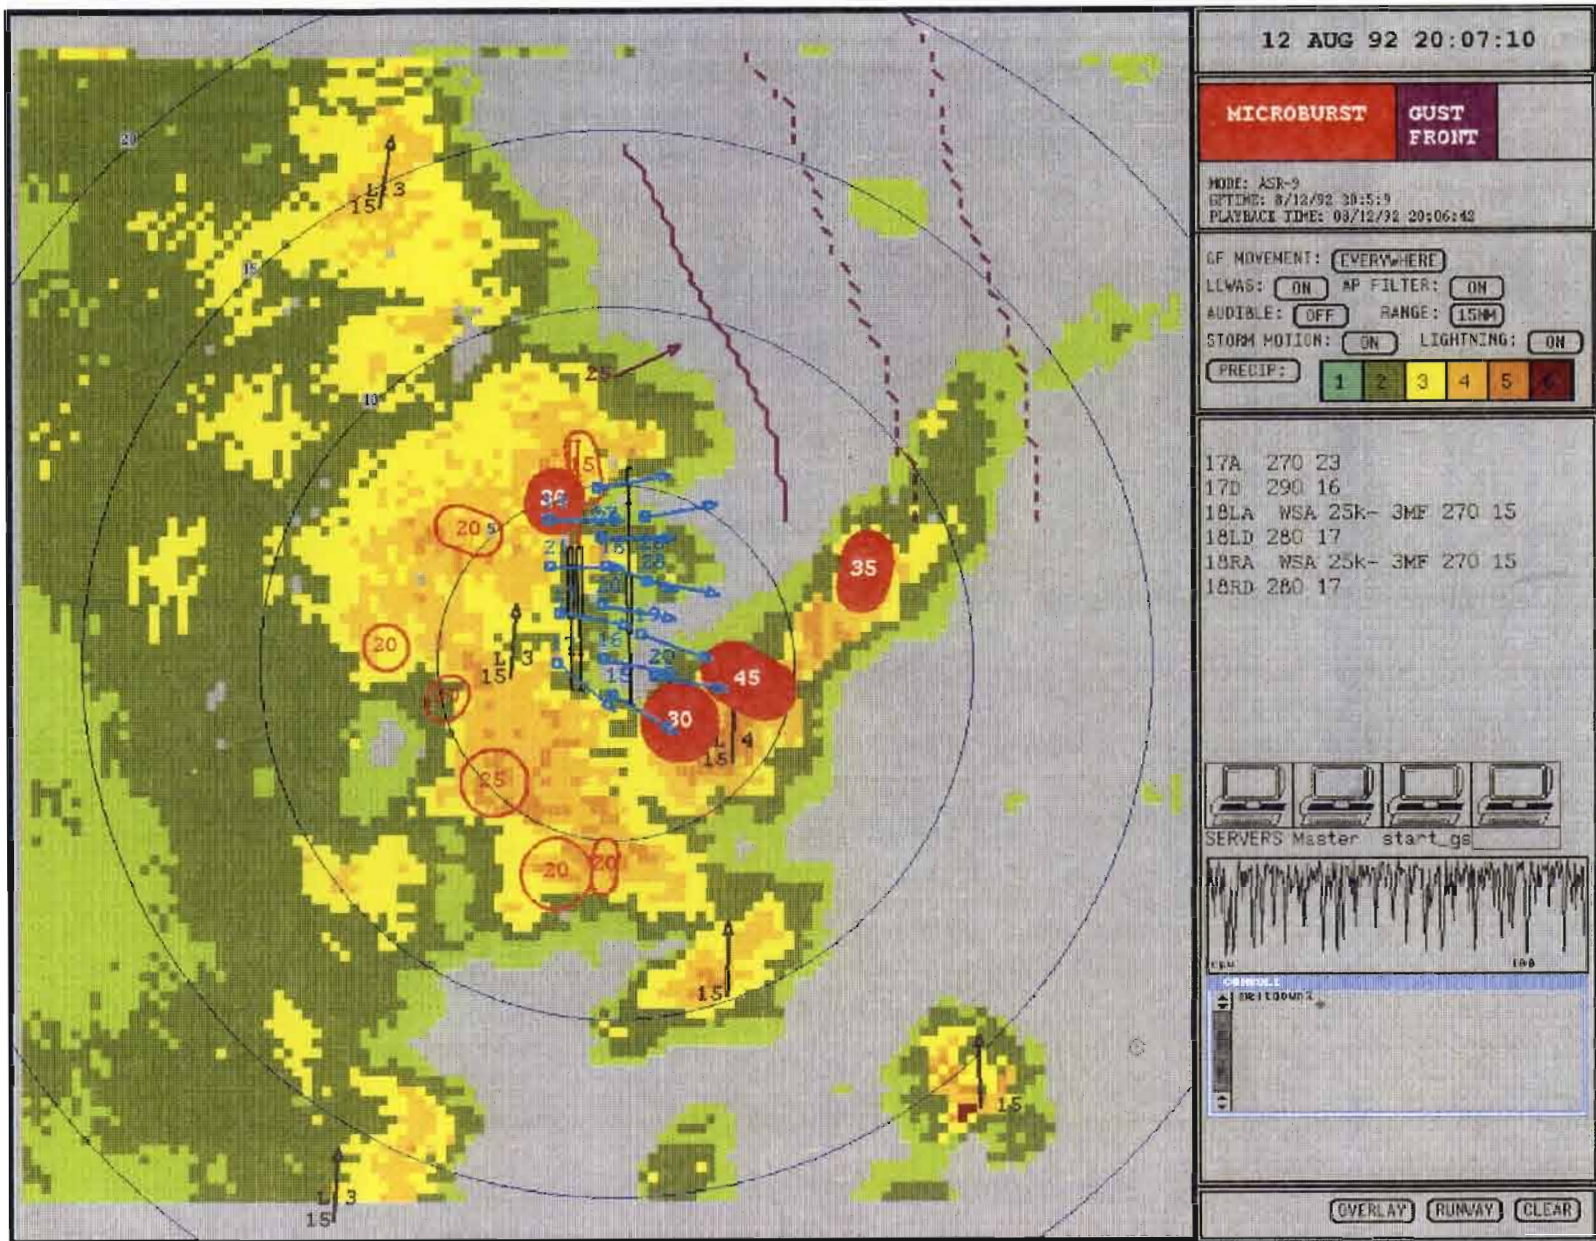


Figure 7. Example GSD showing the current position of a gust front (solid purple curve) along with 10 and 20 minute forecast gust front positions (dashed purple curves). The winds behind the front are indicated by the purple numbered arrow (25 knots from the southwest).

The location of a detected gust front is depicted on the GSD as a solid purple curve. The expected wind speed and direction behind the gust front are indicated by a purple numbered arrow displayed behind the gust front. Ten and twenty minute forecast locations of the gust front are represented by dashed purple curves. Gust front locations and wind shift estimates are displayed on the GSD regardless of whether they are impacting the runways or are strong enough to require a wind shear alert to be issued. Weather information generated by other hazardous weather detection algorithms are also displayed on the GSD. These include: microburst regions, 6-level precipitation intensity, and storm motion vectors.

A wind shear alert is issued when a gust front crosses the runways or is within three miles of ends of the runways, and the air speed gain exceeds a threshold, nominally 15 knots. An alphanumeric alert message which indicates the affected runway and the strength of the shear is displayed on the controller's RDT. These messages are relayed to pilots by the controller when aircraft are cleared for takeoff or final approach. An alert message issued for a gust front might look something like:

17A WSA 20K+ 2MF 330 30

This message indicates that a gust front with a wind speed gain of 20 knots is located on the 2 mile final approach to runway 17. Wind speed and direction (from LLWAS) at the approach end of the runway are from 330 degrees at 30 knots.

In order to generate the gust front product displays, the following information needs to be supplied by the algorithm for each detected gust front:

1. Date and time of the current detection.
2. An ordered list of points (x,y) defining a curve that marks the location of the detected gust front with respect to the radar origin.
3. Wind shear hazard number (differential velocity across the front).
4. Gust front propagation speed and direction.
5. Speed and direction of winds behind the gust front.
6. Ordered lists of points (x,y) defining curves that mark the forecast locations of the currently detected gust front at T+10 and T+20 minutes. Additional forecasts for other time intervals (e.g., T+1, T+11, T+21) may be required if such forecasts are to be used to update the position of the current detection and the 10 and 20 minute forecast positions in between actual algorithm updates (See Section 4.11).

3. PRIOR RESEARCH AND DEVELOPMENT

Automated radar gust front detection algorithms have been under development and evolution for almost 10 years. Uyeda and Zrnica [6] described an automated detection algorithm developed for the Next Generation Weather Radar (NEXRAD) that was based solely on detecting gradients of radial velocity, i.e., radial convergence. Their algorithm was successful in locating and tracking strong gust fronts that commonly occur in Oklahoma in the spring.

An improved version of the initial algorithm, known as the Gust Front Detection Algorithm (GFDA), reduces false alarms by requiring vertical association of gust front signatures from two low-elevation scans (typically 0.5 and 1.0 degrees). The GFDA also incorporates a technique for estimating horizontal winds ahead of and behind detected gust fronts [7, 8, 9]. Like its predecessor, the GFDA detects velocity convergence in radials. This algorithm is the one currently intended for use in the initial operational deployment of TDWR systems.

Figure 8 illustrates the steps used by the GFDA to recognize gust fronts. Radials are searched for runs (segments) of decreasing radial velocity, indicating convergent shear. Segments whose shear values exceed a pre-determined threshold are logically grouped into features on the basis of end point proximity and overlap tests. Feature attributes are then tested against a number of thresholds and are either kept, discarded or combined with other features. After separately processing each of the two full-circle scans, the algorithm tests for vertical continuity of features between the scans. Features that exhibit vertical continuity and exceed a minimum length threshold are declared as gust fronts. The reported location of the detected gust front is determined by fitting a curved line through the peak shear of each segment in the gust front feature. Sequential detections are associated over time to build detection histories for each gust front upon which propagation speeds are estimated and forecasts are generated.

Lincoln Laboratory, in conjunction with NSSL, has developed an Advanced Gust Front Algorithm (AGFA) [11, 12], which further improves upon GFDA by augmenting radial velocity pattern recognition with azimuthal shear and/or reflectivity thin line detection. Azimuthal shear features are recognized by searching across azimuths along arcs of constant range for variations in the velocity field. These azimuthal variations may be the only available evidence for fronts that are nearly radially aligned and thereby suffer from poor Doppler viewing angles. Extraction of azimuthal shear features proceeds in similar fashion to that described above for extraction of radial shear features.

The AGFA thin line feature extraction algorithm uses a successive thresholding scheme to extract segments of consecutive data elements associated with local "ridges" in the reflectivity field (see Figure 9). Segments that are too short or too long to belong to a gust front thin line are immediately discarded. Segments having data elements that are also members of a segment extracted at the next higher threshold level are eliminated from further consideration. In this fashion, segments corresponding to the crest of the reflectivity thin line ridge are identified.

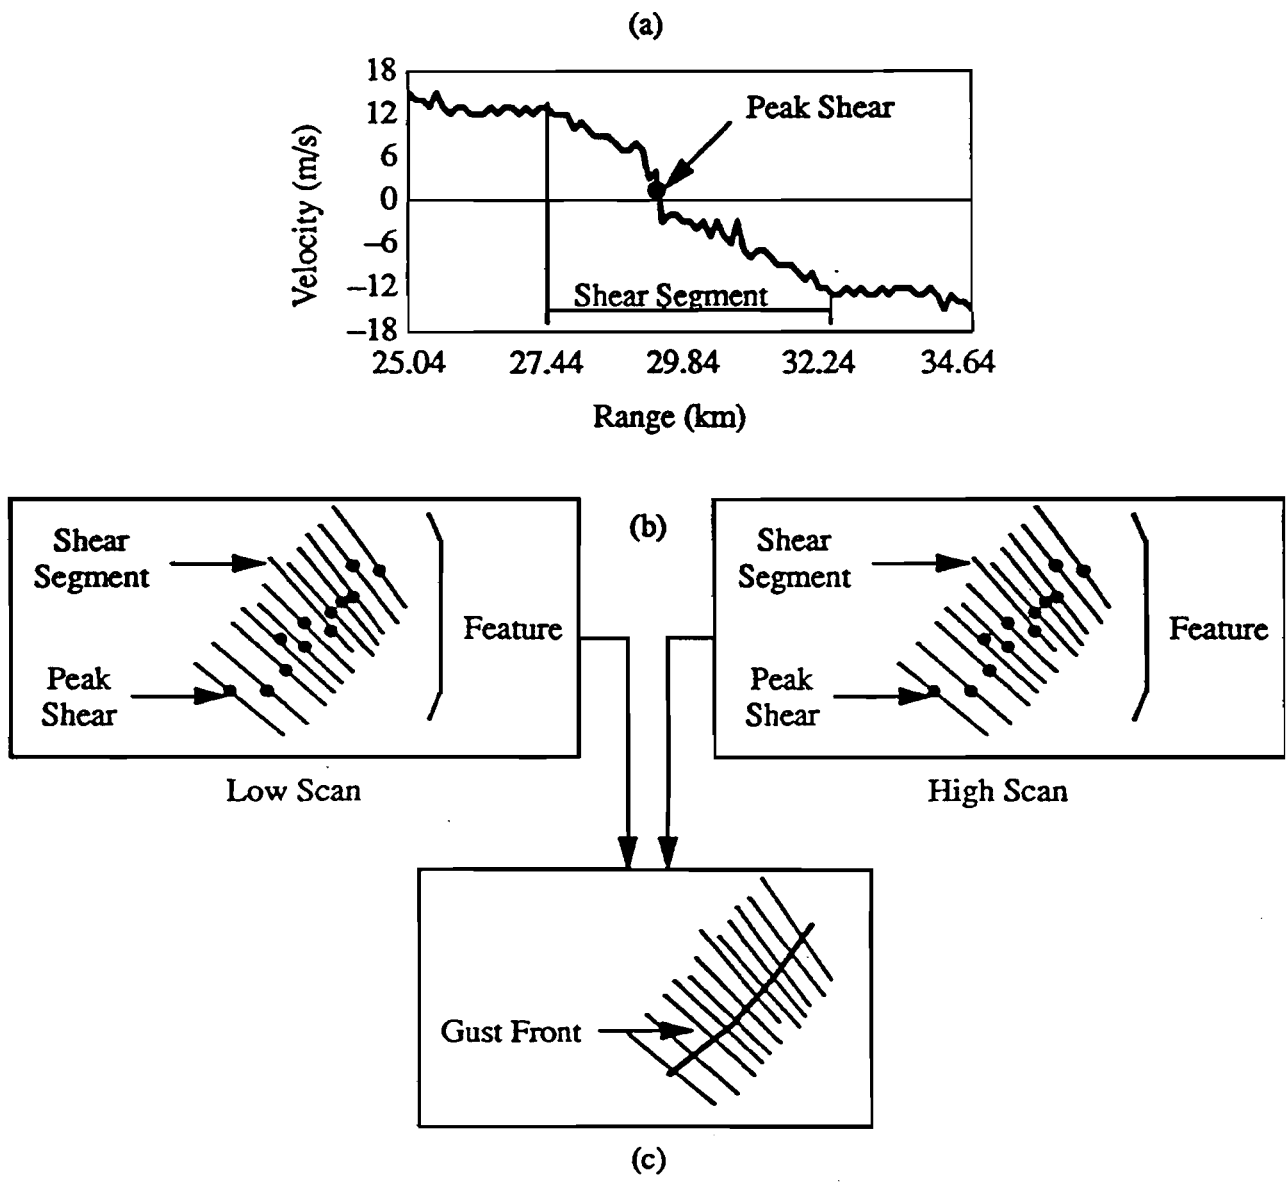


Figure 8. Principal steps of the GFDA algorithm. (a) Radial velocity data are searched for segments of decreasing Doppler velocity. The dots represent the locations of the peak shears along each segment. (b) Segments from the low and high scans are associated into features, and (c) features are grouped together based on spatial proximity into gust fronts. (From Klinge-Wilson et al., 1992 [10]).

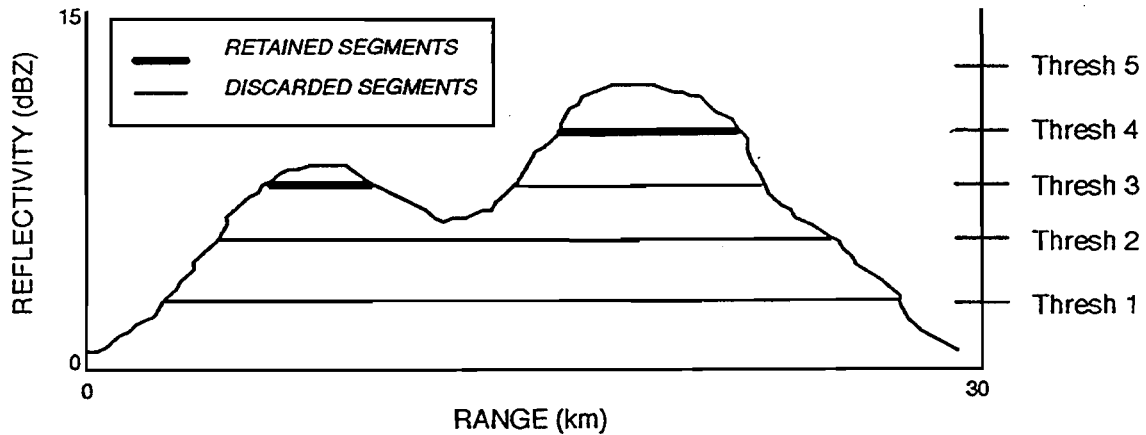


Figure 9. Thin line thresholding procedure applied along a radial of reflectivity data.

Thin line segments are generated in two passes through the reflectivity field. The first pass constructs segments over all range gates along a radial. The second pass constructs segments across radials along arcs of constant range. Segments built at each pass are associated to form thin line regions by requiring some degree of end point overlap along with spatial proximity. The association scheme permits bridging of small gaps across one or more missing segments, which often occurs with fragmented thin line signatures. The final thin line features consist of point lists generated by connecting the midpoints of each of the segments in the features. In AGFA, azimuthal shear and reflectivity thin line features are associated with radial convergence features on the basis of end point proximity and orientation.

Neither GFDA nor AGFA uses motion as a signature for detecting gust fronts. In these algorithms, motion is used only in heuristics that reject gust fronts after they have been extracted.

Klinge-Wilson et al. [10] compared performance of the GFDA with AGFA configured to use reflectivity thin lines but not azimuthal shear. Using gust front data collected with the Lincoln Laboratory TDWR testbed radar (FL-2) in Denver, Kansas City, and Orlando, FL, they found that AGFA provided an improvement of about 37% in overall gust front length detected, confirming the usefulness of thin line signatures (scoring methods are described in Section 5.1).

During field testing in 1990 and 1991, a customized version of AGFA was used for gust front detection on an ASR-9 WSP [2, 3]. Because of the lack of reliable velocity convergence features, the ASR-9 version of AGFA was configured to operate in a thin-line-only detection mode. The

algorithm was successful in detecting the gust fronts that had good quality thin line signatures, but had difficulty detecting gust fronts with weak or fragmented reflectivity thin line signatures. Not surprisingly, without convergence the algorithm was prone to false alarms triggered by elongated, low reflectivity weather echoes, such as are sometimes associated with stratiform rain. Installing sub-optimal detection thresholds in order to reduce the false alarm rate further reduced detection probabilities.

In the above study, the scoring was done against human interpretations of the same images. In part the discrepancy between human and AGFA performance is due to the fact that AGFA does not make full use of a variety of evidence and information available in the ASR-9 WSP data. These include (but are not limited to) velocity thin lines, thin line motion against a relatively static background, and high velocity variance of range-ambiguous echoes. Moreover, algorithms such as GFDA and AGFA, which rely on sequentially applied thresholds to discriminate gust fronts from background, do poorly when faced with weak, uncertain, contradictory, or ambiguous information. In particular, various forms of threshold-based editing used to remove “bad” data often discard useful information. This kind of binary decision making at early stages of processing can result in the elimination of potentially relevant information, unnecessarily limiting detection performance. Finally, these early gust front algorithms have no systematic means of conditionally fusing information from various sources. Different signatures can have varying reliability that depends on situational context. Knowledge of how reliability varies with context is an important tool in assimilating and combining sources of evidence.

4. ALGORITHM DESCRIPTION

4.1 Design Philosophies

MIGFA is an adaptation of the eXperimental Target Recognition System (XTRS) developed at Lincoln Laboratory [13, 14]. MIGFA has inherited the development environment, control structure, knowledge-based signal processing, and several other important attributes from XTRS.

The conventional wisdom in computer vision/object recognition research has been to use general image processing operations, ideally devoid of object- and context-dependent knowledge, at the initial stages of processing. Such operations might include edge detection, segmentation, cleaning, or optic flow estimations. And yet, the ideal is never really achieved in practice. For example, in order to effectively detect edges, some knowledge of the sensor and the expected scene contents are (implicitly) encoded in the form of thresholds or other such parameters. From the results of such general operations, image characteristics are extracted and represented symbolically. Artificial intelligence is then applied, as if by definition, only on the symbolic representations at “higher” levels of processing.

In contrast, sensor-, object-, and context-dependent knowledge is applied in the earliest (image processing) levels of XTRS and MIGFA processing. Knowledge of the problem is used in three ways. Knowledge is used to select from a library those feature detectors that are *selectively indicative* of the object being sought. Knowledge is also incorporated within feature detectors through the design of matched filters that are customized to the physical properties of the sensor, the environment, and the object to be detected. Finally, knowledge of the varying reliability of the selected feature detectors is used to guide data fusion.

With multiple, independent feature detectors, XTRS and MIGFA have the capability for easily adapting to different contextual conditions. At the beginning of processing of each scan, a rule-based expert examines contextual information to select a set of feature detectors known through experience to be most effective for a given set of circumstances. For example a different set of feature detectors is already used for TDWR and ASR-9 WSP data. As the system is tested at sites other than central Florida or during different seasons other than mid-summer, we may find that conditions are sufficiently different to warrant the use of alternative sets of feature detectors.

Because of the redundancy inherent in the use of multiple feature detectors, MIGFA tends to be robust: The malfunction of a feature detector or even the absence of one data source does not necessarily halt processing and may have only minor effects on detection performance. A second reason for robustness is delayed thresholding. Many computer vision systems apply sequential thresholds to discriminate an object from background. However, especially when given ambiguous data, each applied threshold closes off options for detecting an object. A better strategy, one attempted in XTRS and MIGFA, is to apply one threshold only after evidence from all sources of information have been meaningfully fused into a single map of evidence.

4.2 Knowledge-based Signal Processing

Knowledge-based signal processing is the product of two specific techniques: a generalized matched filter and a means of fusing the evidence generated by multiple matched filtering operations into a single map of evidence.

4.2.1 Functional Template Correlation (FTC)

Functional template correlation (FTC)[15], is a generalized matched filter incorporating aspects of fuzzy set theory. As a basis for understanding, consider the basic image processing tool *autocorrelation*. Given some input image I , an output image O is generated by matching a kernel K against the local neighborhood centered at each pixel location I_{xy} . The match score assigned to each pixel O_{xy} is computed by multiplying each element value of K by the superimposed element value in I and summing across all products. If the shape to be matched can vary in orientation, then the pixel I_{xy} is probed by kernel K at multiple orientations. The score assigned to O_{xy} is the maximum across all orientations.

FTC is fundamentally the same operation with one important exception: Where the kernel of autocorrelation is an array of image values (essentially a subimage of the image to be probed) the kernel used in FTC is an array of *scoring functions* that return scores reflecting how well image values match expectations at each element of the kernel. The set of all returned scores are averaged and clipped to the continuous range $[0,1]$. The output of FTC is a map of these values, each of which reflects the degree of belief that the shape or object implicitly encoded in a FT is present at that image location. As an example, consider the FT implementation of a simple matched filter shown in Figure 10, which is designed to detect gust fronts in reflectivity data. Gust fronts are observed as thin lines of moderate reflectivity (approximately 0 to 20 dBZ), that are flanked on both sides by low reflectivity values (approximately -10 to 0 dBZ). On the left is the template kernel consisting of integers, corresponding to the 2 scoring functions shown on the right. Elements of the kernel that do not have an index form guard regions in which image values are ignored and have no effect on match scores. Scoring function 0, corresponding to the flanking regions of low reflectivity, returns a maximal score of 1.0 for image values in the interval of -20 dBZ to -5 dBZ, a gradually decreasing score for image values in the interval -5 dBZ to 10 dBZ, and a score of -2.0 for image values larger than 10 dBZ. Scoring function 1, corresponding to the center of the kernel where moderate reflectivity values are expected, returns maximal scores in the interval between 5 and 12.5 dBZ with gradually decreasing scores for both higher and lower image values. Note that while very low image values can generate scores of -1.0, a slower decline in score with a minimum score of 0.0 is evident for image values above the maximal scoring interval. This asymmetry reflects the belief that gust fronts are often obscured by storm cells. The relatively greater tolerance to high reflectivity values increases the chance of matching gust fronts that are partly obscured.

In general, by increasing or decreasing the interval over which affirming scores (i.e., scores > 0.5) are returned, scoring functions can encode varying degrees of uncertainty with regard to what image values are allowable. But in addition, knowledge of how a feature or object appears in sensor




Figure 10. Example functional template for feature detector TL-DZ.

imagery can be encoded in scoring functions. With various design strategies, the interfering effects of occlusion, distortion, noise, and clutter can be minimized[16]. As a consequence, matched filters customized for specific applications using FTC are generally more robust than classical signal processing operations. In the thin line matched filter example shown above, the filter does not simply find thin lines, but can select those thin lines having reflectivity values within a particular range. Furthermore, the matched filter can display differential tolerances to image values that are higher or lower than the expected range of values. In our automatic target recognition (ATR) systems, FTC has been used primarily as a direct, one-step means of 3-D object detection and extraction. In MIGFA, FTC is used more as a signal processing tool for edge detection, thin line filtering, thin line smoothing, shape analysis, and thinning of shapes.

It is obvious that if FTC were implemented literally as described, the computational expense would be prohibitive for all but the most trivial problems. However, by scaling input data to a fixed integer range (say 0 to 255) and implementing the scoring functions as a precomputed 2-D lookup table indexed by a scoring function number and an image value, FTC is actually faster than autocorrelation.

4.2.2 Interest Images

The second major tool is the use of “interest” as a medium for data fusion and for assimilating evidence at the pixel level[16]. An interest image is a map of evidence (values in the range [0,1]) for the presence at each pixel location of some feature that is selectively indicative of an object being sought (note that the output of FTC is an interest image as long as the functional template encodes a relevant feature). Higher pixel values reflect greater confidence that the intended feature is present at that location. Using interest as a “common denominator”, data fusion can be accomplished by combining interest images derived from various pixel-registered sensory sources. Interest images can also be conceptualized as maps of evidence for the presence of a targeted object at each image location. Using simple or arbitrarily complex rules of arithmetic, fuzzy logic, or statistics, pixel-level evidence from several coregistered sources can be assimilated into a single combined interest image. Clusters of high interest values are then used to guide selective attention and can serve as the input for object extraction. If done effectively, the combined interest image provides a better representation of object shape than is evident in any single sensor modality.

4.3 Overview

The system block diagram in Figure 11 illustrates MIGFA as configured for the ASR-9 WSP version of the system. In preparation for processing, input images V (Doppler velocity image) and DZ (reflectivity image) from the current radar scan are converted from polar to Cartesian representation and scaled to a useful resolution. Image SD is a local standard deviation texture map of V. These images are then passed to multiple, simple, independent feature detectors that attempt to localize features that are selectively indicative of gust fronts. The output of these feature detectors, most based on some application of FTC, is expressed as an *interest image* that specifies evidence for where and with what confidence a gust front may be present. These interest images are fused to form a combined interest image, providing an overall map of evidence for where gust fronts are believed to exist.

From the combined interest image, fronts are extracted as chains of points. The chains extracted from a radar scan, collectively called an event, are integrated with the prior history by establishing point-to-point correspondence. Heuristics are applied at this point to reject chain points that have apparent motion that is improbable. The updated history is used to make predictions of where points along the front will be located at some future time. Such predictions are used in the processing of subsequent images, specifically in the feature detector called ANTICIPATION, which places high interest values wherever fronts are expected to be and by so doing selectively sensitizing the system to detecting gust fronts at specific locations. Anticipation is tuned so that it will not by itself automatically trigger a detection, but when averaged with other interest images will support weak evidence that would otherwise be insufficient to trigger a detection.

LLWAS anemometer data are used in conjunction with wind parameters estimated from radar data to compute a wind shift and a wind shear hazard for each detected gust front.

Figure 12 is a summary of the processing steps for a single example.

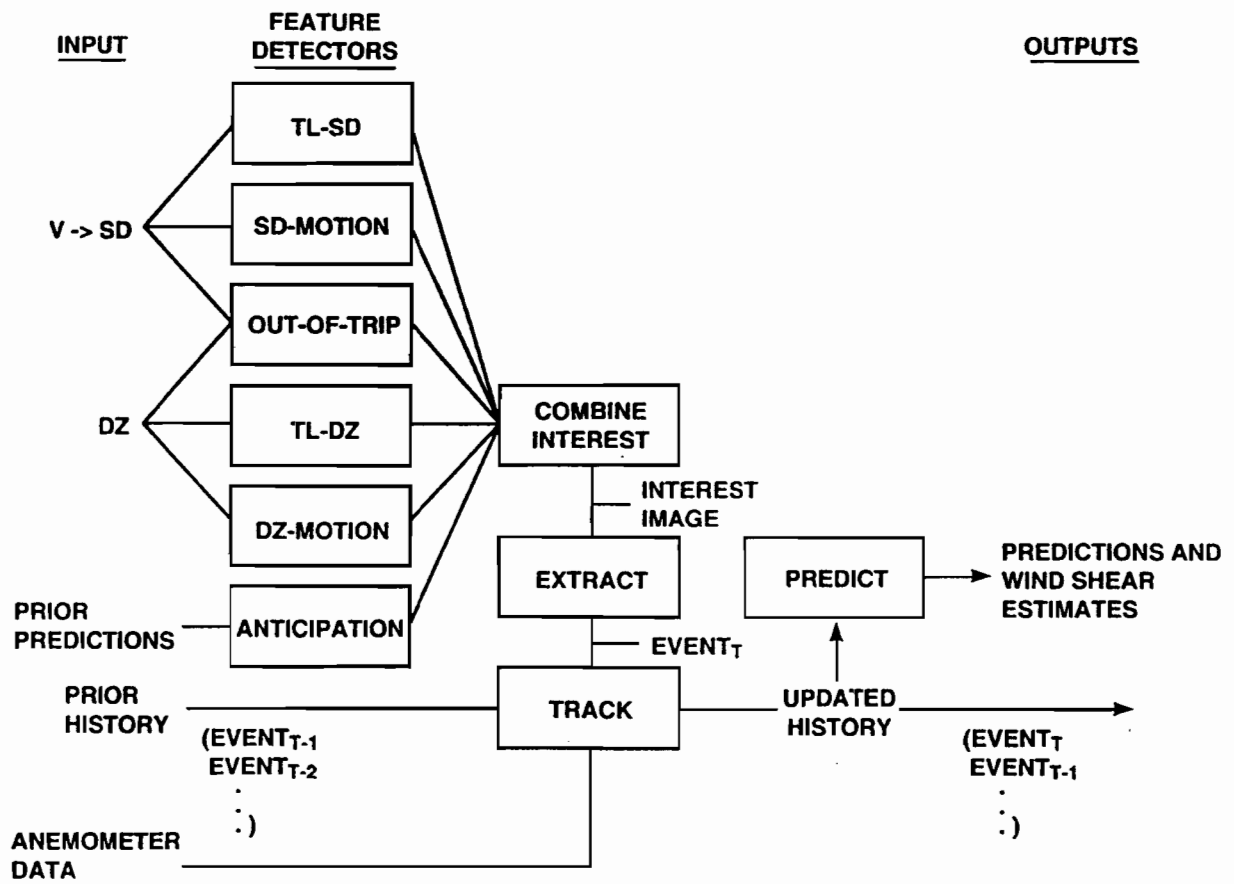


Figure 11. MIGFA block diagram.




Figure 12. Processed scan summary: On the first row are DZ, SD and the combined INTEREST image computed from them. The second row begins with the extracted EVENT. In the frame marked HISTORY, the current indexed chain (white) is shown in context with preceding scans (darker shades of gray indicate more distant events in time). In the frame marked PREDICTIONS, gray pixels show the 10 and 20 minute forecasts of where the fronts are expected to be.

4.4 Image Preparation

As discussed in Section 2.3, velocity convergence is an unreliable signature for gust front detection in ASR-9 data. However, gust fronts are detectable in velocity images. Because of the tendency for high pass clutter-filtered pulse-pair Doppler estimates in a velocity image to assume the +/- Nyquist value in regions of low signal to noise ratio, the local velocity variance is higher for an area of clear air than for an area associated with slightly higher reflectivity values. This information is translated into a usable form by transforming the velocity image V into a texture map of local standard deviation (image SD). At each pixel of V , the standard deviation was computed in a 5×5 window and assigned to the corresponding pixel in SD . An example transformation from V to SD (image DZ is shown for reference) is shown in Figure 13. Note that although the images are shown in Cartesian format, the computation of standard deviation was done in polar format.

In order to support subsequent FTC operations on the input imagery, pixel values for all images are scaled such that the interval between the lowest and highest meaningful values are assigned values between 0 and 255, respectively. Each image is tagged with the scaling factor and offset necessary to translate scaled values back to the original physical values.

Finally, the images DZ , V , and SD are converted from polar arrays (240 range bins \times 256 radials) to Cartesian arrays (130×130). Mapping is done by computing for each element of the Cartesian array the range bin and radial at which the value is to be found in the polar array. An implicit subsampling of the data occurs during remapping. From an initial radial resolution of 115.75 meters per range bin and a minimum azimuthal resolution of 680 meters at 28 kilometers, the final Cartesian image has a pixel resolution of 480 meters per pixel.

4.5 Feature Detection

Given contextual information of the sensor being used, the location of the sensor, and environmental conditions, a rule-based expert selects an appropriate set of feature detectors to be applied to the input data. For the moment, the only rule chooses between two sets of feature detectors: one set customized for the ASR-9, the other customized for the TDWR. These two sets alone may be sufficient. However, as MIGFA testing is expanded, additional rules adapting the feature detector set for sites other than central Florida and seasons other than summer may be added as needed. Although we do not anticipate the need, the set of feature detectors can be reconfigured dynamically from one scan to the next as changes in weather conditions are detected.

The set of feature detectors that follow are the ones that were tested for the ASR-9 WSP in Orlando during the summer of 1992.

4.5.1 Thin line Filtering

Two simple thin line feature detectors are used, one for each of the input images DZ and SD (TL- DZ and TL- SD in Figure 11). The feature detector TL- DZ is based primarily on the application of the functional template illustrated in Figure 10. An example DZ image and the

interest image resulting from the application of this template are shown in Figure 14. Note that the thresholding of the interest image in Figure 14 is only for the purpose of visualization and is not part of the algorithm. Regions of red pixels in interest images are areas of nil (missing) values where a feature detector is considered to be unreliable regarding the presence of a gust front. For example, the interest image shown in Figure 14 has a small red circle of missing values over the radar site. Because gust front observability decreases dramatically as gust fronts cross over the radar site, any negative evidence located in the area over the radar site is unreliable.



Figure 13. Standard deviation texture map.




Figure 14. DZ thin line detector: (A) The input DZ image. Red pixels indicate bad or missing values. (B) The interest image generated by FTC. Red pixels indicate locations for which the detector is unreliable and therefore does not express an opinion. (C) The output interest image thresholded at the level of ambiguity (0.5).

The functional template applied to SD (Figure 15) has the same kernel as the one for the TL-DZ. Only the scoring functions have been changed. In image SD, thin lines associated with gust fronts have low standard deviation values within the front and high values ahead and behind the front. Consequently, scoring function 1 (the center strip) returns maximal scores for low values while scoring function 0 (the flanking regions) returns maximal scores for high values. The result of applying FTC to the SD image corresponding to the DZ example in Figure 14 is shown in Figure 16.

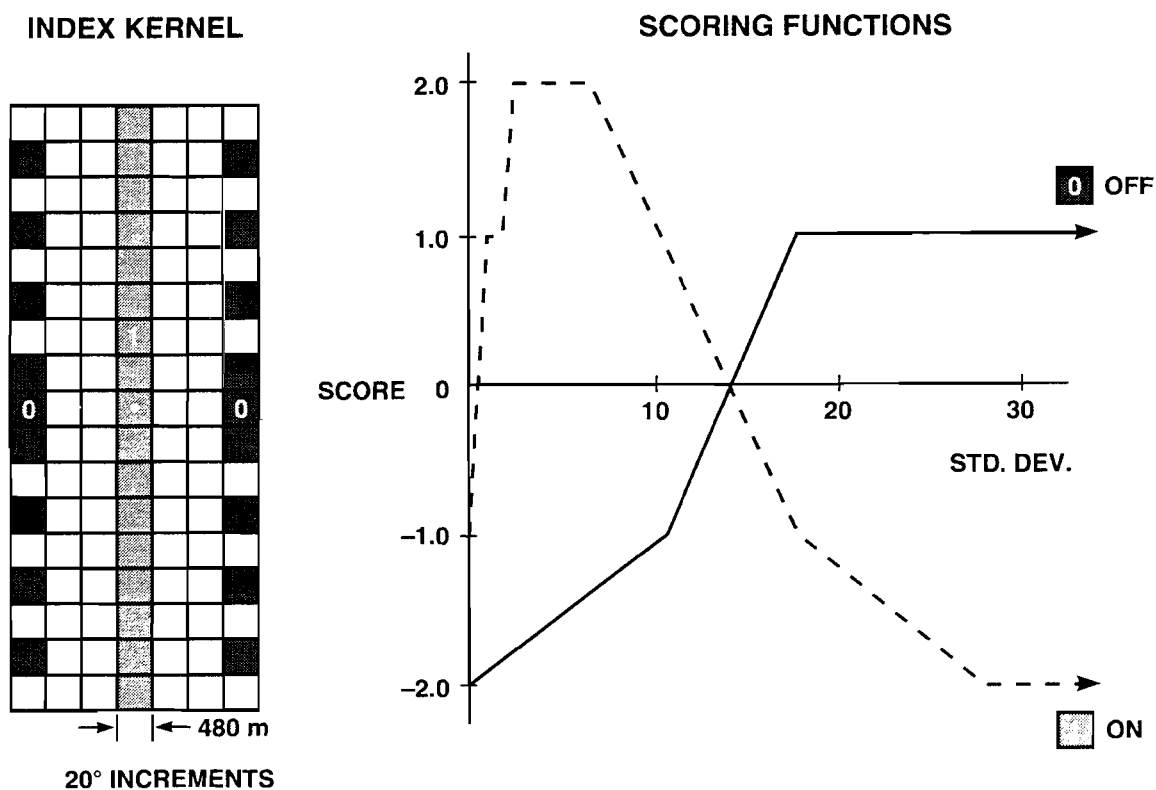


Figure 15. Functional template for feature detector TL-SD.




Figure 16. SD thin line detector: (A) The input V image. (B) The translated SD texture map. (C) The interest image generated by FTC. (D) The output interest image thresholded at the level of ambiguity (0.5).

4.5.2 Thin line motion

Two thin line motion detectors, which are very similar to the basic thin line detectors described above, are also used. Motion is based on simple differencing. In the case of feature detector DZ-MOTION (Figure 19), the DZ image from some previous scan, typically 4 minutes before the current scan, is subtracted from the current image DZ. As can be seen in the frame labeled DIFFERENCE, gust fronts appear as white lines (positive values where the front is now) being trailed by a parallel dark line (negative values where the front was in the previous scan). Although functional templates that scan simultaneously for parallel white and dark thin lines are feasible, these have so far proven to be too computationally expensive to operate within the realtime constraints of the available computer resources. Consequently, the existing motion detectors simply look for thin lines of positive values. The functional template used has a kernel that is identical to the one shown in Figure 10. However, as can be seen in Figure 17, the scoring functions are somewhat different, reflecting various effects of differencing. The output of FTC is the frame labeled MATCH in Figure 19. Thin lines that do not move do not appear in the DIFFERENCE image and therefore




Figure 17. Functional template for feature detector DZ-MOTION.

receive low interest values, reflecting a belief that a thin line that is not moving is either not a gust front or is a gust front that can be ignored. Because the background in differenced images is reduced to near 0 values, the two motion detectors tend to be more sensitive than the thin line detectors described earlier.

One disadvantage of the implemented motion detectors is that as a storm moves, the leading edge of the storm may appear in the DIFFERENCE image as a thin line of positive values. Consequently, both motion feature detectors tend to produce false alarms on moving storms (examples can be seen in the frame labeled MATCH). In order to reduce the likelihood of such false alarms, an interest image of precipitation (the frame marked CELLS) is generated using a functional template whose kernel is round with a diameter of 13 pixels (6.25 km) (see Figure 18). Large patches of reflectivity values roughly greater than 2 dBZ will result in match scores greater than 0.5 in the CELLS image. As is seen in Figure 19, the DZ thin line motion detector subtracts pixel-values in CELLS from the MATCH image, resulting in the frame labeled INTEREST. The final frame shows those pixels whose final interest values are greater than ambiguity (again, the thresholding at 0.5 is done only for visualization). The feature detector SD-MOTION is analogous to DZ-MOTION:

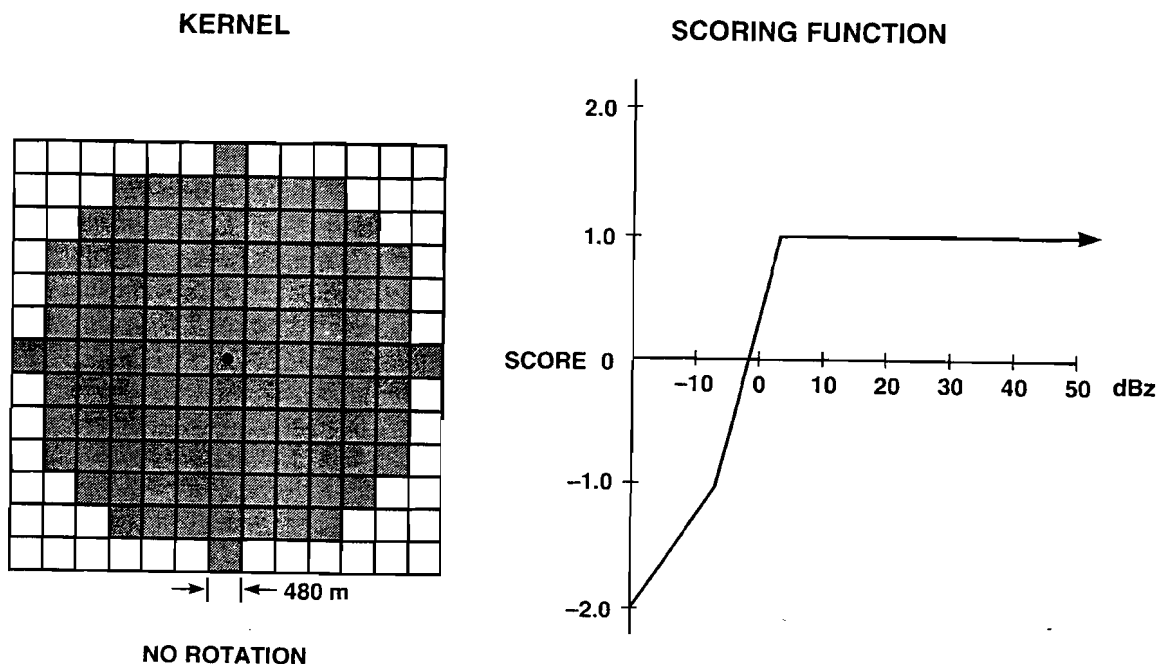


Figure 18. Functional template for detecting patches of precipitation in reflectivity images.

a thin line filter is applied to the difference of two sequential SD images (see Figure 20). However, in this case the current SD image is subtracted from the previous SD image so that where the gust front appears in the current scan appears as a thin line of increased values. In SD-MOTION, wherever the CELLS image has interest values above 0.5, the corresponding values of INTEREST are set to nil, indicating that the detector does not have an opinion for or against the presence of a gust front and defers to the evidence generated by the other feature detectors.




Figure 19. Feature detector DZ-MOTION.



Figure 20. Feature detector SD-MOTION.

4.5.3 Out-of-trip Weather

As discussed in Section 2.3, range ambiguous echoes have a distinctive appearance as reflectivity thin lines that are radially aligned, with corresponding Doppler values having high local variance. Because of the thin line appearance, TL-DZ and DZ-MOTION both tend to incorrectly give high interest values to them.

Out-of-trip weather detection is done using two functional templates applied simultaneously. One looks for radially aligned thin lines in DZ, while the other requires that the corresponding SD values are high. The result is an interest image where out-of-trip signals are highlighted. Following the combination of all other interest images, the out-of-trip interest image is subtracted from the combined interest image to selectively suppress evidence for the presence of gust fronts where out-of-trip signals are found. An example of out-of-trip weather detection is shown in Figure 22, using the dual FT shown in Figure 21. FTC is done in polar format so that the functional templates need to be applied only in the vertical orientation (radial alignment). Subsequent frames show the interest image resulting from FTC and the mapping of the interest values from polar to Cartesian format.

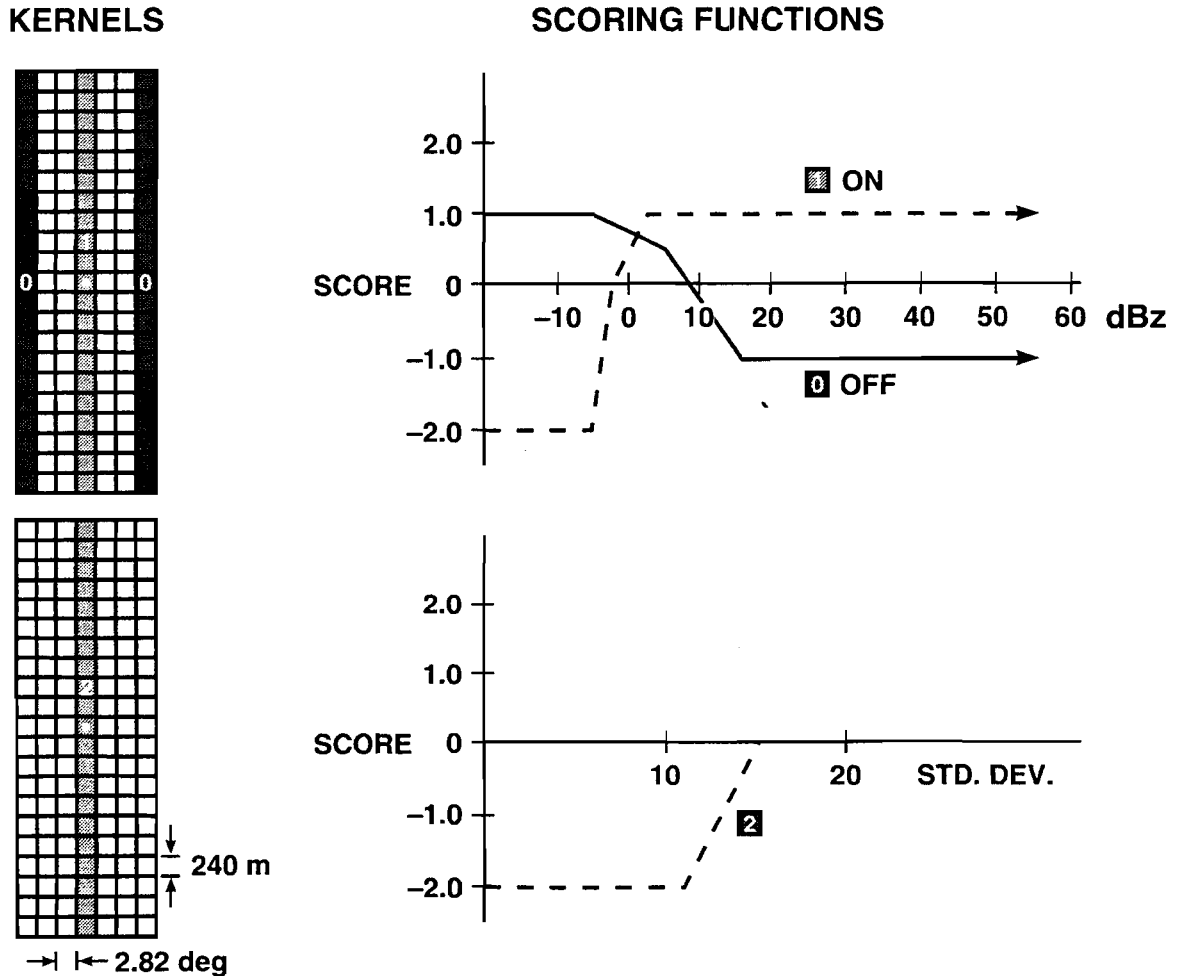


Figure 21. Dual functional template for out-of-trip weather.




Figure 22. Feature detector for out-of-trip weather.

4.5.4 Anticipation

The anticipation interest image provides a mechanism to spatially adjust the detection sensitivity of MIGFA, based on situational context. High anticipation values get averaged with interest values from other feature detectors to increase the likelihood of detection at specific locations. Low anticipation values suppress the likelihood of detection.

The most important use of anticipation is as a replacement for *coasting*. Simply put, coasting is the continued tracking of a target on a radar screen for some time interval after the target disappears (i.e., its signal falls below threshold). Coasting assumes that the loss of a target's signal is not due to a change of behavior (e.g., a change in velocity or perhaps the disappearance of the target). Gust fronts do change behavior, as in cases of collisions of gust fronts. Consequently, blindly coasting a signal after its loss is a potential source of false alarms. Instead of blind coasting, anticipation provides a mechanism for progressively increasing the sensitivity of a detection system, supporting weak evidence that would otherwise be below detection thresholds. In MIGFA, prior history of gust front behavior is used to make a prediction of where the gust front is expected to be in the current scan (prediction tools are described in Section 4.9). These predictions are used to create a band of elevated interest values; typically not so high as to trigger a detection by themselves, but high enough to raise colocated weak signals above threshold. In general, as the length of time a gust front has been tracked increases, the anticipation interest values can also be increased. If absolute coasting is desired, interest values are increased to a level high enough to trigger a detection without any other supporting evidence. Examples of anticipation interest images are shown in Figure 23 and Figure 24.

Other uses of anticipation are to adjust the sensitivity of gust front detections on the basis of contextual knowledge. Some examples are as follows:

1. Many gust fronts are not visible when directly over the radar site, owing to obscuration by intense ground clutter. Even with anticipation of where the front is expected, gust fronts are often lost crossing over the radar site. Absolute coasting over the radar site is accomplished by setting interest values within 2 km of the radar site to nil (i.e., missing value) in all other interest images. Consequently, the anticipation interest image is the only detector allowed to have an opinion of what exists directly over the radar site.
2. Gust front false alarms often occur in thin, wispy bands of low reflectivity stratiform rain. In central Florida at least, gust fronts are seldom associated with the stratiform rain that follows intense storm activity. Under these conditions, the background anticipation interest values are suppressed.
3. In the absence of any precipitation in the radar screen, false alarms are rare in occurrence. Under these conditions, the background anticipation interest values can be safely raised, increasing the likelihood of detecting an incoming gust front generated by a distant storm.

4.6 Combining Evidence

At the same time that the set of feature detectors is selected, a *rule of combination* is also selected to govern the combination of evidence (data fusion). In principle, the rule of combination can be as simple as the averaging of pixel values across all interest images. However, for the set of ASR-9 WSP feature detectors described above, a somewhat more complicated one has been used.

The four interest images generated by TL-DZ, TL-SD, DZ-MOTION, and SD-MOTION are averaged together (missing values are ignored). The resulting averaged interest image and the ANTICIPATION interest image are combined as a weighted average; the ANTICIPATION is given a weight of 0.25 while the average of the first four interest images is given a weight of 0.75. Finally, elements of the OUT-OF-TRIP interest image are multiplied by 0.25 and subtracted from the elements of the weighted average. The resulting image is called the combined INTEREST image.

Figure 23 shows an example ASR-9 WSP scan, the outputs of each feature detector, and the final INTEREST image. In this case, strong evidence for the two fronts is visible in each of the component interest images (except, of course, for the OUT-OF-TRIP image). Clearly, any one of the feature detectors would have provided an adequate detection mechanism. Now consider Figure 24, summarizing the evidence for the presence of the same two gust fronts in a later scan. Detection is becoming more difficult as developing storm cells occlude the fronts. Notice that different parts of the gust fronts are highlighted in different interest images, and that no one feature detector (except ANTICIPATION) sufficiently discriminates the entire gust fronts from background. This example illustrates how evidence derived from multiple feature detectors can be combined such that the weaknesses of one detector can be compensated by others. Alternatively, all feature detectors might generate weak interest values along the entire length of the front. In such a case, the feature detectors can mutually support each other, triggering a detection that would be missed in a detection system based on only one of the feature detectors. No one feature detector is meant to be a perfect, or even a necessarily good discriminator of gust fronts and background. Several weakly discriminating feature detectors, used as a set, can achieve robust performance, depending on how they are combined.




Figure 23. Combining interest: strong evidence. Regions of red pixels are areas where specific feature detectors do not express an opinion regarding the presence of a gust front, deferring to the evidence generated by other feature detectors.




Figure 24. Combining interest: weak evidence.

4.7 Extraction

Algorithms that track gust fronts as entities, such as AGFA, must identify gust fronts, assigning each individual front a unique label that remains constant across time. Gust front statistics, such as propagation speed and location, are computed for the front as a whole. In simple cases, this approach is adequate. But gust fronts are elastic, changing shape through time. And inevitably, complex rules are required to handle labeling, correspondence, and tracking as a single front fragments into disjoint pieces, or as multiple fronts merge or collide. Given the variable nature of gust front behavior, the construction of a fully comprehensive set of rules, correct for all possible circumstances, is a difficult problem.

This problem is bypassed in MIGFA, where the goal of extraction is to identify the set of points (collectively called an *event*) that lie in *any* gust front. Certainly, some chains of points are spatially segregated or have different velocities. For purposes of reporting, such chains can be inferred to belong to separate gust fronts. However, there is no concerted attempt to label or track gust fronts as entities. Instead, individual points are tracked across time; the fact that a point belongs to one gust front or another is irrelevant to processing. Unlike AGFA, which assumes that predicted fronts will stay rigid across time, MIGFA predictions are elastic. The variable velocities of different points along gust fronts are incorporated to make more accurate predictions of what the gust front appearance will be at some time in the future.

Figure 25 illustrates the main steps of extraction discussed in the following sections.




Figure 25. Extraction steps.

4.7.1 Thin Line Smoothing

Thin lines in the combined interest image can be fragmented where gust fronts intersect with out-of-trip weather or are obscured by storm cells. Thin line smoothing is used to bridge the gaps between colinear fragments and to suppress random, unaligned points of high interest values. A “bow-tie” functional template (Figure 26), inspired by the receptive field of the cooperative cell of the Boundary Contour System developed by Grossberg and Mingolla[17] is used as the basis for thin line smoothing. By placing more kernel elements at the ends than at the center, the bow-tie weights the influence of the end regions over that of the center. Consequently, the bow-tie generates high output interest scores for an image point between two colinear high interest segments, even if that point itself has a low input interest value. Also, because of how the scoring functions are designed, the bow-tie filter suppresses and amplifies colinear interest values that are below and above the level of ambiguity (0.5), respectively. As a consequence, the boundaries between gust fronts and background are sharpened, resulting in cleaner shapes for subsequent processing. An example input and output of thin line smoothing is shown in frames COMBINED INTEREST and SMOOTHED INTEREST, respectively, in Figure 25.

In addition to a smoothed interest image, the bow-tie functional template generates a map of orientations. In the ORIENT image, each element is the orientation associated with the highest scoring bow-tie rotated at 10 degree increments from 0 to 170 degrees. Black corresponds with best matches at 0 degrees, while white corresponds with best matches at 170 degrees.

4.7.2 Generating thin line chains

A threshold of 0.5 is applied to the smoothed interest image to create a binary image of candidate fronts. The lengths of resulting elongated shapes are then computed. The elements of binary shapes that are too short (< 6 km. for the ASR-9 WSP) are set to 0. The results of thresholding and pruning are shown in Figure 25. Note that these are the first discriminating thresholds applied during the processing of a scan.

Binary elongated shapes are thinned using an FTC implementation of a modified version of Levaldi’s homotopic thinning [18]. End points are found and marked. The result of these operations is shown in frame MARKED THINNED in Figure 25. End points are shown as white pixels, while internal points are shown as mid-level grey pixels.

Next, the chains of points resulting from thinning are extended along ridges of relatively high interest, using what is essentially a road-following algorithm. At each marked end point, the 8-connected space is examined, looking for the pixel with the maximum interest and whose orientation (found in the map ORIENT) is within 40 degrees of the initial end point. The 40 degree orientation limit was empirically determined to allow sufficient tolerance for curved fronts while prohibiting unrealistic extension of the chains. When the maximum interest score of a new point falls below 0.30 or when no new point has an orientation consistent with the initial end point, extending halts. The process is repeated a second time. Once the new end points have been marked, the chains are

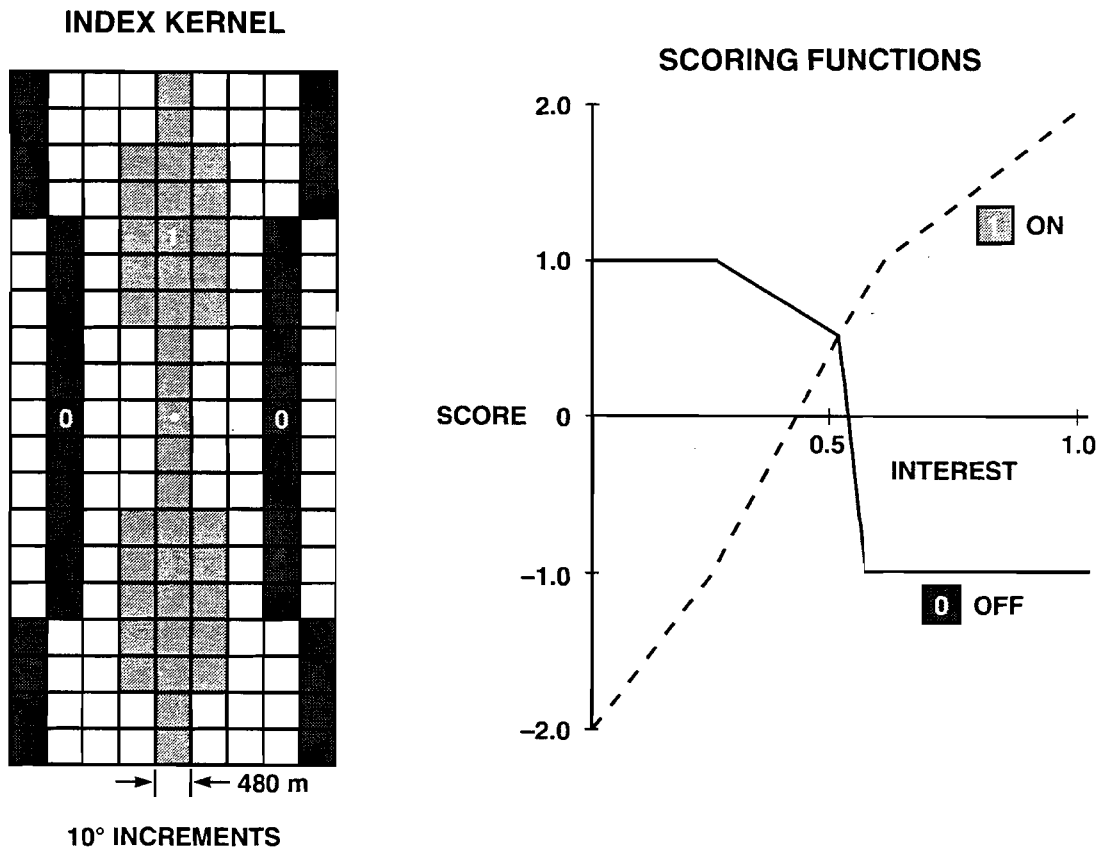


Figure 26. Bow-tie filter for thin line smoothing.

extended as long as the orientation of the extended point is within 11 degrees of that of the newly marked end point. Extending stops when the maximum interest score falls below 0.20.

4.7.3 Finding the strongest chains

After chains have been extended outward from high interest areas, the result may be highly branched and may contain loops. Chains can be assigned scores based on the summed interest value in the smoothed interest image for each point in the chain. Each network of chains can be described as a graph: a set of nodes (junction points) that are linked together by edges. Nodes that have only one edge are called terminal nodes, while nodes with two or more edges are internal nodes. An edge having at least one terminal node as an end point is a terminal edge; otherwise it is an internal edge. Two nodes are neighbors if there exists an edge having the two nodes as end points. One standard application of graph theory is to try to minimize the costs of traversing edges in the process of visiting all nodes, a problem known as the "Traveling Salesman Problem".

In MIGFA, for each disjoint network of chains, the goal is to find the combination of non-looping chain segments with the highest summed interest score (the highest scoring combination is usually, but not always, the longest combination). Consider the simple example shown in Figure 27. The desired combined chain is the one consisting of segments AB, BD, DE, and EF. One way to

reach the goal would be to combinatorically try all subsets of any number of chains and find the subset with the maximum summed interest score. However, with large numbers of nodes and edges, such a combinatoric solution is prohibitively expensive. Instead, the most interesting combination of edges is discovered by iteratively pruning away terminal edges. The basic operation is to find an internal node having two or more terminal edges (e.g., node E in Figure 27). The terminal edge with the highest score (edge EF) is retained in the graph; the remaining adjoining terminal edges are removed (edge EG). If the resulting pruned graph (Step 2) has an internal node with only one internal edge (e.g., node E), then the internal node is removed and the internal and single remaining terminal edges are combined (Step 3). This process continues until no more internal nodes with two or more terminal edges are found. If at this point a network still contains more than one edge, then a loop exists somewhere in the network (as in BCD in Step 3). A loop is broken by removing the internal edge with the lowest summed interest score (CD). Pruning of terminal edges resumes, resulting in the removal of edge BC in Step 5, the fusing of edges BD and DF in Step 6, and the fusing of edges AB and BF in Step 7. The process is repeated until only one edge remains (edge AF in Step 7). If the resulting combined edge has a score greater than a minimum threshold (12.0 interest points [e.g., 12 pixels with interest scores of 1.0 or 24 pixels with interest scores of 0.5, etc.] in the ASR-9 WSP version), then that edge is included in the event for the current scan.

Once the highest scoring combined chain is discovered, its constituent edges are removed from the original graph. The remaining graph structure is examined again (Step 8), looking for another maximally scoring combined chain. Processing continues as long as combined chains with interest scores greater than the threshold are extracted (edge EG with an interest score of 20 is retained, the product of fusing edges BC and BD is rejected). The frame labeled SELECTED CHAINS in Figure 25 shows 2 above threshold combined chains extracted from frame EXTENDED.

4.8 Tracking/Heuristics

As stated above, each point in the extracted event is individually tracked. In order to track a point, the corresponding point in the immediately prior event must be found. Correspondence can be difficult to establish when several gust fronts collide; in such cases, the point in the prior event that is closest to a point in the current extracted point is not necessarily the correct corresponding point. Consequently, a point in the current extracted event corresponds to the closest point in the immediately prior event for which orientation and speed are consistent. If no such point in the prior event is found, then the current point is assumed to correspond to the closest point in the prior event. Once correspondence for a point is established, the point is indexed by creating a pointer linking that point to the corresponding point in the immediately prior event. If the distance between the two corresponding points is too large, or if the distance moved for the two points is inconsistent with prior history, then the point is unindexed (the link is broken). Through the index links, a point can be tracked backwards in time to a first recorded instance. The number of prior events through which a point can be tracked is called its depth. A depth of 0 means that the point is unindexed. Once indexed, each point is assigned the following attributes: coordinates, distance moved, direction moved, depth, Doppler value, interest value, and propagation speed.




Figure 27. Example extraction of most interesting chains from a simple graph structure.

After indexing, each extracted chain is edited:

1. If the direction a single point moves is opposite (approximate difference of 180 degrees) from its neighbors, the direction is reversed.
2. Single chains may be divided into two subchains if a persistent change in velocity or a persistent change in orientation is detected at some point along the chain.
3. Various parameters, such as propagation speed, Doppler value, and direction moved, are smoothed along the length of each chain.
4. Heuristics are applied that, when satisfied, unindex (break the link and set the depth to 0) individual points in a chain. If more than half of any chain's points become unindexed, all points in the chain are unindexed.

Heuristics are used to reduce the number of false alarms, making use of knowledge of how false alarms can be distinguished from real gust fronts. Some examples are as follows:

1. If the direction a point moves is inconsistent with the measured Doppler value, the point is unindexed.
2. If the point is approaching the radar site and moving in the same direction and no faster than the winds measured by anemometers at the radar site, there is no implicit convergence between the gust front and the radar site. Therefore, the point is unindexed.
3. If all points in a chain form a concave arc with each point moving directly towards the radar site, the detected signal is probably an antenna-sidelobe generated artifact and all points are unindexed.

The final stage of tracking is to make a binary decision for each chain of points whether or not to report the chain as a gust front detection. Whether or not particular points of detected chains are included in the report depends upon the summed interest score of the chain and the individual depths of its constituent points. In chains with high summed interest scores (reflecting a higher degree of confidence), points with lower depths can be included. Conversely, chains that have low summed interest scores are less likely to be gust fronts. Consequently, the constituent points are required to accumulate higher depths before being included in the announced gust front detections. In frame INDEXED EVENT in Figure 12, the set of all extracted points are shown. Those points shown have sufficient depth and interest scores to be reported. In frame HISTORY of Figure 12, the reported points are shown in context with previously reported events.

4.9 Prediction

The current extracted event, indexed into the prior history, is used to make predictions of where the points having sufficient depth and interest are likely to be at some time in the future. Given the direction moved, the propagation speed, and the current coordinates of a point, a new coordinate is computed for some time in the future. Gaps can arise between the projected future coordinates of two adjacent gust front points when the orientation and velocity of the adjacent gust front points are not identical. Such gaps between projected points are filled in. An example, showing the reported chains and their expected locations in 10 and 20 minutes in the future are shown in Figure 12 (PREDICTIONS).

The prediction module is also used to generate the next ANTICIPATION interest image. Based on the integrated history, the expected locations of gust front points in the next scan are computed and used as described in Section 4.5.

4.10 Computing Wind Shift and Wind Shear Hazard

The wind shift product is an estimation of the wind velocity expected after the passage of a gust front. Doppler estimates associated with the gust front thin line are used to determine the magnitude of the wind vector behind the gust front. Of course, the Doppler values are an accurate representation of the wind speed only when the winds are nearly radially aligned. In order to obtain the wind speed estimate associated with the most favorable viewing angle, the near-maximum Doppler value found along the reflectivity thin line is identified. The near maximum value is obtained by first sorting the Doppler values coincident with the reflectivity thin line, and then selecting the value corresponding to the 90th percentile of the ordered distribution. Choosing the 90th percentile guards against outliers that may be generated by aircraft or poor quality velocity estimates. The direction of the wind behind a gust front is assumed to be the same as the propagation direction of the front, which in turn is assumed to be normal to the orientation of the thin line signature. Although deviations from this assumption do occur, its validity has been supported by comparisons of gust front propagation direction with post-frontal winds measured by LLWAS anemometers. This approach has been used for TDWR data in situations where data quality cannot support more sophisticated wind shift computation methods [8]. In other words, the wind behind the gust front is defined as the vector V_b with a magnitude equal to the Doppler measured speed behind the front and a direction equal to the direction of propagation.

The wind shift estimate is depicted graphically on the GSD as a purple arrow pointing in the direction of gust front propagation. The base of the arrow is placed 6 km. behind the midpoint of each distinct gust front chain.

The mean wind over the airport is computed using the LLWAS network mean wind algorithm [19]. This algorithm uses temporal smoothing, producing a new time-weighted average as new data arrives every 10 seconds. Statistical outliers are rejected using a chi-squared test. Every sixth iteration, approximately once a minute, the current mean wind value is added to a 2-hour long history. The latest mean wind value is provided for tracking heuristics relying on the detection of implicit convergence to confirm that moving thin lines are indeed gust fronts. The mean wind value is also used in the computation of the wind shear hazard.

The wind shear hazard ΔV_h is the estimate of the change in velocity expected with the passage of a gust front. When this value exceeds 15 knots (7.7 m/s), an alarm is sent to air traffic controllers. Estimates of the winds on both sides of a gust front are needed to compute ΔV_h . The vector V_b described above can be used as an estimate of the winds behind the front. However, the winds ahead of the gust front usually cannot be determined from ASR-9 WSP data. Fortunately, airports receiving ASR-9 WSP systems are also equipped with LLWAS. Anemometer data from this network can be used to derive estimates of the winds ahead of gust fronts approaching the airport.

Figure 28 illustrates how ΔV_h is computed as the difference between the magnitude of the wind behind the front V_b and the magnitude of the component of the vector of the wind ahead of the front V_a that is parallel to V_b . More precisely,

$$\Delta V_h = |V_b| - |V_{a'}|,$$

where $V_{a'}$ is the dot product of V_a and the unit vector having the direction of V_b :

$$V_{a'} = V_a \cdot \frac{V_b}{|V_b|}.$$

As long as an incoming gust front is outside the LLWAS anemometer network, the current mean wind estimate can be used as the value for V_a . However, as a gust front impacts the airport and the anemometer network, the mean wind value no longer reflects the winds ahead of the front. Using the affected current mean wind value to compute ΔV_h will cause an underestimation of ΔV_h during airport impact. This problem is avoided by using the most recent unaffected mean wind estimate as obtained from the stored history for V_a instead of the current mean wind. This is done by estimating the most recent time for which the gust front was outside of a circle circumscribing the anemometer network (the radius and center of this circle are site-adaptable parameters). The mean wind associated with this time is used as V_a . No attempt is made to determine if a front that is outside the LLWAS circle is inbound or outbound. For an outbound front that is beyond the LLWAS circle, the use of the current network mean estimate will obviously lead to incorrect estimates of V_a , but since the wind shear hazard estimate is only examined for alert purposes during impact with the airport, the incorrect value is ignored and therefore has no consequence.

Accuracy of pre-frontal wind estimates may be compromised if two or more gust fronts impact the LLWAS network at nearly the same time. Since both fronts are inside the LLWAS circle, MIGFA will assign unaffected pre-frontal wind estimates to both fronts corresponding to the time when each first crossed the circle. If one front passes through the network a little ahead of the other, the algorithm may incorrectly continue to use the older network mean wind estimate indicating winds prior to the passage of the first front. Improvements to better handle this scenario are currently being investigated. However, given the limited area being affected, this problem should occur infrequently.

4.11 Gust Front Update Task

The need for timely displays of current and forecast gust front positions dictates a display update rate (DISPLAY_RATE) that may actually exceed the algorithm throughput from MIGFA

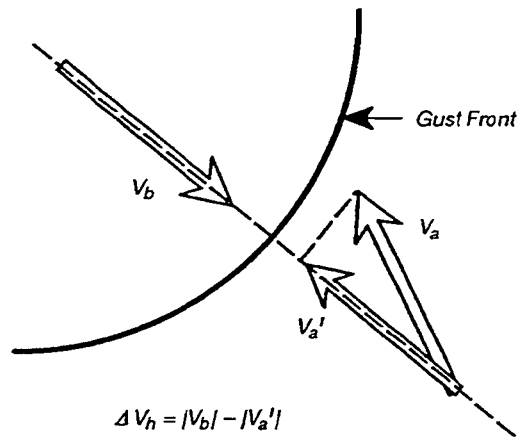


Figure 28. ASR-9 WSP gust front wind shear hazard (ΔV_h) estimation. ΔV_h is computed as the magnitude difference between the winds behind the front V_b and the vector component of the wind ahead of the front V_a that is parallel to V_b , which is labeled V_a' .

(during 1992 tests, the MIGFA processing cycle was two minutes, while the display update rate was one minute). The gust front update task (GFUPDATE) receives gust front detection and forecast data from MIGFA and uses an internal timer to schedule data deliveries to the GSD at the desired display update rate. Forecasts of future gust front location are used in the absence of current MIGFA reports to update the gust front position on the GSD.

Briefly, the algorithm works as follows. On receipt at time T of one or more detection records and associated time series of forecast records, any pending GFUPDATE output is flushed and the internal timer ΔT is reset to 0. The new detection, along with forecast data records corresponding to times $T + 10$ minutes and $T + 20$ minutes, are sent on to the GSD display process. Unless new data arrives from MIGFA to reset the ΔT , each signal from the timer (occurring every DISPLAY_RATE seconds) triggers the transmission of a new packet of data. The generated packet contains a new gust front detection record containing gust front points copied from a forecast record corresponding to the current time ($T + \Delta T$). Wind shift and wind shear information, saved from the previous detection record, are included in the new detection record. New forecast records are created from saved records forecasting gust front locations 10 and 20 minutes from the current time: $T + \Delta T + 10$ and $T + \Delta T + 20$. Note that the time resolution of the forecast record sequence must match DISPLAY_RATE to ensure that appropriate forecast records are available for use in between actual algorithm reports.

5. RESULTS

5.1 Truthing and Scoring

Algorithm performance is scored against human interpretations of the input radar data. Hidden in this statement is the assumption that human interpretations are 100% accurate. As will be seen later, this assumption is not always correct.

The human interpreter has access to both Doppler and reflectivity images for an entire sequence of ASR-9 WSP scans, which can be viewed separately or in sequence as a movie. For each scan, a description of "truth" is stored in a table as a list of coordinates marking the end points of the gust front and an intermittent sampling of points in between. Also stored is the estimated maximum wind shear in the zone of convergence. This scoring exercise is intended to measure MIGFA detection performance and not end-to-end gust front detector capability for the ASR-9 WSP. Consequently, the human interpreter was restricted to including in the truth set only those gust fronts that had some visible signature in ASR-9 WSP imagery, however subtle. Other data sources, such as matching TDWR data and anemometer measurements of winds over the radar site, were used to confirm or deny the existence of gust fronts having an ambiguous appearance in ASR-9 WSP data. In cases where MIGFA detections in ASR-9 WSP data were scored against a human interpreter looking at TDWR data, the same procedures were used to generate TDWR truth tables.

An automatic procedure for scoring computed gust front detections against human generated truth is described in detail in Klinge-Wilson et al. [10]. Briefly, the scoring algorithm draws lines that connect the sequence of coordinates encoding the human estimated limits of a gust front. The lines are then expanded to a 5 km wide region, called in this report a truth box. Computed gust front detections that overlap with some truth box are counted as successful detections. Those that do not overlap are counted as false alarms. A Probability Of Detection (POD) is computed by taking the ratio of successfully detected fronts to the number of fronts identified by the human interpreter. The Probability of a False Alarm (PFA) is the number of detections that fall completely outside of any truth box divided by the total number of detections (false and valid). In addition to the hit-or-miss POD and PFA scores, scoring is also done in terms of percent overlap of algorithm detections and truth boxes. To do this, truth boxes are subdivided along their length into 1 km bins. The Percent of Length Detected (PLD) is then defined as the number of truth bins overlapped by a detection divided by the total number of bins in the truth box. The Percent of False Detection length (PFD) is expressed as the total detection length that falls outside of any truth box divided by the total length of the detection.

One improvement to this method is the use of a "MAYBE" category of truth. Often, gust fronts or parts of gust fronts are weak or only marginally discernible, forming a "gray area" in which the human observer is undecided or uncertain. If an algorithm detects a weak gust front that the human observer believed to be too weak to be detectable, the detection should not be

counted as a false alarm. Conversely, if the algorithm misses a gust front that is too weak to have any operational significance, its POD and PLD should not be affected.

The human analyst identified regions of “definite” truth as well as regions of “MAYBE” truth for each scan of data. Each truthed region is called an event. The following scoring equations treat “maybe” truth categories in accordance with the philosophy described above and were used to score AGFA and MIGFA outputs:

$$\begin{aligned}
 POD &= \frac{N_{ed}}{N_e} \\
 PFA &= \frac{N_f}{N_d} \\
 PLD &= \frac{L_{ed}}{L_e} \\
 PFD &= \frac{L_f}{L_{ed} + L_f}
 \end{aligned}$$

where

N_{ed} = Number of “definite” truth events overlapped by detections

N_e = Number of “definite” truth events

N_f = Number of detections falling outside of both “definite” and “maybe” truth events

N_d = Number of detections

L_{ed} = Total “definite” event length detected

L_e = Total length of “definite” truth events

L_f = Total length of detections falling outside of “definite” or “maybe” truth events

5.2 Comparison of AGFA and MIGFA

Table 2 compares performance of MIGFA against the previously constructed AGFA described in Chapter 3, which uses more conventional methods of signal processing and computer vision. A test set of ASR-9 WSP data collected in Orlando, Florida during field testing in 1991, contains 9 different, moderately strong gust fronts tracked through 15 hours (372 images). A human interpreter looking at the same data detected 280 instances of the 9 gust fronts. The POD and PLD scores indicate that MIGFA performed substantially better than AGFA in detecting gust fronts. The POD evident for MIGFA (88%) is an improvement over that shown for AGFA. But when the rate of missed detections, determined by subtracting the POD scores from 100%, the results are more impressive. While MIGFA missed only 12% of the gust fronts identified by a human interpreter in the data, to a miss rate of over 40% for AGFA. This same relationship can be seen in the relative

TABLE 2
AGFA and MIGFA performance on ASR-9 WSP data as scored against human interpretations.

	Gust Fronts		Gust Front Length	
	POD	PFA	PLD	PFD
Baseline (AGFA)	56.7	4.6	38.9	12.9
MIGFA	88.1	0.6	86.2	33.4

PLD scores. At the same time, the PFA for MIGFA is quite low, more than 7 times smaller than that shown for AGFA. However, the increased PFD (from 12.9% for AGFA to 33.4% for MIGFA) suggests that MIGFA does a worse job of discriminating the extent of individual fronts than does AGFA. In order to better understand why MIGFA was extending fronts beyond what the human interpreter believed appropriate, we examined several cases where the PFD was high. In most of these cases, we found that the extra points that MIGFA included in gust front detections were believable. For an example, Figure 29 shows the truth box for a gust front overlaying a MIGFA-generated detection. The human interpreter was reluctant to include the extreme ends of the front; the ends were nearly radially aligned, had weak reflectivity values, and might represent out-of-trip weather. However, since the extended thin line moved consistent with the center of the front and because the extended thin line had Doppler variance too low to be out-of-trip weather, MIGFA probably gave the more plausible interpretation of the scene.

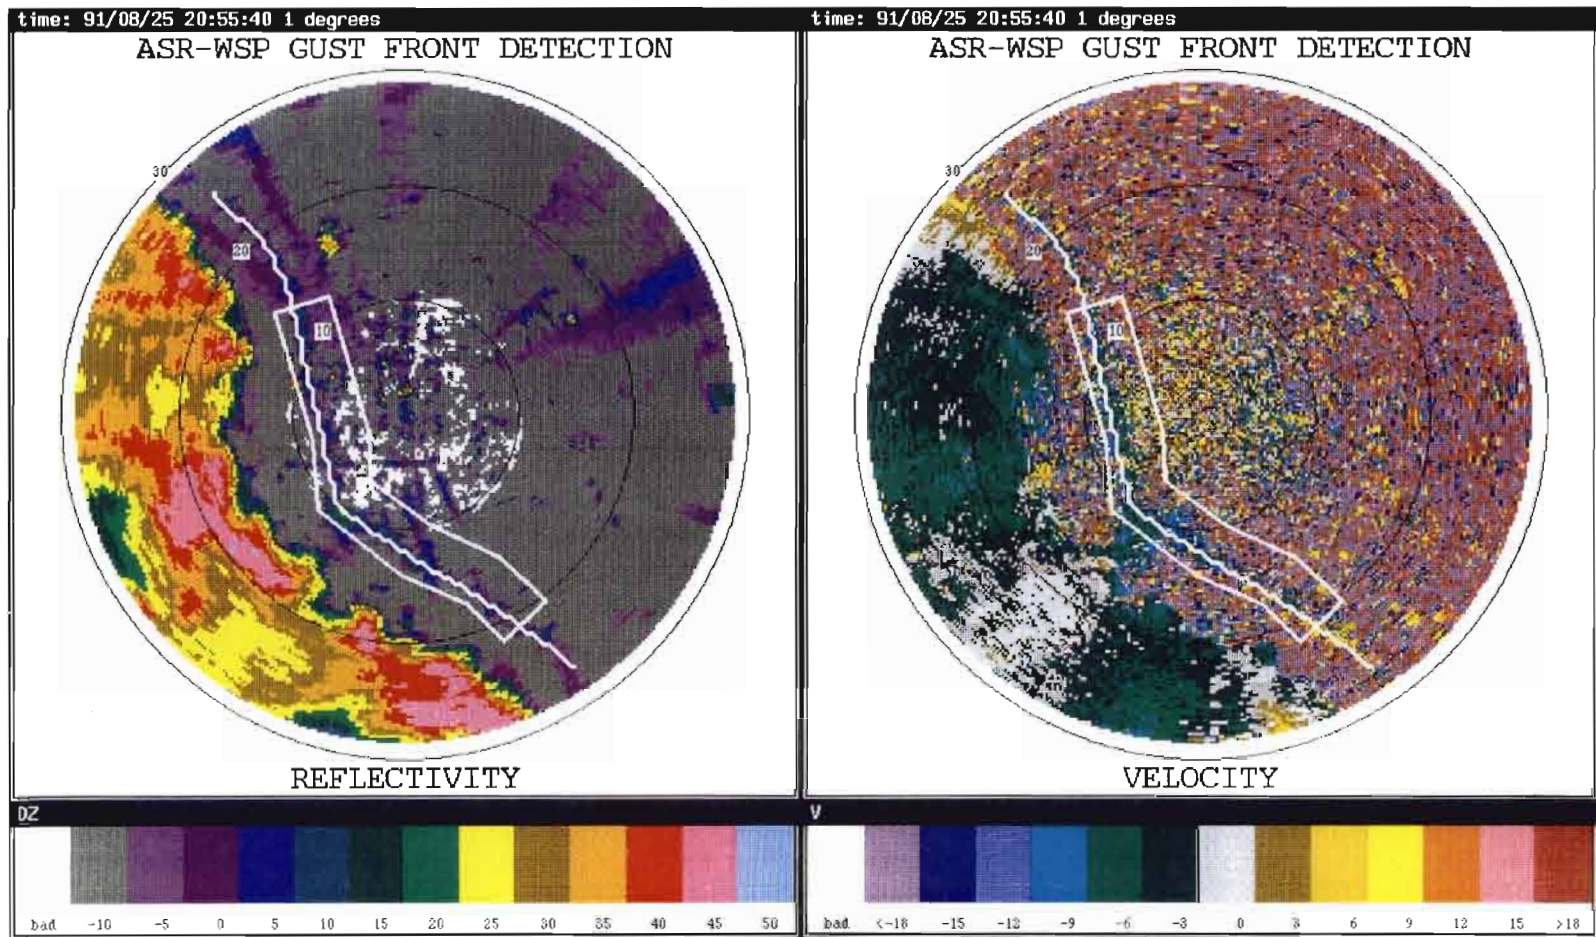


Figure 29. Example of MIGFA gust front detection (white line) extending beyond region identified by a human interpreter (white box). The entire detection was determined to be valid based on prior track history and supporting evidence from corresponding TDWR images. Range rings are in km, reflectivity (DZ) in dBZ, and radial velocity (V) in m/s.

To substantiate these anecdotal observations, gust fronts detected in ASR-9 WSP data by MIGFA and AGFA were scored against human interpretations of TDWR data collected at the same time. These scores, shown in Table 3, indicate that while the general trend of the first 3 figures of merit shown in Table 2 are substantiated, the PFD for MIGFA (6.4%) is approximately the same as for AGFA (4.2%). We assume that the TDWR “truth” is more accurate than the ASR-9 WSP truth, given the improved quality of gust front signatures in TDWR data. Therefore, the drop of PFD that is evident when the results are rescored against TDWR truth instead of ASR-9 WSP truth suggests that the MIGFA results were more accurate than the ASR-9 WSP truth. To better illustrate this point, the difference between the PFDs as scored against ASR-9 WSP and TDWR truth crudely approximates the percentage of detected gust front points missed by the human interpreter. For MIGFA, the sum of this difference ($33 - 6 = 27\%$) and the PLD scored against ASR-9 WSP truth (86%) is 113%, better than the human by 13%. For AGFA, the comparable result is $(13 - 4 + 39) = 48\%$. Given matching ASR-9 WSP data as input, MIGFA appears to be able to discriminate the marginally visible ends of gust fronts better than the human interpreter.

TABLE 3
AGFA and MIGFA performance on ASR-9 WSP data as scored against human interpretations of matching TDWR data.

	Gust Fronts		Gust Front Length	
	POD	PFA	PLD	PFD
Baseline (AGFA)	42.6	3.2	21.0	4.2
MIGFA	75.1	0.0	58.7	6.4

5.3 Results of 1992 Field Testing

MIGFA was installed at the ASR-9 WSP site at Orlando International Airport in the spring of 1992 and was part of a formal operational test from 8 July to 20 September. During this time, gust front detections and predictions were relayed to air traffic controllers for use in planning air traffic. During July, several problems and algorithm deficiencies surfaced. The most significant problem was that the estimate of winds behind the gust front and predictions of eventual wind shifts and wind shears were highly erratic due to noise in the Doppler image. Based on this and other findings, several fixes and enhancements were added throughout July. Because of these fixes and enhancements, MIGFA output collected prior to 20 July have been excluded from our evaluation.

During the period from 20 July to 20 September, 46 weather events were declared by MIGFA to be gust fronts. Forty of the events were verified from logs, LLWAS measurements, TDWR data or ASR-9 data as actual gust fronts. Trackings of actual gust fronts tended to persist for long periods of time, often up to over an hour. In the same interval, six distinct weather events were believed to be false alarms. In all cases, the false detections were short lived with three of them reported for only a single scan and the longest duration being 10 scans. False detections occurred in only 19 scans out of over 14000 scans processed.

An alternative score is based on the probability of detecting and providing advance warning on the most operationally significant gust fronts, those that eventually triggered actual LLWAS alerts with convergent wind shears greater than 15 knots (7.7 m/s). During the period from 20 July to 20 September, 14 convergent wind shears of greater than 15 knots were recorded on the anemometer network at the airport. Two of these wind shears were the result of convective circulations beneath storm cells directly over the airport. The cause of a third wind shear is unknown, but was probably due to a microburst that was reported immediately south of the airport just as the wind shear was recorded. In none of these 3 cases could human interpreters find evidence of gust fronts in the ASR-9 data.

Of the 11 remaining verified gust fronts, MIGFA correctly tracked eight at least up to (but not always over) the airport. In these cases, air traffic controllers were given initial warnings from 18 to 79 minutes prior to gust front arrival. Of the three missed gust fronts, one was a gust front that was occluded by fast moving storm cells following the gust front. The second missed gust front had a very weak, fragmented thin line signature that was missed both by MIGFA and by the operators at the radar site who were logging weather and system activity. The third missed case was a gust front that had just recently been generated by a large microburst only 5 km away from the runways. Being so young, the gust front had not yet developed a thin line signature. Human interpreters studying the data offline could find no evidence for the gust front in ASR-9 WSP data, but could see a small zone of convergence without a corresponding thin line signature in TDWR data. In summary, 8/11 or 73% of gust fronts impacting airport operations with wind shears greater than 15 knots were correctly detected and tracked by MIGFA up to the airport. For comparison, human operators at the radar site logged 9/11 or 82% of the same gust fronts.

False alarms that are reported to be approaching the airport could also impact airport operations. If a false alarm were trusted, causing inappropriate changes in airport operations planning, the resulting delays could be just as bad as when a gust front is missed. During the test period three incoming events, covering a combined time period of 24 minutes (12 scans), were scored as false alarms. All three were probably the result of thin lines of stratiform rain. In each case, tracking was dropped while the estimated time of arrival at the airport was more than 40 minutes in the future. Consequently, none of these incoming false alarms would have influenced airport operations planning. MIGFA results were scored in detail against scan-by-scan truth generated by a human interpreter looking at ASR-9 WSP data for the period from 1 August to 20 September. As with offline testing, performance was scored against human interpretations of the same data. Table 4 shows the performance statistics for the test period. In general, the results substantiate

TABLE 4

MIGFA results on ASR-9 WSP data for the period 1 August to 20 September, 1992 in Orlando, Florida. Results are scored against human interpretations of the same ASR-9 WSP data.

	Gust Fronts		Gust Front Length	
	POD	PFA	PLD	PFD
MIGFA	75.4	1.8	80.8	21.1

the off-line test results. The POD (75%) and PLD (81%) are somewhat lower than those shown for the 1991 offline test set (Tables 2 and 3). But, this small decrease is not surprising given that MIGFA has been tested repeatedly and tuned to work well on the selected offline data set.

Most of the missed detections were due to two problems. First, several gust fronts had reflectivity values at or below the sensitivity limits of the ASR-9. Of course, those fronts with reflectivity values below the ASR-9 limits were not detected by either MIGFA or the human interpreter. However, there were a few cases of marginal contrast in which the human could detect a gust front, but for which MIGFA never accumulated enough confidence to declare a detection. However, note that the human interpreter had the opportunity to examine the sequence of radar images repeatedly and could use information from scans late in the sequence to confirm or deny the existence of the gust front in early scans. Not much can be done to overcome the sensitivity limits of the ASR-9. In most (but not all) cases, these gust fronts with marginal reflectivity levels were associated with weak wind shears. Since these weak fronts had minimal impact on airport operations, a failure to detect them should not be a significant liability.

The second problem was one of obscuration. In several cases, storm cells or out-of-trip weather were extensive enough to hide or fragment the thin line signatures so that some gust fronts were detected late, dropped early, or sometimes missed altogether. To improve performance in these cases, changes to the ANTICIPATION feature detector have been proposed. The interest values of anticipated gust front locations would be increased wherever potential obscuration is detected. By raising interest values high enough within storm cells, detections could be triggered without supporting evidence from other sources (i.e., coasted).

The high PFD (21.1%) is almost entirely because MIGFA extended gust fronts beyond the ends delimited by the human interpreter. But with the use of anticipation based on prior tracking data, MIGFA was able to extend the detected gust front length through areas where the signatures were ambiguous in appearance. As was seen with the offline testing described earlier, a case-by-case analysis indicates that many of these extensions are in fact believable and inappropriately scored

as false lengths. Rescoring the results against TDWR will be necessary to improve the accuracy of the PFD score.

5.4 Accuracy of Wind Estimates

Wind estimates at the beginning of the test period were highly erratic, at one point reporting 50 knot winds for a gust front that actually had winds behind the front of less than 10 knots. The consequence was that many incorrect wind shifts and several false wind shear hazard alerts were indicated early in the test period. The problem was discovered to be a side effect of the high variance of Doppler values. As discussed in Section 4.10, winds behind the front were estimated using the 90th percentile Doppler value within the gust front. However, gaps in thin line signatures that were filled in by MIGFA contained highly variable Doppler values, some of which could be quite extreme. In such cases, the 90th percentile Doppler value would likely be a noise value and not an accurate measurement of the winds within the front. The problem was fixed late in July by temporally averaging Doppler values.

Time series of LLWAS network mean winds recorded during 19 gust front events that occurred after the above problem had been corrected were examined in order to validate wind shift forecasts produced by MIGFA. This set included only gust fronts that were correctly detected crossing over the radar site, allowing direct comparisons of MIGFA and LLWAS estimates. Complex events that appeared to consist of 2 or more merged gust fronts, each with distinct velocities were also removed from the sample set. For each gust front in the remaining sample set, MIGFA-generated estimates of wind velocity were averaged over a period of twenty minutes preceding passage overhead. These estimates were compared with LLWAS network mean wind data averaged over a twenty minute period beginning immediately after the wind shift was evident in the LLWAS network mean data. Table 5 lists the results of these comparisons.

The results show that wind direction forecasts produced by the ASR-9 WSP gust front algorithm are in generally good agreement with measurements made by the LLWAS anemometers, with a mean difference of -8.9 degrees ± 4.34 standard error of the mean. Some of the differences can be attributed to real changes in gust front direction between the pre-frontal and post-frontal measurement periods. However, four cases (07/30, 08/12, 08/31, 09/15) had differences in excess of 30 degrees, which could not be attributed to changes in gust front direction alone.

The most common cause for the larger discrepancies was the use of a single vector to represent the movement of a complex frontal boundary. Some "single" gust fronts were actually comprised of two or more merged outflows sharing a common leading edge. Such merged fronts often showed noticeable bends marking the transition from one outflow regime to another, each having different outflow strength and propagation characteristics. In many of these cases, the average propagation vector assigned by MIGFA to the detected front was not aligned with the movement of the portion of the front that passed over the airport. Similarly, in young gust fronts that have a relatively small radius of curvature, the outflow boundary tends to expand radially from the storm source.

TABLE 5
Observed vs. Forecast Gust Front Wind Shift

		LLWAS Mean Wind		ASR-9 WSP Forecast		Difference	
Date	Time (GMT)*	Speed (m/s)	Dir	Speed (m/s)	Dir	Speed (m/s)	Dir
07/30/92	1945	9.4	350	9.2	315	-0.2	-35
07/30/92	2042	6.3	184	4.3	178	-2.0	- 6
07/31/92	1824	15.6	296	10.4	297	-5.2	+ 1
08/02/92	2127	11.3	295	8.6	279	-2.7	-16
08/03/92	2110	8.3	048	5.4	057	-2.9	9
08/04/92	1945	10.9	324	6.9	311	-4.0	-13
08/05/92	1826	8.3	324	6.8	305	-1.5	-19
08/06/92	2041	18.5	153	9.7	162	-8.8	9
08/08/92	1857	14.0	096	7.5	099	-6.5	- 3
08/09/92	2020	9.8	119	7.5	136	-2.3	17
08/11/92	2118	18.2	137	9.9	146	-8.3	9
08/12/92	1950	17.2	282	9.6	245	-7.6	-37
08/14/92	1932	14.4	284	10.1	304	-4.3	20
08/15/92	1930	13.9	274	7.8	273	-6.1	- 1
08/17/92	2312	11.0	111	5.9	091	-5.1	-20
08/28/92	1638	13.7	287	10.9	293	-2.8	5
08/31/92	2025	7.5	104	4.3	073	-3.2	-31
09/15/92	2016	10.5	041	5.0	003	-5.5	-38
09/19/92	2100	10.0	119	5.2	098	-4.8	-21

* Time given in table corresponds to time at which gust front was directly over the radar.

The propagation vector estimated at the midpoint of the gust front may have a markedly different orientation from the part of the gust front heading for the airport.

Forecast gust front wind speeds were in all cases lower than those measured by the LLWAS anemometer network, with decreases ranging from 0.2 m/s to 8.8 m/s. The mean decrease (-4.4 m/s \pm 0.56 standard error of the mean) was significant with a confidence interval larger than 0.999. Some of the underestimation may have been due to dampening of the sharper wind speed increases because of temporal smoothing of the Doppler values. This smoothing was implemented to eliminate erratic estimates of wind shifts that caused several false alerts early in the test period. Other possible causes include poor Doppler viewing geometry, and winds within a gust front thin line echo (where reliable Doppler values are available) having speeds less than those a few kilometers behind the front. A correction factor based on empirical or statistical relationships between velocity estimates in the thin line echo region and wind speeds behind the gust front is being considered.

The accuracy of MIGFA-generated wind speed gain (ΔV_h) reports was assessed by comparing these estimates against LLWAS ΔV_h values that resulted in wind shear alerts (i.e., wind speed gains greater than 15 knots or 7.7 m/s) between 20 July and 20 September. Since the majority of Orlando gust fronts were weak, the number of alerts was limited. Of eleven such gust fronts, only nine were detected and tracked during approach. Of these nine, 2 detections were dropped about 10 minutes prior to crossover, one because of a system crash. Only seven were tracked over the radar site, allowing direct comparisons with LLWAS alerts. In one case, the LLWAS network was inoperable.

Table 6 compares the maximum ΔV_h reported by LLWAS and ASR-9 WSP for these six gust fronts. These results indicate that ASR-9 WSP gust front gain estimates are reasonable. ASR-9 WSP ΔV_h estimates were within 3 m/s of LLWAS values for each of the six cases. The mean difference was -1.1 m/s \pm 0.77 standard error of the mean, suggesting that the MIGFA and LLWAS values were not significantly different.

The apparent good agreement is actually somewhat surprising, since the ΔV_h estimate depends in part on the wind shift estimate, which was found to be consistently underestimated; MIGFA shear estimates should show the same pattern. One possible reason for the good agreement may be the small sample size (6) of wind shear cases. Another more likely reason has to do with the way in which the LLWAS ΔV_h reports are produced: The ΔV_h reported by LLWAS is the along-runway component of the ΔV_h and is therefore likely to be smaller than the ΔV_h measured between any two anemometer station pairs. The amount of reduction due to taking the along-runway component of the ΔV_h for the LLWAS alerts may be comparable to the biases in MIGFA cross-front ΔV_h reports caused by using the underestimated winds behind the front.

Although not borne out by these results, MIGFA might produce wind shear alerts on some fronts that LLWAS would reject, even in cases where both systems report identical above-threshold ΔV_h values. The reason for this is that the LLWAS algorithm actually checks the velocity gradient between anemometer stations and does not report an alert if the shear (velocity difference divided by the distance between the two stations) fails to exceed a threshold (0.001 s⁻¹). MIGFA, on

the other hand, does not have enough information available to distinguish between weak or strong velocity gradients across gust fronts (see Figure 30), and cannot downgrade an alert on this basis.

TABLE 6
Observed (LLWAS) vs. Reported (ASR-9 WSP) Gust Front ΔV_h

LLWAS ALERT				ASR-9 WSP Report
Date	Start Time (GMT)	End Time (GMT)	ΔV_h (m/s)	ΔV_h (m/s)
07/30/92	20:28:21	20:48:11	12.0	13.2
07/31/92	18:18:14	18:21:14	8.9	9.4
08/02/92	21:27:28	21:30:38	9.7	7.2
08/03/92	21:04:17	21:15:27	10.3	10.0
08/06/92	20:40:48	20:41:38	10.0	7.5
08/14/92	19:32:30	19:38:10	10.3	7.6

5.5 FAA Survey of Air Traffic Controllers

The FAA Technical Center Engineering, Integration, and Operational Evaluation Service conducted a survey of 24 Orlando air traffic control personnel regarding the accuracy of ASR-9 WSP wind shear products, how well the products were displayed, and how useful the products were in runway management [20]. In particular, air traffic controllers were asked to rate the accuracy of the gust front information displayed on the GSD according to the following 5 point scale: very poor (-2), poor (-1), fair (0), good (1), and very good (2). The mean rating fell between fair and good at 0.69 with a standard deviation of 0.76. Based on these results, along with comments and answers to follow up questioning, the report concluded:

The general feeling is that the system still produces a slightly high amount of false alarms, especially gust fronts.

The general finding that air traffic controllers thought the system was producing too many gust front false alarms stands in marked contrast to the low false detection rate presented in this report. Out of 14000 scans processed in the period from 20 July to 20 September, only 19 scans contained false detections. Of these 19 scans, only 12 contained false detections with an indicated motion towards the airport. In all cases, the alarm was dropped more than 40 minutes away from




Figure 30. Example radial velocity profiles across three fronts having different cross-front gradients but the same end-to-end velocity change. In MIGFA, ΔV_h would be reported as 8 m/s for each front. In contrast, the ΔV_h computed using the standard definition (change of velocity over 1 km) would be quite different for the 3 cases.

the airport. Out of 19 detected gust fronts approaching the radar site in the interval from 30 July to 20 September, the wind shift estimates were consistently underestimated. Of the 6 gust fronts that did generate wind shears greater than 15 knots, the average wind shear values reported by MIGFA were consistent with wind shears later measured by LLWAS.

The reason that air traffic controllers complained about the false alarm rate is likely due to the timing of the survey. According to the survey report, air traffic controllers filled out the questionnaires in a seven day period sometime in August. Wind shear values generated by MIGFA were indeed highly erratic during the first three weeks of the test period (8-29 July). If air traffic controllers were not informed of the fix, their assessment would understandably have been biased by the early performance. Consider the following recommendation from the survey report:

The gust front algorithm needs to be refined, especially the gust front prediction feature; the system is still generating too many false alarms.

Given that the “gust front prediction feature” is singled out as needing more work, “false alarms” probably refers not to incorrect gust front detections, but to incorrect predictions, which would include the wind shift and wind shear products. Comments from individual air traffic controllers quoted in the survey report also single out gust front prediction as needing improvement. Despite subsequent good performance displayed by MIGFA during the first weeks of August, there might not have been enough good performance to overcome initial bad opinions. The earlier in August the survey was conducted, the greater the likelihood of this explanation.

Air traffic controllers might also report a high false alarm rate if they expected the gust front algorithm to display only those gust fronts that were operationally significant. For example, the reporting of a real gust front that weakens and dies before reaching the airport might be perceived as a false alarm. Or, when a weakening gust front crosses over the airport, the impact on winds at the airport may be slight. This too might be perceived as a false alarm. If air traffic controllers do not want to be notified about weaker gust fronts, then a wind shear threshold might be used to automatically remove reports of weak gust fronts from the display.

6. EVALUATION

As we have shown with direct comparisons using the same input ASR-9 WSP data, MIGFA provides a substantial improvement over AGFA in detection performance. We also provide indirect evidence suggesting that, given the same input data, MIGFA may be nearly as good as human interpreters.

However, the absolute reported probabilities of detection in this dataset (88% when scored against ASR-9 truth, 75% when scored against TDWR truth) are potentially misleading and should be regarded with caution. The data set used for comparison testing is relatively small and from only one season at one site. The offline test is likely to be too small to contain a good representative sampling of gust fronts, but does provide a reasonable basis for comparing MIGFA against the older algorithm AGFA.

On the other hand, remember that the purpose of the expanded, operational test has been to learn about the phenomenology of gust fronts as well as about the limitations and properties of MIGFA. Bear in mind that we are comparing the performance of an algorithm that has only had the benefit of a single test season against an approach that has been under development for nearly 10 years and has been operationally tested for several years. Several adjustments and enhancements, based on experience gained during the 1992 operational test period, have since been added to MIGFA.

6.1 MIGFA Versus Human Performance

As discussed in Chapter 5, MIGFA appears to display higher detection performance than a human interpreter looking at same ASR-9 WSP data. One reason for this is that the human was constrained to delimit gust fronts only where there was an observable signature. Through the use of anticipation, MIGFA is able to follow a detection through areas where only very weak signatures are evident. Also, remember that during the extraction process, gust front chains are extended from thin lines of strong interest values outward along the thin line to less interesting points. Even with no other evidence, anticipation can be sufficiently high for the extension algorithm to follow. Consequently, if tracking history for a gust front exists and anticipation has accumulated, MIGFA can track the gust front through regions where there is no observable signature at all.

How much of the PFD score for the 1992 operational test period is accounted for by such situations is unknown without thoroughly verifying the existence of gust fronts using other data sources, such as the TDWR. However, several gust fronts were individually verified using anemometer data or by examining individual TDWR scan images. Given the pattern of results in the offline comparisons of MIGFA with AGFA and examples such as the one illustrated in Figure 29 in Section 5.2, some portion of the PFD reflects correct detections that the human observer did not include in the truth.

6.2 Failure modes

Some of the failure modes identified during the 1992 operational test period are described in the following subsections.

6.2.1 Weak signals

In many gust fronts, especially weak ones, the reflectivity levels are simply too low to be detectable with the ASR-9. Figure 31 compares matching TDWR and ASR-9 WSP data collected at Orlando at a time when several weak fronts existed. While at least four fronts can be readily identified in the TDWR images, only two of the fronts (one to the north and one to the southeast) can be seen with any degree of confidence in the ASR-9 imagery. Most of these fronts are too weak to have significant impact on airport operations. Consequently, failing to detect such weak gust fronts is not considered to be a problem.

However, even moderately strong gust fronts can be invisible in ASR-9 imagery. Consider a particularly frustrating case recorded at Orlando on July 10, 1992 and shown in Figure 32. The sequence of four ASR-9 WSP reflectivity images show a gust front at different times. The first and second are of the gust front approaching the airport when it was 7 and 2 km east of the airport, respectively. Even knowing where the gust front is, having verified its existence in TDWR data, it is difficult to detect in these first two images. The third and fourth images are of the same front at later times after having crossed over the radar site. The gust front is much more visible, having reflectivity levels well above background. MIGFA picked up the detection at the third frame, only after the airport was visited by a 110 degree wind shift and a ΔV_h of 8 m/s.

Several similar cases of "invisible" inbound gust fronts, for the most part associated with weak sea breeze fronts, occurred during the month of July. The causes of this observability phenomenon are unknown, but do not appear to be related to radar hardware malfunctions. One possible explanation is that outbound gust fronts are more easily detectable due to the radar beam curvature (anomalous propagation) in the shallow temperature inversion at the radar site produced in the wake of a gust front or sea breeze front passage. The slight downward bending of the beam may produce a more favorable centering of the beam on the shallow low-altitude layer of scatterers responsible for generating the thin line signature.

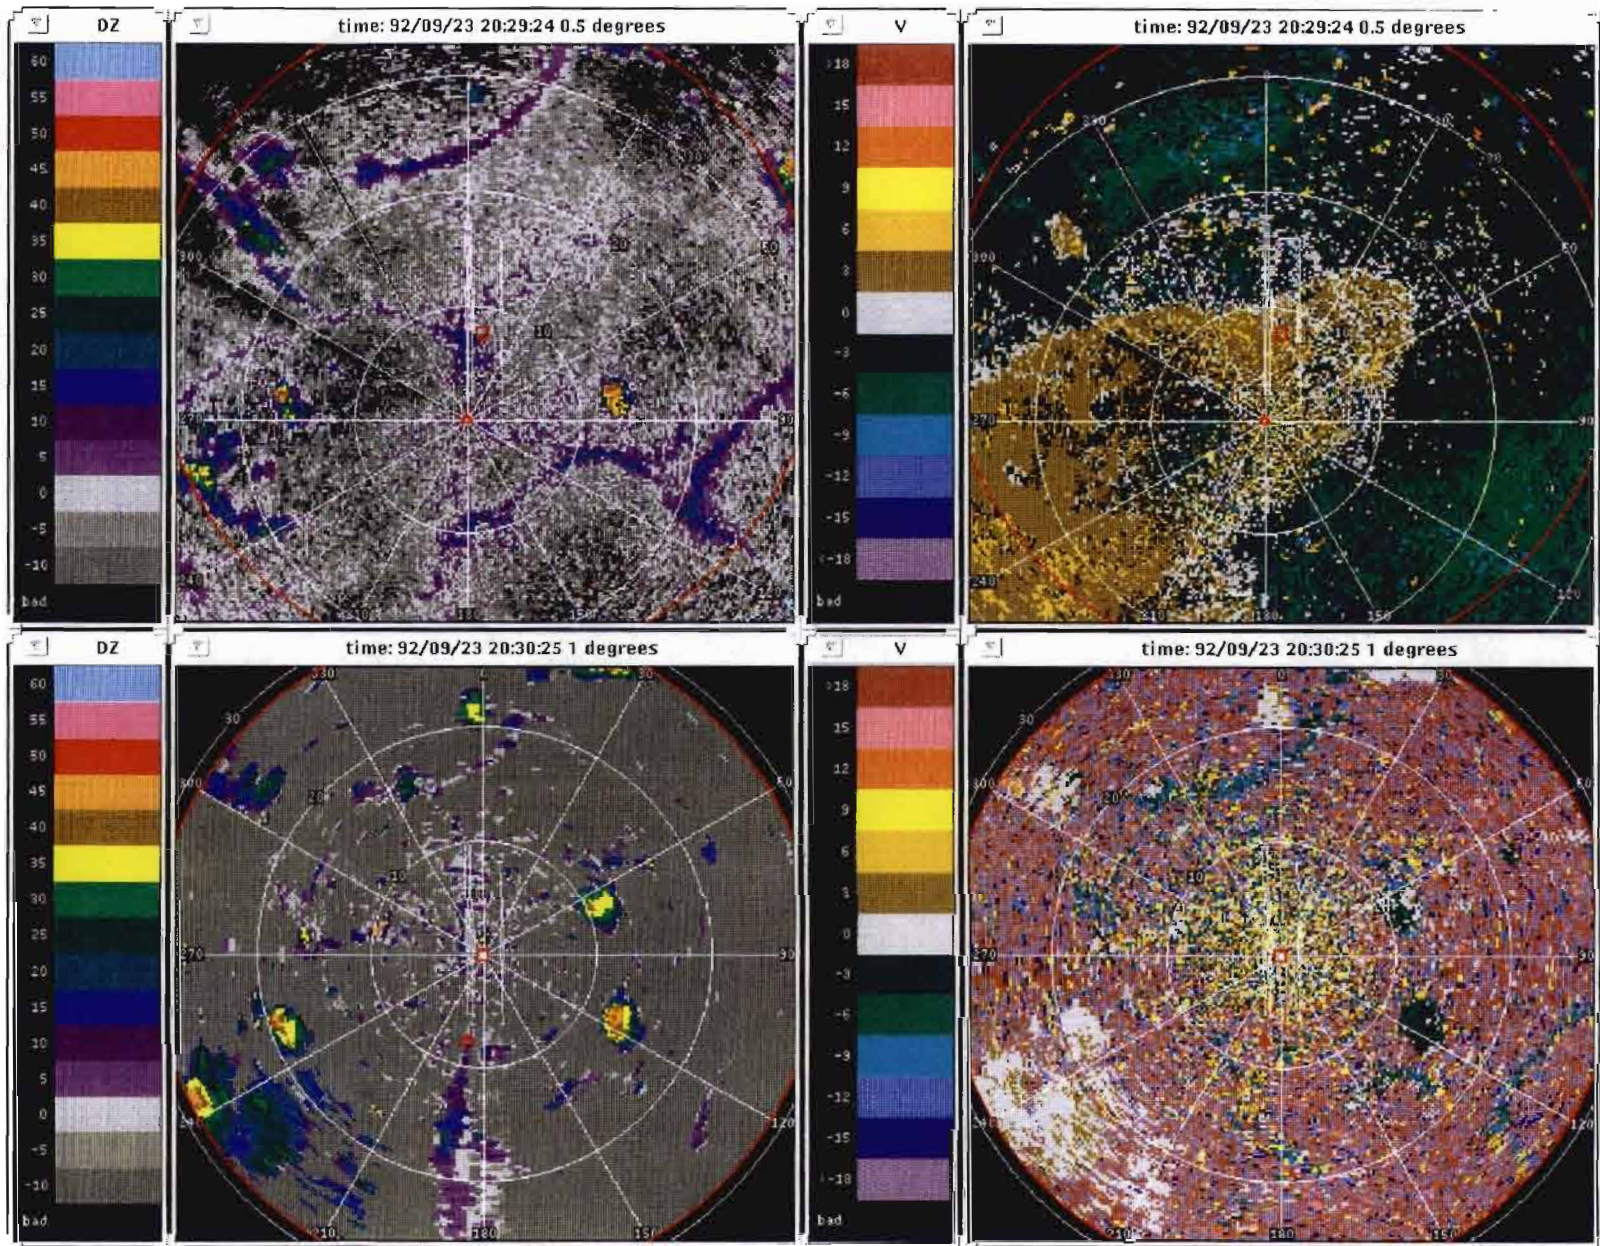


Figure 31. Comparison of gust front visibility in matching TDWR and ASR-9 WSP images. The upper two images are TDWR reflectivity (left) and radial velocity (right). The lower two images are the corresponding ASR-9 WSP images. Range rings are in km, reflectivity (DZ) in dBZ, and radial velocity (V) in m/s.

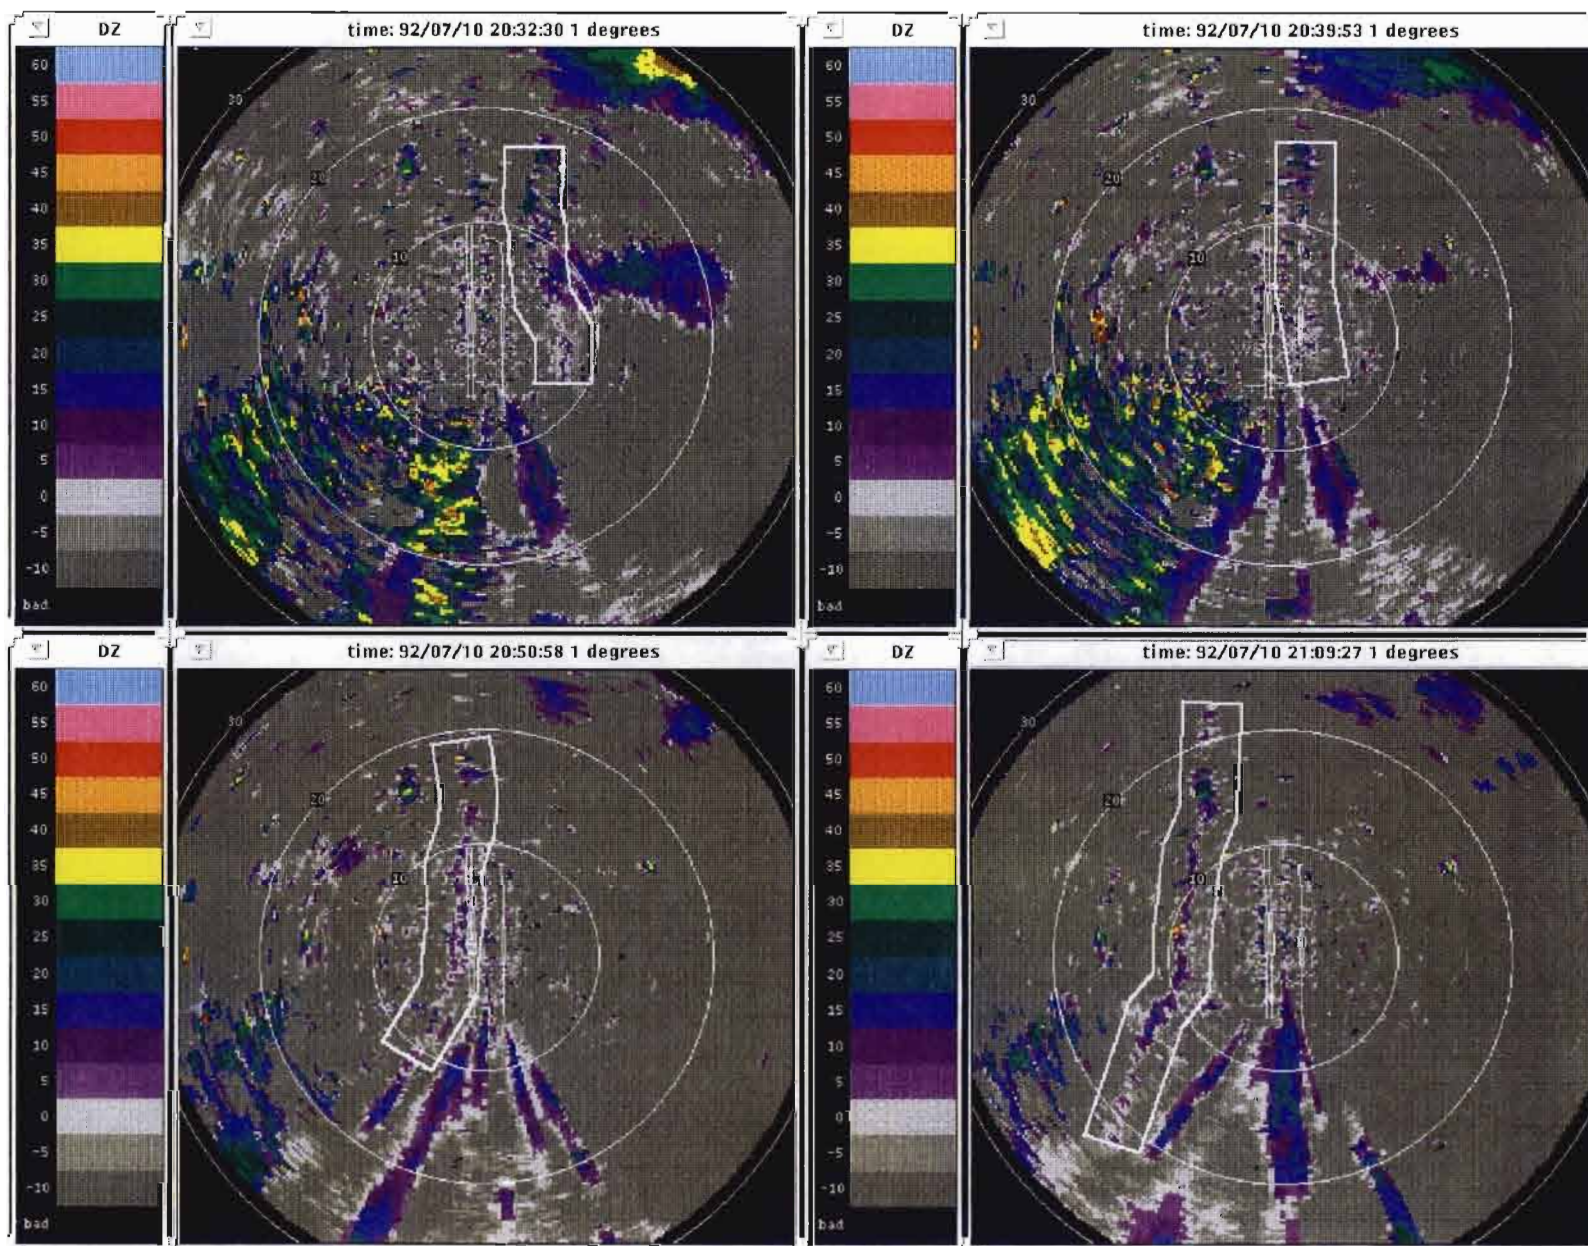


Figure 92. ASR-9 WSP reflectivity image sequence illustrating contrast between nearly invisible inbound gust front signatures and more distinct outbound gust front signatures. Progression is from left to right, top to bottom. White boxes show locations of the front, which passes over the radar while moving from east to west. Range rings are in km, reflectivity (DZ) in dBZ.

6.2.2 Lack of reliable convergence signatures

As discussed in Section 2.3, convergence signatures are unreliable indicators of gust fronts in ASR-9 Doppler imagery. Consequently, the ASR-9 version of MIGFA must rely heavily upon the thin line signature of gust fronts in reflectivity data. However, not all thin lines are associated with gust fronts. Figure 33 shows an elongated, low reflectivity storm echo associated with an extensive area of stratiform rain that is moving with the ambient wind from the SW to the NE. Since the reflectivity levels of light rain echoes overlap with the range of reflectivity levels exhibited by gust fronts, high interest values are produced by the thin line feature detectors. Since in this particular case the stratiform rain was moving in a direction perpendicular to the orientation of the thin line, both motion detectors also generated high interest scores. In most cases, thin line features associated with stratiform rain are transient and can be discriminated from real gust fronts on the basis of persistence through time. In this particular case however, the thin line was tracked long enough for the algorithm to briefly announce it to be a gust front.

Another false alarm encountered during the 1992 operational test period is shown in Figure 34. The strong thin line evident to the south of the airport moved from west to east very slowly, at about 3 m/s. This behavior is consistent with numerous sea breezes that crossed central Florida during the summer. However, this particular reflectivity thin line was not evident in the low elevation tilts used by TDWR for gust front detection. Because of its fan-beam design, ASR-9 imagery includes signal contributions from high as well as low altitude features. Consequently, anything that shows up using ASR-9 but does not show up using TDWR is likely to be occurring at high altitudes. We have no idea what this particular event actually is.

If either of these weather patterns had been more directly approaching the radar site, the false alarms might have been rejected. The ambient winds over the radar site, measured by the LLWAS anemometer network, are used by MIGFA to confirm that a convergent wind shear exists somewhere between an incoming candidate gust front and the radar site. However, this test has only limited reliability, since there may be significant differences between the ambient winds and those in the vicinity of the detected feature.

For example, one incoming false detection, which persisted for 10 scans, was accompanied by a ΔV_h estimate of between 8 and 11 m/s. The false detection was triggered on an area of elongated, low-reflectivity storm echoes that developed 15 km west of the airport in the wake of a severe weather passage. Winds at the airport were gusty and direction was variable due to the influence of several storm cells over and around the airport. In this case, extrapolating the airport winds over a distance of 15 km to generate a convergence measure for the detected feature was inappropriate. Note that outside of the airport vicinity, the convergence estimate is only used for purposes of rejecting detections that do not exhibit some degree of convergence. The ΔV_h estimate used for declaring wind shear alerts becomes more trustworthy as the front moves closer to the airport (and the LLWAS network). Only after the detected front has impacted the airport is the ΔV_h value examined for possible alert generation.

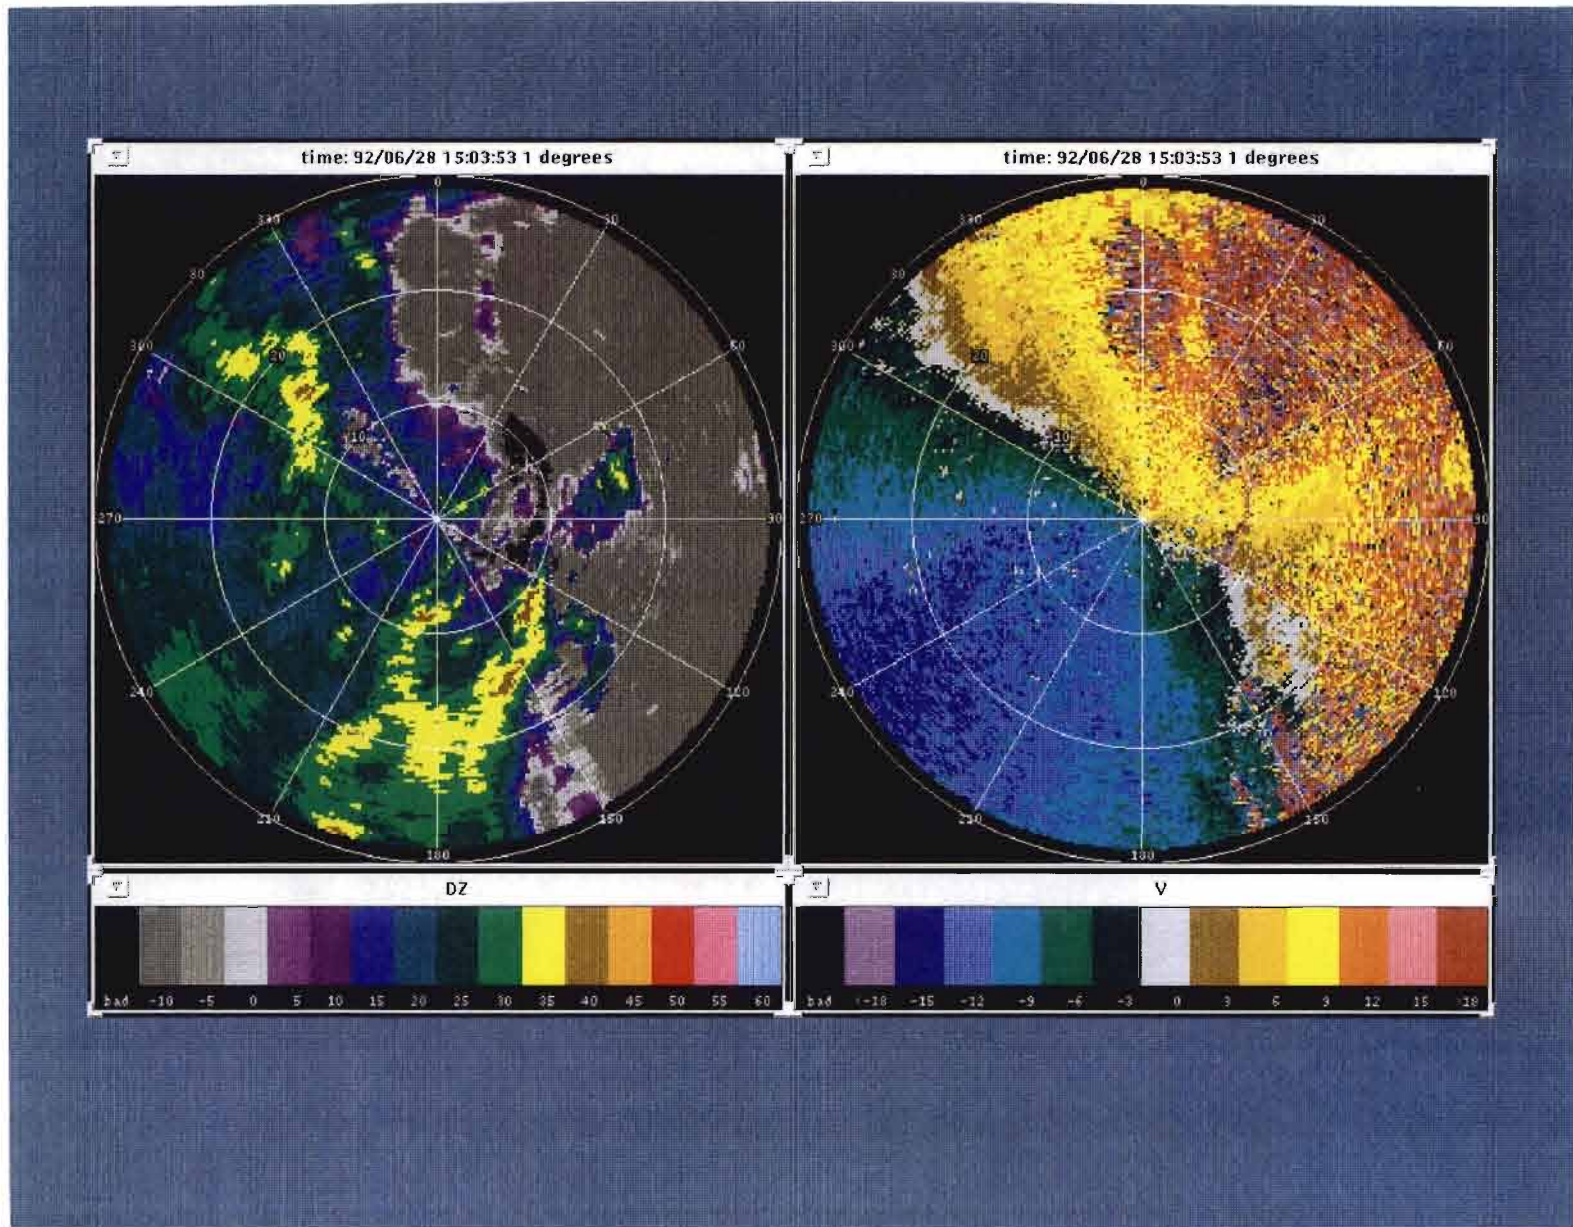


Figure 33. Example of a thin line echo associated with stratiform rain. The elongated low-reflectivity weather echo centered at 15 degrees azimuth, 20 km range triggered a false detection. Range rings are in km, reflectivity (DZ) in dBZ, and radial velocity (V) in m/s.

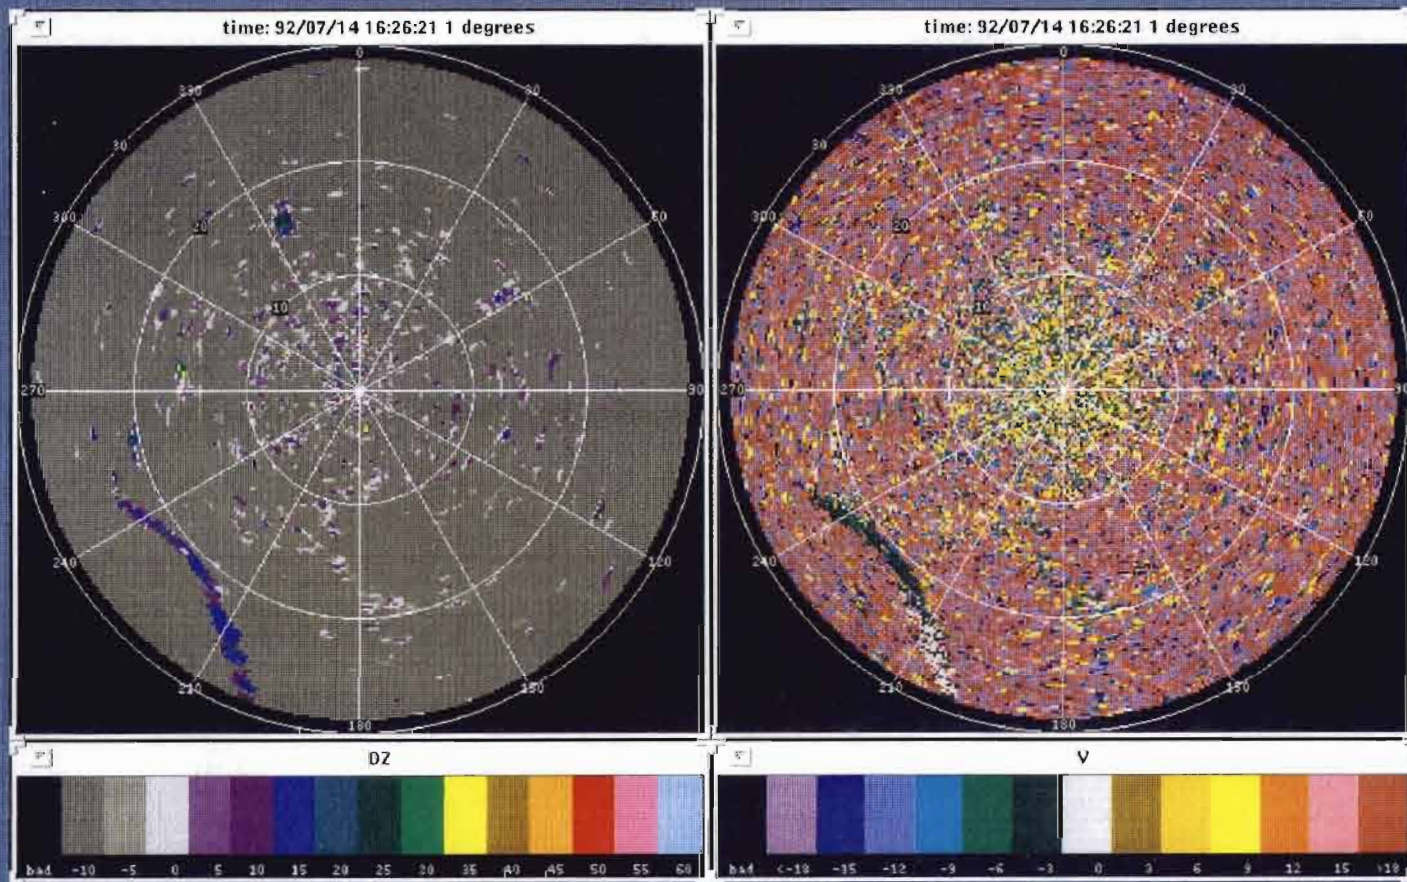


Figure 34. Example of a high altitude thin line echo not associated with a gust front. The echo moved at about 3 m/s and triggered a false detection. Range rings are in km, reflectivity (DZ) in dBZ, and radial velocity (V) in m/s.

6.2.3 Sidelobe or Backlobe Contamination

High reflectivity weather at close range can produce strong return signals which show up in sidelobes or through the backlobes of the ASR-9 antenna pattern. This produces ghost echoes located at identical range intervals, but at azimuths which are offset significantly from the primary weather echo. It has been observed most often in the low beam velocity data as a broadly smeared azimuthal arc of coherent velocity estimates having the same Doppler velocity sign and magnitude as the primary weather echo source. The arc may subtend more than 180 degrees and is generally limited to weather reflectivities greater than 40 dBZ occurring within 20 km of the radar. The sidelobe returns are typically less than 0 dBZ, but have been observed as strong as 5 dBZ; because of the low reflectivity, they are often not apparent in the the reflectivity images.

Examination of image sequences for 25 gust fronts recorded during 1991 operations in Orlando revealed that sidelobe/backlobe echoes were fairly common, occurring at least once in 40% of the 25 gust front event periods examined. Persistence was variable, ranging from 5 minutes to one half hour, depending on movement and intensity of the high reflectivity storm region. The reflectivity and range extent of the sidelobe returns are comparable to those expected for gust front reflectivity thin lines and can therefore be a potential source of false detections by the gust front algorithm. Figure 35 shows a particularly strong occurrence of sidelobe/backlobe contamination. In this case, the contamination is apparent in both the reflectivity and low-beam velocity data. Note that the velocity estimates in the contaminated region are incorrectly reported as ranging between -9 and -12 m/s. (Examination of corresponding TDWR velocity data showed estimates ranging from +3 to +6 m/s.)

When a gust front passes through one of the sidelobe echoes, a false convergent velocity signature is often produced. This can actually make the location of the gust front in velocity data more apparent, but it would lead to incorrect convergence estimates if the contaminated velocity data were used to compute the shear across the front.

When sidelobe-generating weather is approaching the radar site, sidelobes can be discriminated from gust fronts by the fact that the sidelobe thin line is converging on the radar site while gust fronts tend to be outwardly expanding. This means of discrimination has been exploited as a heuristic in the tracking module of MIGFA. Outward moving sidelobes cannot be discriminated. Since a false alarm on an outward moving sidelobe would not impact airport operations, discrimination is not as critical.

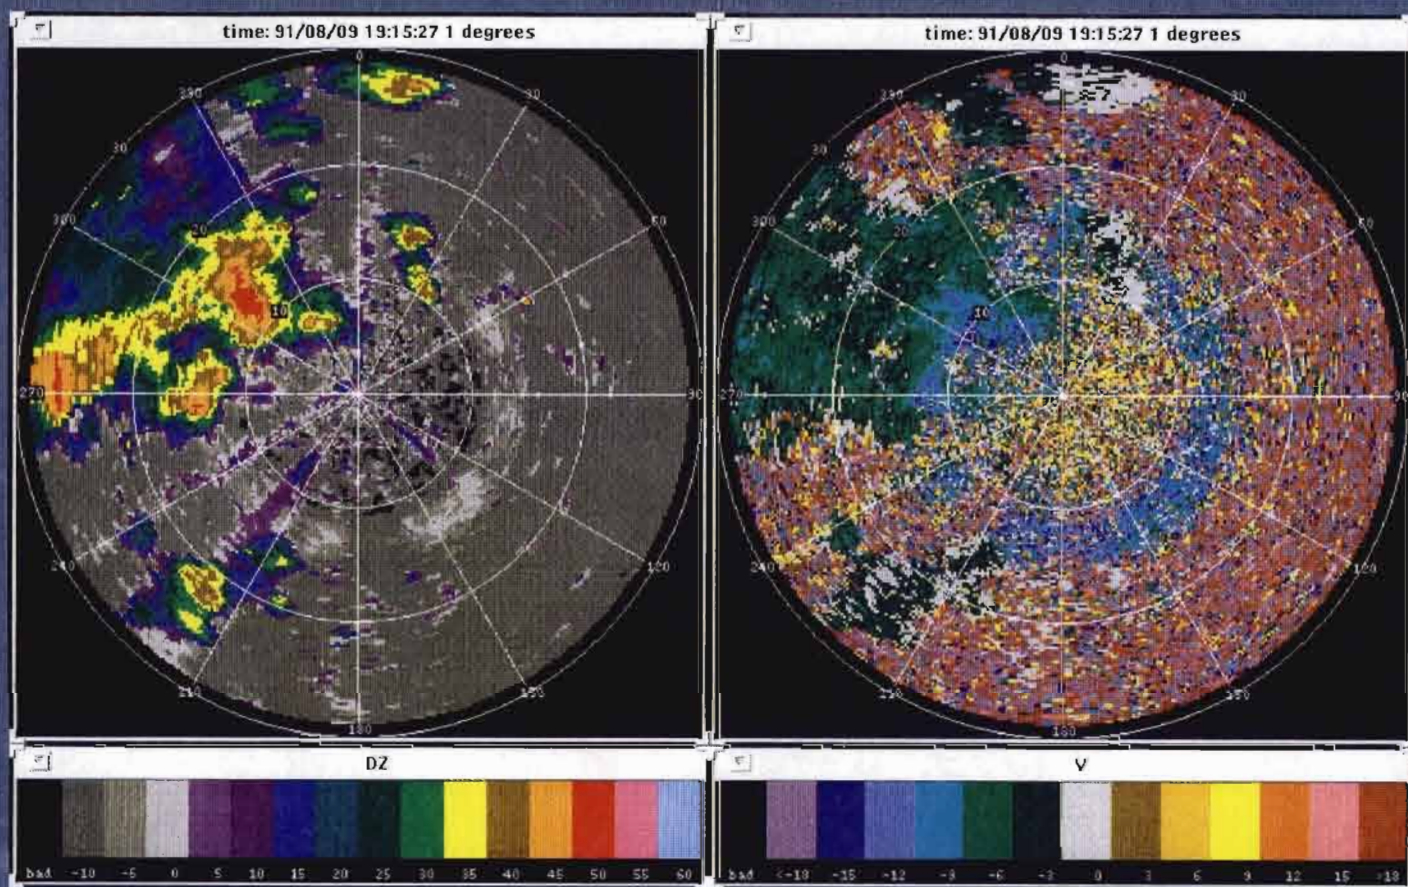


Figure 35. Example of sidelobe contamination in ASR-9 WSP reflectivity and velocity images. Sidelobe contamination appears as a broad ring of slightly elevated reflectivity values and a corresponding ring of lower Doppler variance. In this example, the rings can be seen at a range of 13 km. Range rings are in km, reflectivity (DZ) in dBZ, and radial velocity (V) in m/s.

6.2.4 Obscuration

During the 1992 operational test period, several gust fronts were either detected late, prematurely lost, or not detected at all due to obscuration by anomalous propagation (AP), storm cells, or out-of-trip weather. Some examples and possible algorithm enhancements to improve performance are presented below.

Anomalous Propagation: A strong moisture gradient coupled with a temperature inversion produces conditions most favorable for downward ducting of the radar beam, also known as anomalous propagation (AP). Since the ASR-9 WSP relies on a static clear day map of the ground clutter for clutter removal, an enhanced level of clutter breakthrough in unmapped regions results during AP episodes. The AP returns often resemble normal weather returns and could produce narrow reflectivity features which might be mistaken by the algorithm for gust front thin lines. An example of a large number of false weather echoes attributable to AP can be seen in the bottom half of Figure 36. AP is responsible for the large regions of mottled reflectivity northwest and southeast of the airport. These reflectivity regions are absent from the unaffected TDWR data shown in the top half of the figure. Note the zero Doppler magnitude velocities associated with the AP signals.

Figure 36 shows a gust front amidst AP. Because the gust front thin line signature is distorted, the TL-DZ detector does not generate much interest. Recent ASR-9 WSP processor enhancements permit front-end recognition of AP signals by flagging sample volumes that have near-zero Doppler magnitude along with a corresponding low spectrum width. One possible way of mitigating the obscuration is to set to nil (missing) any pixel location in the TL-DZ interest image where AP had been identified, thereby deferring to the evidence generated by other feature detectors. In particular, the motion detectors are relatively immune to the effects of AP. AP echoes, being generated by returns from ground clutter, do not move. Consequently, the differencing of consecutive images, which is the basis of motion estimation in the motion feature detectors, tends to subtract out AP as background, thereby revealing the moving gust front thin line.

Out-of-trip Weather: Figure 37 shows an example of a gust front reflectivity thin line obscured by out-of-trip weather. Our initial concern was to prevent false alarms on the reflectivity thin lines that typify out-of-trip weather. So, wherever out-of-trip weather was identified, the associated interest values were reduced. As long as interest values were suppressed only along a narrow wedge, gust fronts that overlapped with the out-of-trip weather could still be detected. However, when the area of distant precipitation is extensive, relatively broad wedges of out-of-trip weather can show up on the radar screen. In such a case, a whole sector of interest values is suppressed, effectively preventing the detection of the gust front overlapping the out-of-trip weather.

The effects of this problem can be reduced by using computed thin line orientations more effectively. The map ORIENT generated during thin line smoothing indicates the best orientation of thin lines for each image pixel. Knowing that any thin line resulting from out-of-trip weather will always be radially aligned, ORIENT can be used to selectively suppress interest scores only on radially aligned thin lines. Thin line signatures that are not radially aligned, even when embedded in out-of-trip weather, would not be affected.

Storm cells: Gust fronts can initiate convection along their boundaries, causing storms to form along or immediately behind them. Or, gust fronts can propagate into neighboring storm cells. In either case, the high reflectivity storm cells can occlude parts or all of the less reflective gust front. Even though there exists sufficient reflectivity in the storm cells to generate reliable Doppler estimates, these estimates may reflect motion in the upper regions of the storm and not be indicative of the low-altitude winds. Once again, convergence is not a viable signature for gust front detection in imagery from the present ASR-9.

Certainly, storm cells will interfere with establishing an initial gust front detection. But, given a sequence of images, there usually exists some time interval when a significant part of the gust front is unobscured and when tracking can be initiated. What has been particularly frustrating in the 1992 operational test period is that occlusion by storm cells several times prematurely terminated gust front tracking. Figure 38a shows an example of a gust front during the time it was being tracked. Figure 38b shows the same gust front 4 minutes later at the time tracking was halted due to obscuration. One way of addressing this problem may be to adjust the ANTICIPATION interest image by selectively boosting the interest value of predicted gust front locations wherever storm cells are detected. The boosted anticipation levels should help carry a gust front detection through occluding storm cells.




Figure 36. Example of gust front obscuration due to clutter breakthrough from anomalous propagation (AP). Corresponding TDWR reflectivity and velocity images (top) are provided for comparison with ASR-9 WSP images (bottom). Widespread AP is responsible for the much of apparent weather reflectivity seen in the ASR-9 WSP data northwest and southeast of the airport. A gust front is clearly visible extending from west to east, passing 18 km to the south of the airport in the TDWR images, but is obscured by the AP in the ASR-9 WSP data. Range rings (in km) are plotted with respect to the ASR-9 WSP location. Reflectivity (DZ) is in dBZ, and radial velocity (V) in m/s.

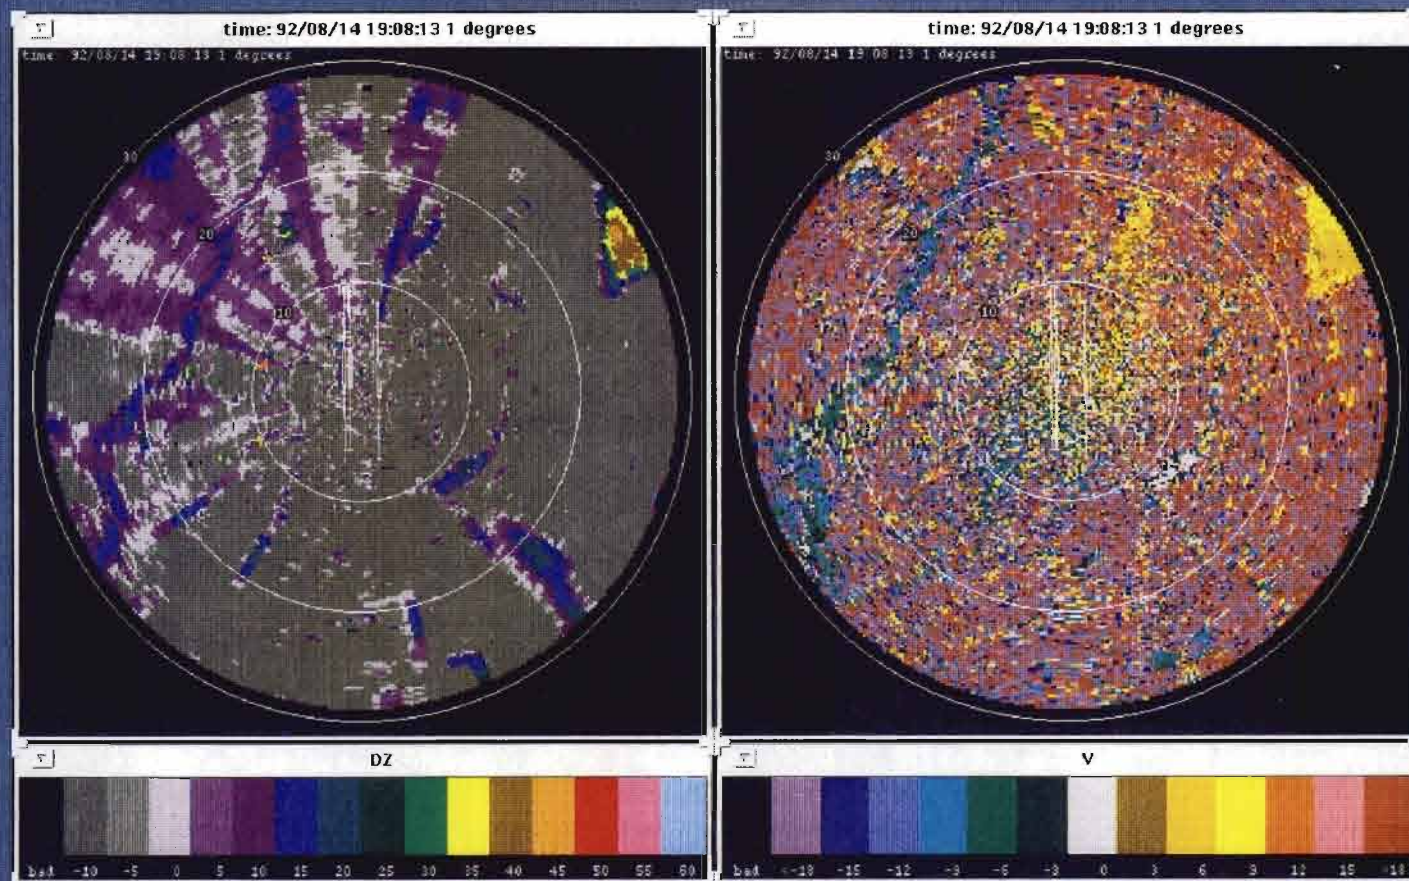


Figure 37. Obscuration of a gust front by range ambiguous (out-of-trip) weather echoes. The gust front is located 18 km northwest of the radar, extending from northeast to southwest. The large grouping of long, radially-oriented reflectivity wedges in the northwest quadrant are out-of-trip echoes that are partially occluding the gust front thin line in the reflectivity image. Note the lack of Doppler velocity coherence in the out-of-trip echo regions. Range rings are in km, reflectivity (DZ) in dBZ, and radial velocity (V) in m/s.

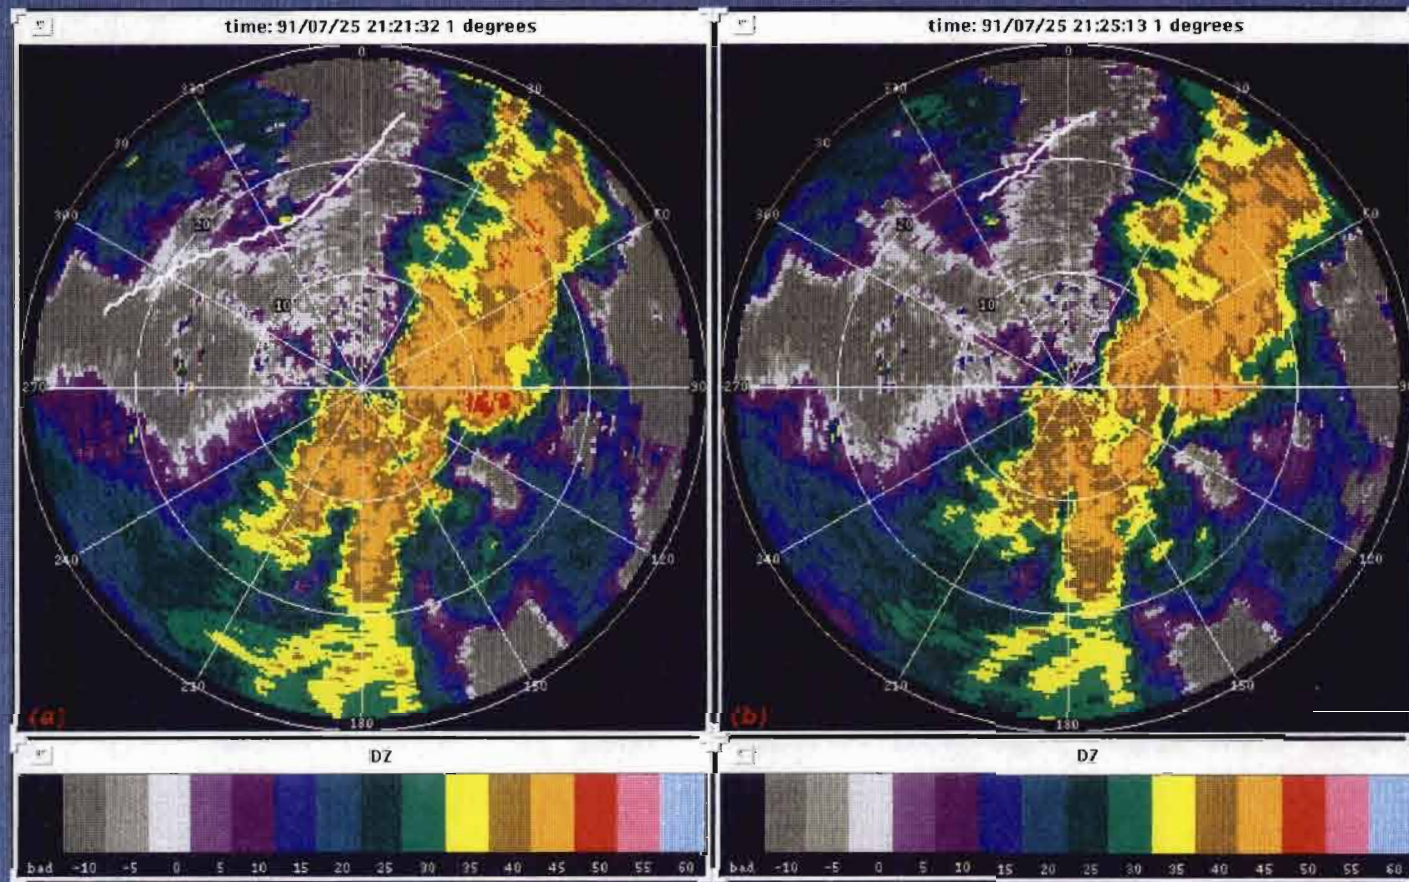


Figure 38. ASR-9 WSP reflectivity images illustrating obscuration of a gust front by storm cells. (a) Gust front thin line echo with MIGFA detection overlaid (white curve). (b) Final MIGFA detection of the same gust front 4 minutes later as tracking was being dropped due to obscuration by precipitation. Range rings are in km, reflectivity (DZ) in dBZ.

6.3 Site dependent knowledge

MIGFA has been designed such that knowledge can be easily exploited to improve gust front detection performance. A rule-based expert decides which feature detectors should be selected from a library of available detectors. By means of this mechanism, MIGFA automatically configures itself to either ASR-9 WSP or TDWR processing. The same mechanism can also be used to implement site-specific knowledge, although our intention is to use site-specific knowledge within MIGFA as little as possible. The more site customizations that are required, the higher the expenses of deployment and maintenance will become.

Two sources of variance, the ASR-9 itself and the geography, are not expected to require site-specific modifications. With proper calibration, each deployed ASR-9 WSP should provide reflectivity and velocity estimates of comparable quality to those used in our tests. No ASR-9 calibration-related adjustments to MIGFA should be required. The collective impact of miscellaneous signal losses (e.g., transmitter, receiver, microwave plumbing) could reduce the overall sensitivity of the radar. Although this would likely reduce end-to-end system detection performance, tuning of the MIGFA algorithm would probably not compensate. As for variations in geography, clear day clutter maps used by the ASR-9 to filter ground clutter are imperfect; changing topography and atmospheric conditions can cause clutter breakthrough. The impact on MIGFA performance is believed to be minor, since the motion detectors discriminate against stationary echoes characteristic of clutter breakthrough.

A third source of site variance, the dominant weather patterns, may require some site-specific modifications. At each location, depending indirectly on the geography and climate, gust fronts may exhibit behaviors different from those observed at other sites. From past experience accumulated during AGFA development, gust fronts in the midwestern U.S. tend to move faster on average than those in the southeast and Florida peninsula. Gust fronts tend to have lower reflectivity values in the High Plains than in the Southeast. While gust fronts tend to trigger storm cells in the unstable, humid air of central Florida, gust fronts do so less in Denver where the air is drier. Gust front acceleration can be affected by the funneling effects of hills and mountains.

In general, because of the tolerance to uncertainty built into functional templates and the delay of discriminatory thresholding until after all data have been translated into a single map of interest, MIGFA should be relatively robust in the face of most of these site variations. So far, only one feature detector includes site-specific knowledge that will probably require changes when moved to other sites. The modulation of background interest in the ANTICIPATION interest image, based on the relative amounts of stratiform rain and storm cell activity, will probably not work in all places. Gust fronts are almost never observed in central Florida in the presence of diffuse stratiform rain, which is typically represented by relatively low reflectivity values ranging from 15 to 30 dBZ. In contrast, strong dry microbursts (and subsequently gust fronts) occurring near Denver can be generated by cells with reflectivity values in the same 15 to 30 dBZ range.

6.4 Resource Requirements

The test version of MIGFA has been implemented using SKETCH, a computer vision development environment implemented using Austin Kyoto Common Lisp and C at MIT Lincoln Laboratory. The primary SKETCH design goal is to facilitate the process of software development; run-time efficiency is a secondary consideration. The exception to this is FTC. Over half of the processing time required by MIGFA is spent in FTC, which has been optimized for speed. With this information in mind, the following statistics give a crude idea of computer resource requirements.

The average processing time per ASR-9 WSP scan is about 40 seconds on a Sparc2 rated at 28 MIPS equipped with 32 Mbytes of memory. The worst case processing time observed was under 55 seconds. In addition to MIGFA, other processes running to support MIGFA are data collection, display monitoring, LLWAS processing, and result reporting. All together, the average processing time was about 90 seconds. During operational testing, 120 seconds were allocated for processing each radar scan.

7. CONCLUSIONS

Algorithms to detect thin line signatures in reflectivity images and convergence signatures in Doppler images are conceptually easy to design and implement. And yet, several research groups have worked collectively for nearly 10 years to develop reliable automatic gust front algorithms. Despite the effort, no algorithm has demonstrated performance comparable to human performance.

The problem is that automatic gust front detection is deceptively a much harder problem than simply detecting one or both of these signatures. In order for human interpreters to detect and track gust fronts, they make use of knowledge about the radar and the weather. They use spatial and temporal context. And, they have the ability to deal with uncertainty while assimilating ambiguous or even contradictory evidence. The large performance gap between algorithmic and human gust front detection probably reflects the lack of these perceptual skills in previous algorithmic approaches. With the machine intelligence techniques developed as part of XTRS, MIGFA displays levels of performance that are competitive with human interpreters.

Possible improvements to the ASR-9 WSP version of MIGFA have been identified, based on feedback from the 1992 operational test period at Orlando. Continued refinements are likely to be required as problems are uncovered with testing in more varied weather at other sites. In some cases, alternative customized sets of feature detectors may be required for different radar sites. Hopefully, sites will fall into a few broad classes so that only a few alternative feature detector sets will be necessary.

Although some improvements will undoubtedly be made, it is important to remember that no interpreter, algorithmic or human, can do more than what the capabilities of the sensor allow. The relatively low sensitivity of the ASR-9 means that some gust fronts will not be detected. The lack of reliable convergence signatures means that some false alarms will occur. Given data from a high resolution weather radar such as the TDWR, detection performance ought to be much better. A TDWR version of MIGFA, which has in fact already been assembled and is being tested, confirms these expectations. The TDWR version uses a different set of feature detectors in order to exploit the higher sensitivity reflectivity data and reliable convergence signatures in the Doppler data. The control structure and the extraction, tracking and prediction modules are identical. The TDWR version of MIGFA has been tested offline with data collected in previous years. Although a direct case-by-case comparison of AGFA and MIGFA has not yet been run for TDWR data, MIGFA again appears to substantially outperform AGFA. The TDWR feature detectors, as well as complete test and evaluation results will be presented in another report.

Another candidate radar system is the Next Generation Weather Radar, which goes by the acronym WSR-88D (formerly NEXRAD). A recent feasibility study published by the National Severe Storms Laboratory concluded that, with some modifications regarding image resolution, GFDA should work with WSR-88D data[21]. A sequel to this report went on to suggest that AGFA could be used for detecting a variety of meteorological phenomena besides gust fronts, including synoptic fronts, sea breeze fronts, gravity waves, terrain-induced convergence boundaries,

old thunderstorm outflow boundaries, and convergence bands within cyclones[22]. Assuming that GFDA can be extended to WSR-88D data, our belief is that MIGFA can also.

MIGFA is based on XTRS, a general system for automatic object detection and recognition. In conventional hierarchical computer vision design, machine (artificial) intelligence techniques are applied only in the higher (more abstract) levels, manipulating symbolic representations of the data. In contrast, XTRS provides a framework for applying machine intelligence at the earliest stages of detection processing, manipulating raw pixel data. It is true that this approach will not solve the general vision problem of understanding a complex scene, i.e., understanding the relationships of several instances of hundreds of possible objects in an unconstrained contextual environment. But for problems in which the goal is to detect one particular type of object and discriminate it from an understood background, XTRS provides an effective, straightforward means. In this report, we have described an algorithm based on XTRS that represents a significant improvement in performance over previous efforts at automated gust front detection. With these techniques, XTRS should also be applicable to a variety of other meteorological detection problems. For example, XTRS is currently being used to develop a microburst prediction algorithm for TDWR data as part of the Integrated Terminal Weather System (ITWS).

GLOSSARY

A/D	Analog-to-Digital
AGFA	Advanced Gust Front Detection Algorithm
ANTICIPATION	Feature detector highlighting regions of expectation
AP	Anomalous Propagation
ASR-9 WSP	Airport Surveillance Radar 9 with Wind Shear Processor
ATC	Air Traffic Control
ATR	Automatic Target Recognition
dBZ	Decibel (referenced to reflectivity factor Z)
DZ	Reflectivity image
DZ-MOTION	Motion detector for DZ features
FAA	Federal Aviation Administration
FT	Functional Template
FTC	Functional Template Correlation
GFDA	Gust Front Detection Algorithm
GFUPDATE	Gust Front Update Task
GSD	Geographical Situation Display
I/Q	In-Phase/Quadrature
ITWS	Integrated Terminal Weather System
LLWAS	Low Level Wind Shear Alert System
MIGFA	Machine Intelligent Gust Front Algorithm
MIPS	Machine Instructions Per Second
NEXRAD	Next Generation Weather Radar
NSSL	National Severe Storms Laboratory
OUT-OF-TRIP	Feature detector for range ambiguous echoes
PFA	Probability of False Alarm
PFD	Percent of False Length Detected
PLD	Percent of Length Detected
POD	Probability of Detection
PRF	Pulse Repetition Frequency
RDT	Ribbon Display Terminal
SD	Velocity standard deviation texture map
SD-MOTION	Motion detector for velocity SD features
SNR	Signal-to-Noise Ratio
STC	Sensitivity Time Control
TL-DZ	DZ thin line feature detector
TL-SD	SD thin line feature detector
TRACON	Terminal Radar Approach Control
TDWR	Terminal Doppler Weather Radar

GLOSSARY
(Continued)

V	Velocity image
WSA	Wind Shear Alert
WSR-88D	Weather Surveillance Radar - 88D (formerly NEXRAD)
XTRS	Experimental Target Recognition System

REFERENCES

1. M. Weber and T. Noyes. Wind shear detection with airport surveillance radars. *Lincoln Laboratory Journal*, 2(3):511-525, Fall 1989.
2. T. Noyes, S. Troxel, M. Weber, O. Newell, and J. Cullen. The 1990 Airport Surveillance Radar Wind Shear Processor (ASR-WSP) operational test at Orlando International Airport. Project Report ATC-178, MIT Lincoln Laboratory, Lexington, MA, July 1991.
3. M. Weber. Airport Surveillance Radar (ASR-9) Wind Shear Processor: 1991 test at Orlando, FL. Project Report ATC-189, MIT Lincoln Laboratory, Lexington, MA, June 1992.
4. E. Chornoboy. Optimal mean velocity estimation for Doppler weather radars. *IEEE Trans. Geo. Remote Sensing*, submitted for publication, 1992.
5. D. Klingle, D. Smith, and M. Wolfson. Gust front characteristics as detected by Doppler radar. *Monthly Weather Review*, 115(5):905-918, May 1987.
6. H. Uyeda and D. Zrníc. Automated detection of gust fronts. *J. Atmos. Oceanic Tech.*, 3:36, 1986.
7. A. Witt and S. Smith. Development and testing of the gust front algorithm. Technical Report DOT/FAA/PS-87/4, FAA, 1987.
8. S. D. Smith, A. Witt, M. Eilts, D. Klingle-Wilson, S. Olson, and J. P. Sanford. Gust front detection algorithm for the Terminal Doppler Weather Radar part I: Current status. In *Proceedings of the 3th International Conference on the Aviation Weather System*, pages 31-34, Anaheim, CA, January 1989.
9. G. Hermes, A. Witt, S. Smith, D. Klingle-Wilson, D. Morris, G. Stumpf, and M. Eilts. The gust front detection and wind shift algorithms for the Terminal Doppler Weather Radar system. *Journal Atmos. Oceanic Tech.*, (in press), 1992.
10. D. Klingle-Wilson, M. Donovan, S. Olson, and F.W. Wilson. A comparison of the performance of two gust front detection algorithms using a length-based scoring technique. Project Report ATC-185, MIT Lincoln Laboratory, Lexington, MA, May 1992.
11. M. Eilts, S. Olson, G. Stumpf, L. Hermes, A. Abrevaya, J. Culbert, K. Thomas, K. Hondl, and D. Klingle-Wilson. An improved gust front detection algorithm for the TDWR. In *Proceedings of the 4th International Conference on the Aviation Weather System*, pages J37-J42, Paris, France, June 1991.
12. M.W. Merritt, D. Klingle-Wilson, and S.D. Campbell. Wind shear detection with pencil-beam radars. *Lincoln Laboratory Journal*, 2(3):483-510, 1989.
13. Jacques G. Verly, Richard L. Delanoy, and Dan E. Dudgeon. Machine intelligence technology for automatic target recognition. *Lincoln Laboratory Journal*, 2(2):277-311, Summer 1989.

REFERENCES (Continued)

14. Richard L. Delanoy, Jacques G. Verly, and Bryan Williams. Region-based target recognition from laser radar imagery using appearance models. Technical Report TR-933, MIT Lincoln Laboratory, Lexington, MA, July 1992.
15. Richard L. Delanoy, Jacques G. Verly, and Dan E. Dudgeon. Functional templates and their application to 3-D object recognition. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, San Francisco, California, March 1992.
16. Richard L. Delanoy, Jacques G. Verly, and Dan E. Dudgeon. Pixel-level fusion using interest images. In *Proceedings of the 4th National Symposium on Sensor Fusion*, Orlando, Florida, April 1991.
17. Stephen Grossberg and Ennio Mingolla. Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception and Psychophysics*, 38(2):141-171, 1985.
18. S. Levialdi. Parallel pattern processing. *IEEE Trans. Syst. Man and Cybern.*, SMC-1:292-296, 1971.
19. University Corporation for Atmospheric Research. *Network Expansion LLWAS (Phase III) Algorithm Specification, Version 1990.01*, October 1990.
20. Final report for the air traffic control evaluation of the prototype airport surveillance radar wind shear processor at Orlando International Airport. Survey conducted by the Engineering, Integration, and Operational Evaluation Service, Weather and Primary Radar Division, ACW 200, FAA Technical Center, October 1992.
21. Gregory Stumpf. On the potential use of the Terminal Doppler Weather Radar gust front detection algorithm on the WSR-88D system. Part 1: Impacts of radar system differences. Technical Report DOT/FAA/NR-92/8,I, NOAA/ERL National Severe Storms Laboratory, Norman, Oklahoma, September 1991.
22. Gregory Stumpf. On the potential use of the Terminal Doppler Weather Radar gust front detection algorithm on the WSR-88D system. Part 2: Detecting non-gust front convergent weather phenomena. Technical Report DOT/FAA/NR-92/8,II, NOAA/ERL National Severe Storms Laboratory, Norman, Oklahoma, October 1991.