
Supporting the Deployment of the Terminal Doppler Weather Radar (TDWR)

James E. Evans and David M. Bernella

■ The Terminal Doppler Weather Radar (TDWR) program was initiated in the mid-1980s to develop a reliable automated Doppler-radar-based system for detecting weather hazards in the airport terminal area and for providing warnings that will help pilots avoid these hazards when landing and departing. This article describes refinements made to the TDWR system since 1988, based on subsequent Lincoln Laboratory testing in Kansas City, Missouri, and Orlando, Florida. During that time, Lincoln Laboratory developed new capabilities for the system such as the integration of warnings from TDWR and the Low Level Wind Shear Alert System (LLWAS). Extensive testing with the Lincoln Laboratory TDWR testbed system has reconfirmed the safety benefits of TDWR.

IN THE mid-1980s, the Federal Aviation Administration (FAA) initiated the Terminal Doppler Weather Radar (TDWR) program at Lincoln Laboratory in response to a need for improved surveillance of real-time hazardous weather (especially low-altitude wind shear) at major airports known for frequent thunderstorm activity. The initial focus of the TDWR program was to provide reliable fully automated Doppler-radar-based detection of microbursts and gust fronts and 20-min warnings of wind shifts that could affect runway usage. Figure 1 shows the locations planned for the TDWR systems. An article in a previous issue of this journal [1] described how the algorithms that accomplish fully automated wind-shear detection were developed and validated prior to being implemented in the TDWR system. In this article, we describe recent work at Lincoln Laboratory:

1. further refinement of the baseline technical and operational capability of the TDWR by the testing of initial products in additional meteorological environments,
2. validation of key elements of the production

contractor TDWR design, including the interfaces between contractor- and Lincoln Laboratory-developed system elements,

3. determination of TDWR radar locations to optimize wind-shear detection in difficult clutter environments, and
4. development of new products to meet the additional needs of terminal air traffic users.

The TDWR program at Lincoln Laboratory has provided the opportunity for both technical and technology-transfer innovations. Determining the nature and hazard of the wind-shear phenomena to be detected requires advanced pattern-recognition and pattern-analysis techniques. The signal-to-clutter ratios can be quite low because we are concerned with the detection of wind shear in clear-air radar scattering conditions at quite low altitudes (often at about 100 m above ground level [AGL]) in cluttered urban environments. Further complicating the detection challenge is a high likelihood of range-ambiguous returns because of the need to make observations of the weather at relatively high pulse-repetition frequencies (PRF).

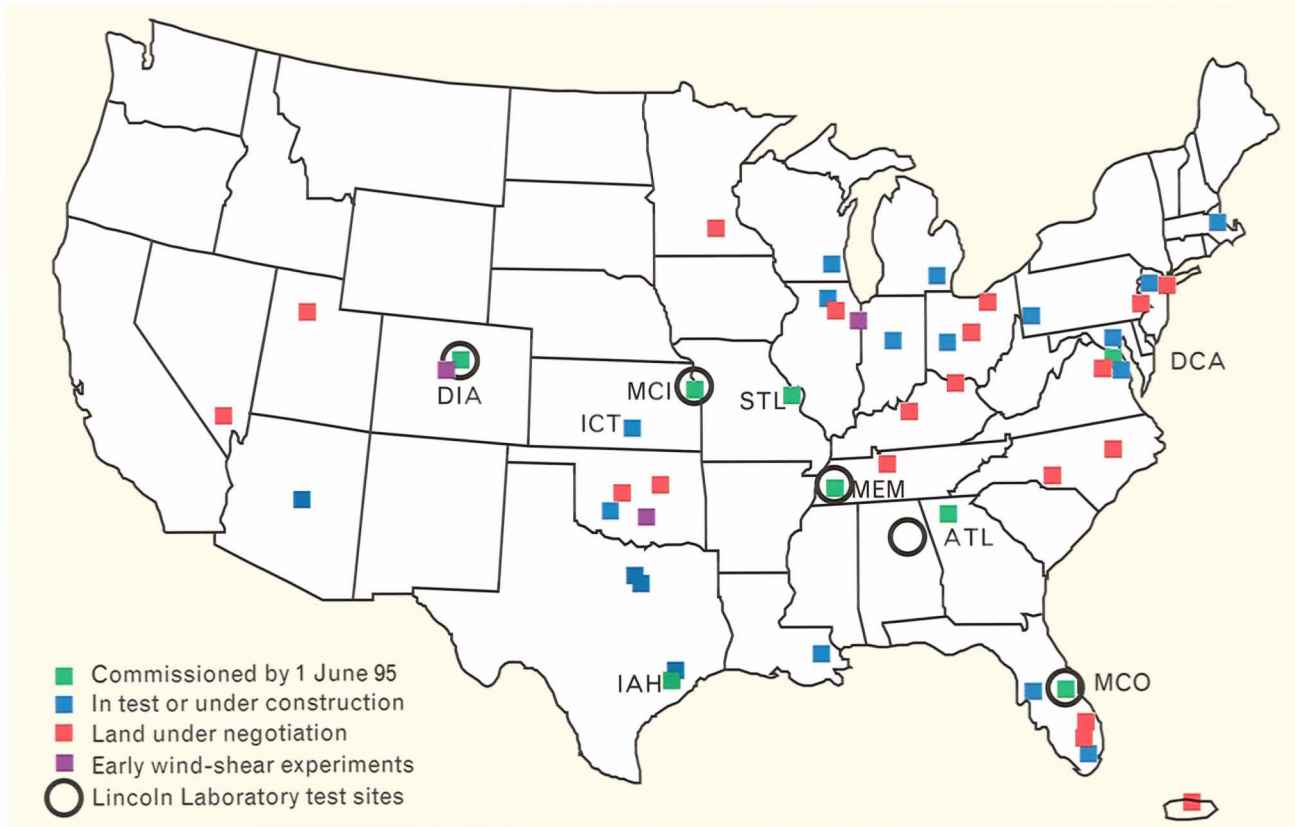


FIGURE 1. Locations and statuses of TDWR installations.

There is also an urgent need to improve the transition of research at organizations such as Lincoln Laboratory into operational systems in the field. Thus the recent technology-transfer approach used for the TDWR program should be of interest to readers who may have only limited interest in the TDWR technology.

Government procurements of high-technology systems have typically proceeded by one of two paths:

1. The government sets specifications for end-to-end performance of the system and the contractors attempt to meet this performance. In some cases, the government specifications will be influenced significantly by the capability demonstrated by a particular firm. Although government research is a guide to the contractors, it need not be utilized to achieve the specified results. An example of this is the development of the Stealth aircraft.
2. The government determines all of the details of the desired system and the contractor is re-

quired to duplicate the system and the specified details. One example of this approach is when the government provides most or all of the software that will be used in the operational system.

The TDWR development with Lincoln Laboratory's involvement used a mixture of the above two approaches. The meteorological/pattern-recognition expertise and data to develop automatic wind-shear detection algorithms resided largely in research organizations such as Lincoln Laboratory. On the other hand, there were elements of the signal waveform design and signal processing in which contractor innovation was very important. This hybrid acquisition approach was very powerful in terms of drawing on the best knowledge of both industry and the government-sponsored research to arrive at a better deployable system. This approach, however, gave rise to some difficult system integration and validation challenges that are described later in this article.

Another important element of improved technology transfer between research organizations and indus-

trial firms is the transition of algorithms that are embodied in operational system software. The Weather Sensing Group at Lincoln Laboratory has focused on rapid prototyping, in which a variety of existing and new software packages are used to achieve rapid-turnaround off-line testing, after which the software is dropped into the existing systems at operational ATC facilities for real-time testing and user-need clarification. Because of the continuing evolution in software engineering technology and language preferences, the resulting software often uses a number of different languages, it may not be documented well, and it frequently is somewhat inefficient in an operational setting. Thus the Lincoln Laboratory prototype code has typically not been suitable for operational use in the FAA systems.

The algorithm transfer approach used to date has been for Lincoln Laboratory to specify the algorithms in a generic computer science language, and for the contractor to develop optimized real-time code from this specification. An important element of the TDWR program's recent work has been the investigation of other ways to transfer the Lincoln Laboratory knowledge to the production contractor so that the operational capability is achieved more quickly and at lower cost.

This article proceeds as follows. In the next section, we provide a background of the TDWR program, including information on the program's motivation, previous testing of the TDWR system, and the nature of the allocation of responsibilities between the contractor and government. The section "Validation of the Production TDWR System Features" describes how the Lincoln Laboratory TDWR testbed was modified to validate key elements of the production system and to assess its sensitivity to various meteorological conditions. As a result of continued testing with the Lincoln Laboratory prototype, several refinements were made to the initial algorithms to address site-specific meteorological issues that are discussed in the section "Refinements of the Initial TDWR Products." Because the physical location of a TDWR system can significantly affect the system's detection of low-altitude wind shear, Lincoln Laboratory was requested to provide technical assistance in finding the optimum locations for a number

of operational systems. This effort is described in the section "Site Selection." The next section, "Product Refinements to Meet Additional FAA Needs," discusses the storm-motion and TDWR/Low Level Wind Shear Alert System (LLWAS) integration algorithms that were developed in response to additional air traffic information needs. The development of these new algorithms presented opportunities to try new approaches for transferring software. This article concludes with a summary of the current status of the Lincoln Laboratory program and a brief description of the future work in support of the TDWR program.

Background

The principal impetus of the TDWR program was a series of major air carrier accidents during the 1970s and 1980s, culminating with the crash, caused by wind shear from a microburst, of an L-1011 wide-body jet at Dallas-Fort Worth airport in 1985. Using results from scientific experiments (notably by T. Fujita of the University of Chicago and by researchers at the National Center for Atmospheric Research [NCAR]), Lincoln Laboratory constructed a Doppler weather radar testbed to obtain low-altitude wind-shear data and to develop fully automated algorithms for the detection of wind shear. Following preliminary tests in Memphis, Tennessee, and Huntsville, Alabama, the Lincoln Laboratory testbed (operating in the FAA-authorized S-band of the RF spectrum) was used in the successful demonstration of initial TDWR capability—including wind-shear and precipitation detection, clutter suppression, and antenna-scanning strategy—during an operational test and evaluation at Denver's Stapleton Airport in 1988 [2].

The TDWR capability demonstrated at Denver consisted of

1. microburst detection (to provide wind-shear or microburst alerts with an estimated wind-speed loss),
2. gust-front detection (to provide wind-shear alerts with an estimated wind-speed gain),
3. the expected gust-front locations 10 and 20 min in the future,
4. the expected wind velocity behind a tracked gust front, and



FIGURE 2. In the TDWR system, the Situation Display (SD), shown in the left, and the Ribbon Display Terminal (RDT), shown in the right, provide key information to air traffic users. The SD shows the locations of wind-shear events on a precipitation reflectivity background, and the RDT provides alphanumeric wind-shear alert messages in a format suitable for direct readout.

5. the precipitation reflectivity according to the National Weather Service (NWS) six-level standard.

The technical performance criteria, which were satisfied during the Denver tests, included a microburst-detection probability greater than 0.9 with a false-alarm probability less than 0.1. It was desired that warnings be reported at least 1 min prior to a plane's encountering a microburst.

This information was provided to air traffic users by the two displays shown in Figure 2:

1. a color Situation Display (SD), which shows the locations of wind-shear events on a precipitation reflectivity background, and
2. a Ribbon Display Terminal (RDT), which provides alphanumeric wind-shear alert messages in a format suitable for direct readout (e.g., "microburst alert at 2 nmi final, expect a 30 knot loss").

Supervisors in the tower and radar control rooms use the SDs for operations planning while the tower

controllers use the RDTs for reporting events directly to pilots.

From 1987 through 1988, the FAA conducted a competitive procurement for the TDWR contract. As a result of the successful Denver tests [2], the FAA awarded a contract to Raytheon Co. (Sudbury, Massachusetts) for the production of forty-seven TDWR systems. Included in the terms of the contract was implementation of the government-supplied specifications for wind-shear detection, PRF selection, ground-clutter filtering and residue editing, and antenna-scanning scenarios. Furthermore, Raytheon would be responsible for developing data-quality algorithms and the signal waveforms and processing necessary to obtain unambiguous Doppler velocities from the weather returns. Because the unfolded, or ambiguity resolved, Doppler velocities were the principal input to the microburst and gust-front algorithms provided by the government, this allocation of responsibilities made Raytheon's work a key element of the overall wind-shear detection process and there-

fore gave rise to a need for careful validation of the interface between Raytheon- and government-developed elements if the first deliveries were to occur in 1992 as planned.

By the time the contract was awarded to Raytheon, the basic algorithms for the detection of wind shear had been successfully tested on small, fast-moving storms near Memphis, on larger storms in Huntsville, and extensively on “dry” microbursts in Denver. (A dry microburst is a microburst in which the associated rain reaching the ground is so slight that the ground remains dry.) These tests, however, did not represent all meteorological conditions at the planned TDWR installations. Furthermore, because of a lack of available RF spectrum allocation in the S-band near airports, the production TDWRs were specified to operate at an RF frequency in the C-band—one-half the wavelength of the Lincoln Laboratory testbed system. This new condition led to serious concern about possible degradation of both the detection algorithms (because of attenuation from dense thunderstorms) and the velocity de-aliasing algorithms. Therefore, a decision was made to modify the Lincoln Laboratory testbed for C-band operation and to conduct further tests of the algorithms in other meteorological environments during the next two years while Raytheon was in the preproduction phase.

After the 1988 operational test and evaluation in Denver, different geographical areas were considered for further experiments. There were two areas of particular interest: (1) the Central Plains (Kansas City, Missouri), where fast-moving squall line storms prevail, and (2) central Florida (Orlando/Tampa), which statistically has the highest density of thunderstorms in the United States. Because the C-band modification could not be completed in time for the 1989 summer storm season, a decision was made to install the testbed in Kansas City in 1989 and then to move the testbed to Orlando, where the C-band modification could best be performed so that the system would be ready for operational use in time for the 1990 summer storm season. Thus, after operations at Kansas City were terminated in the fall of 1989 and the system was moved to Orlando International Airport, Westinghouse Electric Corp. modified the transmitter/exciter for C-band operation and the

Orlando air traffic personnel evaluated the system during the summer storm seasons from 1990 through 1993. During this period, the system acted as a testbed and mechanism for aiding Raytheon during its design and development phase, providing support for algorithm refinement, for further product development, and for user evaluation. The following section describes this work.

Validation of the Production

TDWR System Features

Raytheon was responsible for developing signal waveforms and software algorithms for the unambiguous determination of the Doppler velocities of the weather returns. The Raytheon approach used two features that the Lincoln Laboratory testbed did not:

1. a key elevation angle was scanned twice at different PRFs so that Chinese remainder methods could be used to determine the unambiguous velocity, and
2. a wind-field model was used to resolve ambiguous data in regions that lacked continuity along the radar radial directions.

The Raytheon approach was an area of overall program risk because the Doppler unfolding algorithm could adversely affect the government-specified algorithms for wind-shear detection. The risk was managed by prolonged real-time testing of the Raytheon algorithm with the TDWR testbed radar in conjunction with an off-line analysis of cases in which problems appeared to have occurred. To carry out these studies, Lincoln Laboratory executed a contract with Raytheon in which the latter would provide (1) technical assistance in the rapid prototype implementation of the signal waveforms and unfolding algorithm in the Lincoln Laboratory testbed, and (2) support to Lincoln Laboratory in the analysis of the testbed data. This approach had several benefits:

1. Raytheon personnel could test contemplated changes to their design by using the Lincoln Laboratory testbed system and its database as opposed to having to go through the very lengthy process of making formal changes to the configuration-controlled Raytheon design.
2. The results of field testing (and the problems encountered) were available immediately to

Raytheon. Thus the company did not have to wait until the formal testing reports were completed, reviewed by the FAA, and distributed.

3. In general, problems can arise in using a fixed-price contract to obtain an in-depth technical assessment because the fixed-price basis of the contract does not provide much incentive for the contractor to perform a detailed and thorough analysis. This issue was not a factor in the TDWR work because Raytheon's support to Lincoln Laboratory was separate from the production TDWR contract.
4. Raytheon personnel working at Lincoln Laboratory could use the laboratory's TDWR data analysis facility (e.g., the displays, the software package, and the data processing personnel) as opposed to having to develop such a capability in-house first.

The Raytheon unfolding algorithm was used for real-time operation in Orlando from 1991 to 1993. Off-line studies with recorded data and comparison with the Lincoln Laboratory unfolding algorithms were used to rectify problems identified in the real-time testing and to provide regression tests for any additional refinements. Failures in the unfolding algorithm were responsible for approximately 20% of the microburst-detection failures and approximately 30% of the gust-front failures. Both of these failure levels were better than the performance of the Lincoln Laboratory unfolding algorithm used in the earlier TDWR testbed operation.

A number of other changes were made to the TDWR testbed to facilitate validation of the Raytheon design. These changes allowed early testing of the point-target and pulse-interference detectors, automatic gain control function, and other system features affecting the data quality. Of major importance was validation of C-band performance because the S-band had been used for all testing prior to the contract award. The S-band transmitter, which was derived from an FAA Airport Surveillance Radar-8 (ASR-8), was replaced by a C-band transmitter in 1990 for summer operations at Orlando to verify proper operation of the TDWR algorithms at the designated C-band. The replacement transmitter, built by Westinghouse with a water-cooled klystron tube



FIGURE 3. TDWR testbed located southwest of Kansas City International Airport.

provided by Lincoln Laboratory, verified adequate performance at C-band. When production air-cooled TDWR klystron tubes became available in late 1990, Westinghouse again modified the transmitter to use an air-cooled tube for the 1991 operations. Thus a testbed was provided for the new tube under field conditions fully one year before Raytheon could have a field system ready. As it turned out, the tube was extremely reliable and operated very well up to the decommissioning of the testbed in the fall of 1993.

To date, the TDWR production systems deployed at airports have encountered few if any problems in the design areas validated with the Lincoln Laboratory testbed system.

Refinements of the Initial TDWR Products

As mentioned earlier, the TDWR algorithms for the detection of wind shear had been tested extensively in Memphis, Huntsville, and Denver, prior to Raytheon's receiving the TDWR production contract. In the Memphis and Huntsville tests, the TDWR system encountered extensive air-mass thunderstorm activity with moderate storm movement. In the Denver tests, the system experienced a mixture of low-rainfall microbursts and High Plains thunderstorms. Although testing at these three sites had exposed the TDWR system to different meteorological conditions, the system had yet to be exposed to a Midwest environment, which is characterized by rapidly moving storms, or to a Florida environment, which is charac-

terized by a very high frequency of slowly moving, highly electrified storms. In this section, we discuss additional changes to the TDWR microburst-detection algorithm that arose from the further testing of the system in Kansas City and Orlando.

Reducing False Alarms Not Associated with Storms

Figure 3 shows the TDWR testbed located near the Missouri River, approximately eight miles southwest of Kansas City International Airport (MCI). During the 1989 testing at this location, an excessive number of microburst false alarms caused by transient low-altitude Doppler velocity features occurred in the clear air in the absence of any storms [3]. These transient features typically arose from noise in the velocity field caused by flocks of birds or swarms of insects and from strong southerly winds passing over the rolling terrain near MCI, especially when the winds intensified in advance of an approaching storm. In addition, microburst false alarms were also caused by surface wind diverging behind strong gust fronts that had moved away from storms.

Close inspection of the false alarms showed that the majority of them occurred when there were very low reflectivities in the airspace above the location of the false alarm. The TDWR microburst algorithm uses vertical reflectivity data to identify microburst precursors and hence to provide early warnings. This existing use of storm reflectivity information suggested that storm cells as identified by the precursor algorithm could be a requirement for a low-altitude radial divergence in the velocity field to be identified as a microburst. The possible drawback of this feature was that a microburst with a very low reflectivity might be discarded erroneously if this test was to be used in certain areas, such as Denver, where low-reflectivity microbursts do occur. Thus the use of substantial reflectivity aloft as a necessary condition for microburst declarations has been implemented in the TDWR algorithm as a site-adaptation parameter. (Note: The likelihood of a dry, or very low reflectivity, microburst in Midwest environments is very low. In fact, less than 2% of all Kansas City microbursts have a core reflectivity less than 18 dBz, which is associated with light rain at the ground surface. Also, dry microbursts typically arise when nearby mountains generate

storms in situations in which storms would not have otherwise occurred because of the dryness of the air near the ground.)

After a quick implementation of software changes to incorporate the reflectivity requirement, the performance of the TDWR system improved substantially: the probability of false alarm (P_{fa}) dropped from 20% to approximately 7% with a small decrease in the probability of detection (P_d) from 95% to 94%. From an air traffic perspective, this improvement was more important than the figures alone suggested because the majority of the remainder of the false alarms occurred in complicated severe storms that would have been avoided by pilots using other information. By contrast, the false alarms that occurred in conditions of nominally clear air significantly undermined the overall credibility of the system.

The Orlando Problem:

When Is a Microburst Hazardous?

The Orlando TDWR test site was the first location where the line of sight (LOS) to the airport was significantly obstructed. Because the testbed used an air-inflated radome (for ease in changing sites), the cost-effective solution was to construct a concrete building about 45 ft high on top of which the antenna pedestal and radome could be located to overcome the LOS obstruction from the surrounding trees (Figure 4). TDWR product testing was conducted at this site



FIGURE 4. TDWR testbed located five miles south of Orlando International Airport.

from the summer of 1990 through the fall of 1993.

The early Orlando testing resulted in a searching reexamination of the following question: at what distance from a microburst divergent region should a wind-shear warning be issued? Initially, the TDWR program used results from the NCAR Classify and Locate Wind Shear (CLAWS) experiment in 1986 in which experienced radar meteorologists analyzed Doppler weather radar data in real time and provided microburst locations to the Denver control tower for transmission to aircraft [4]. On the basis of the CLAWS results, the TDWR criterion was that a microburst alert would be issued to pilots when the declared region for a microburst was within 1/2 nmi of a runway, as indicated in Figure 5. To be particularly conservative, the alert issued would be for the worst-




FIGURE 5. Conservative approach for microburst-warning algorithm: (a) flight approach and (b) flight departure. A microburst alert is issued to pilots when the declared region for a microburst is within 1/2 nmi of a runway. To be particularly conservative, the alert issued is for the worst-case intensity within the microburst, even though the most intense shear might not be along the expected flight path. For the flight approach, the warning issued in this example is "Microburst alert, 1 mile final, 40 knot loss." For the flight departure, the warning issued is "Microburst alert, 2 mile departure, 40 knot loss."

case intensity within the microburst, even though the most intense shear might not be along the expected flight path.

Shortly after testing began, however, it became clear that the previous criteria were too operationally conservative for the Orlando environment [5]. Air traffic controllers felt that the traffic flow was being unduly restricted by the microburst alerts, and pilots started ignoring the warnings. An investigation showed why the criteria were too conservative for Orlando:

1. The Orlando microburst-producing storms typically had heavy rain occurring in an otherwise sunny environment so that controllers and pilots could easily see the storm edge (and blowing rain),
2. The Orlando storms, which are typically slow moving, resulted in overalert situations (e.g., a microburst alert issued for a storm 1/2 nmi from the runway, with minimal or no wind shear along the runway) that persisted for many minutes,
3. Microburst events occurred very frequently (e.g., 1600 events were observed in a single Orlando summer versus approximately 400 in Kansas City), and
4. The Orlando air traffic controllers had much past experience in conducting airport operations safely with thunderstorms close to runways and approach/departure corridors.

Consequently, J. Stillson at Lincoln Laboratory developed a different approach in which the pilot warning intensity takes into account the proximity of the microburst to the flight path, as shown in Figure 6 and Table 1. The algorithm computes the average radial velocity shear across the microburst and then integrates this quantity over the portion of the runway warning region that is overlapped by the estimated microburst extent. Thus the resulting *shear integration* alert will be significantly lower than the strength of the microburst when the microburst is just touching a warning region and will rise to full strength as the microburst moves across the runway.

Using dual-radar Doppler wind analyses to determine the wind along the runway, we assessed the performance of the improved algorithm by comparing

the alerts generated to the wind shear along the expected flight path. As shown in Table 2, the shear-integration algorithm made a significant improvement in reducing the number of warnings without adversely affecting the detection of hazardous wind shear. Operationally, the new algorithm has been well received at Orlando. But the previous microburst-alerting algorithm was deemed operationally adequate at locations including Denver and Kansas City, where there are factors that might call for a conservative approach (e.g., rapidly moving storm cells that are common at those locations). Thus the shear-integration method has been implemented in the operational TDWR system as a site-selectable feature.

Site Selection

The selection of a site for a TDWR system (including the choice of antenna height) can significantly affect the overall system performance. The principal technical objectives are that the site

1. be as close as possible to the extended centerlines of the principal runways that are used during convective weather,
2. be outside the airport at a distance from 8 to 12 nmi from the airport reference point,
3. have a clear line of sight (LOS) along/within the approach/departure corridors and over all the runways down to 60 to 100 m AGL, if possible, to measure the most intense wind velocities in a microburst outflow,
4. be at a location on the opposite side of the




FIGURE 6. Example of shear-integration method for microburst-warning algorithm. In the example, all of the five microbursts have a total velocity change of 20 m/sec. This total velocity change of 20 m/sec is divided by the horizontal extent of each of the microbursts to obtain the average velocity shear (the fourth column in Table 1). The average velocity shear is then multiplied by the horizontal overlap of the microburst and the runway to obtain the shear-integration warnings (the last column of Table 1). If the method from Figure 5 had been used, the warning would have been 20 m/sec for all of the microbursts. For the situation in which the five microbursts occur simultaneously, the shear integration warning would be the worst-case value of 20 m/sec.

Table 1. Alerts for Various Microburst Locations in Figure 6

Microburst Region	Velocity Change (m/sec)	Horizontal Extent of Microburst (nmi)	Average Velocity Shear (m/sec/nmi)	Horizontal Overlap of Microburst and Runway (nmi)	Shear-Integration Warning (m/sec)
a	20	2.0	10	1.0	10
b	20	2.0	10	2.0	20
c	20	2.0	10	0.5	5
d	20	0.5	40	0.5	20
e	20	2.0	10	0.5	5

Table 2. Performance of Baseline TDWR Algorithm versus Shear-Integration Algorithm

	Denver, 1988		Kansas City, 1989		Orlando, 1990	
	Baseline TDWR	Shear Integration	Baseline TDWR	Shear Integration	Baseline TDWR	Shear Integration
Probability that a microburst alert is issued when flight path encounters microburst	63%	59%	100%	100%	100%	98%
Probability that a wind-shear alert is issued when flight path encounters microburst	91%	85%	100%	100%	100%	100%
Probability that a microburst alert is issued when flight-path shear is less than microburst intensity	60%	43%	69%	53%	65%	32%

Note: A *microburst* is defined as wind change greater than or equal to 30 kn. *Wind shear* is defined as wind change greater than or equal to 20 kn but less than 30 kn.

airport from the direction of approach of the bulk of the thunderstorms, and


5. have low ground clutter in the critical area of the hazardous-weather scan sector.

In addition, a candidate site must be environmentally acceptable, be responsive to the concerns of the local population, and, finally, be available for purchase or long-term (twenty year) lease at fair market value.

No single location ever meets all of the above criteria unequivocally and, in fact, a number of the criteria often conflict. For example, high antenna heights can assist in obtaining clear LOS to the microburst-detection region, but may also yield a better view of ground-clutter sources. Therefore, the process of selecting a site involves the evaluating and prioritizing

of a suite of trade-offs for each candidate site that is available from a real estate point of view and acceptable from an environmental standpoint. Additionally, the latest research on wind-shear mechanisms, algorithm features, and testbed operational characteristics should be considered in accomplishing these trade-offs.

Because of Lincoln Laboratory's experience in siting and operating the TDWR testbed at several locations, the FAA asked Lincoln Laboratory to locate sites for about thirty TDWR locations that had not yet been identified. To accomplish this task, Sterling Software, Inc. was contracted to develop a measurement van for determining LOS coverage, RF interference, and local-area ground-clutter values at locations



that the Lincoln Laboratory Weather Sensing Group personnel had determined as best satisfying the criteria described above. Figure 7 shows the measurement van with its pneumatic mast in an up position.

The LOS measurements were obtained by both visible and infrared cameras mounted on a servo-operated optical masthead platform. In addition, sensitive tilt meters connected to computers in the van were used to obtain and record the camera tilt at the time of shutter operation. Each camera frame covered an azimuth of approximately 6° and the frames, when stitched together, produced an image of the topology with a precise overlay of the horizon (Figure 8).

A second masthead platform and a Raytheon Marine Radar system were used to take ground-clutter measurements of the surrounding environment. C-band RF interference measurements were taken continuously at the same time as the LOS measurements. Both sets of measurements were taken at various heights above the ground, from 15 to 30 m at 5-m increments.

The measurement van was used to assess the suitability of about thirty-five locations because, in a number of cases, the FAA real estate offices were not successful in making arrangements to use some of the sites. Also, concerns over radiation effects were a factor at several other sites. Indeed, a recent airplane crash at Charlotte International Airport in North Carolina highlighted the problems, and politics, that can arise in obtaining TDWR sites (*The New York Times*, 25 July 1994, p. A19, and *The Washington Post*, 7 July 1994, p. A3). Although TDWR radiation is typically 50 dB below the ANSI maximum permissible level and less than existing local TV radiation levels, concern over such radiation was often voiced at local information hearings. Lincoln Laboratory pro-



FIGURE 7. Measurement van used to assess the suitability of potential TDWR sites. Shown with its RF masthead installed, the van collected data to determine the LOS coverage, RF interference, and local-area ground-clutter values at the candidate locations for the TDWR systems.

vided technical support to the FAA in addressing these issues at a number of hearings.

Product Refinements to Meet Additional FAA Needs

After the TDWR production contract was issued, additional FAA needs for weather information were identified. These needs were investigated off line with recorded TDWR testbed data and then assessed operationally in real-time tests with the TDWR testbed system at Kansas City and Orlando.




FIGURE 8. Digital horizon profile from combined (visible and infrared) camera images at the 20-m elevation that was designated for a TDWR location near Dallas–Fort Worth. The airport reference point is shown with an X, and elevations above ground level are shown with orange + signs.




FIGURE 9. Sensor locations and runways used for the TDWR/Low Level Wind Shear Alert System (LLWAS) integration study at Orlando. Tables 3 and 4 contain the results of this study.

TDWR/LLWAS Integration

In parallel with the TDWR development activity, the FAA has investigated approaches for improving the existing Low Level Wind Shear Alert System (LLWAS)—an anemometer-based system for wind-shear detection. Designed to detect gust fronts, the initial LLWAS was deficient in detecting microbursts in a number of ways, e.g., the sensor spacing, sensor height, and detection algorithms. Work at NCAR has resulted in an improved LLWAS system, which is scheduled for deployment at seven TDWR-equipped airports to provide complementary wind-shear detection.

LLWAS and TDWR are complementary systems in that LLWAS measures the horizontal winds at the sensor locations whereas TDWR detects radial winds throughout the terminal area as well as storm features aloft that will result in future microbursts. Thus LLWAS must infer winds between the sensor locations, including along runways, while TDWR must infer the nonradial component of the horizontal winds.

Because the LLWAS and TDWR systems take dif-

Table 3. Detection Statistics for Separate Wind-Shear Systems and for Integration Algorithm (from Reference 6)

	<i>TDWR</i>	<i>LLWAS</i>	<i>Integrated Result</i>
Probability that an alert was issued when a microburst alert was warranted	99%	97%	99%
Probability that an alert was issued when some type of alert was warranted	92%	76%	93%
Probability that a microburst alert was issued when warranted	97%	90%	97%
Probability that a microburst alert was issued when no alert was warranted	4%	3%	2%
Probability that a wind-shear alert was issued when no alert was warranted	22%	2%	19%
Probability that a microburst alert was issued when a wind-shear alert was warranted	31%	25%	27%

Note: A *microburst* is defined as wind change greater than or equal to 30 kn. *Wind shear* is defined as wind change greater than or equal to 20 kn but less than 30 kn.

ferent types of measurements and use rather different processing algorithms, situations can arise in which the wind-shear warnings provided by the systems differ. The FAA did not want air traffic controllers to make meteorological judgments as to which system was more accurate whenever TDWR and LLWAS disagreed. Hence the FAA assigned Lincoln Laboratory and NCAR the task of developing an algorithm that would use information from both systems to arrive at a single optimized wind-shear warning.

There were two broad approaches that could have been taken:

1. use data from both TDWR and LLWAS to produce optimized estimates of the overall low-altitude wind field and storm location from which wind-shear warnings could be derived, or
2. use the wind-shear outputs of each system in an optimized fashion to minimize false alarms and obtain a better P_d than either system alone.

Because both systems were already quite reliable individually (e.g., $P_d > 90\%$ and $P_{fa} < 10\%$), we concluded that the complexity and development time associated with approach 1 were not warranted. Hence we focused on approach 2, and R.E. Cole of Lincoln Laboratory developed an integration algorithm [6] that has five key features:

1. warnings are provided with the same formats as TDWR or LLWAS alone, with no indication as to which sensor(s) was responsible for a particular warning,
2. a warning of a strong microburst by either system results in a strong microburst warning (thus maximizing the P_d for strong events),
3. a warning of a weak microburst by one system must be confirmed by the other system before a wind-shear alert warning is issued (thus reducing nuisance alerts),
4. gust-front warnings within the LLWAS coverage region are provided by LLWAS alone (because LLWAS can spatially resolve and detect gust fronts easily), and
5. the magnitude of the wind change is the mean of the TDWR and LLWAS alerts.

The performance of this algorithm was assessed at Orlando. In addition to the existing LLWAS system at that site, a number of Lincoln Laboratory meso-

scale network (mesonet) systems [7] with anemometers on high poles were deployed to create a very large pseudo-airport with a wide variety of runway geometries relative to the TDWR testbed (Figure 9) [6]. In the experiment, the warnings from the integration algorithm were compared with the dual-Doppler estimates of the wind field along the runways and approach corridors for all runway orientations. Tables 3 and 4 show the results of this study. We see that the integration algorithm achieved microburst P_d and P_{fa} values that were at least as good as the values obtained by either system individually. Furthermore, the integration algorithm achieved more accurate microburst warning loss values than the TDWR system alone.

The lack of a suitable anemometer array together with dual-Doppler coverage prevented a similar quantitative evaluation for the other TDWR testbed sites. An evaluation by NCAR using Denver data, however, showed much the same qualitative performance for the integration algorithm. Thus Raytheon is currently implementing the algorithm, which should be operational at a number of airports by the summer of 1995.

Storm Motion

Testing in Denver had shown the operational utility of the TDWR precipitation product in anticipating

Table 4. Accuracy of Wind-Shear Loss Values for Individual Systems and for Integration Algorithm (from Reference 6)

	TDWR	LLWAS	Integrated Result
Measurements within ± 2.5 kn	8%	16%	17%
Measurements within ± 7.5 kn	28%	50%	43%
Measurements within ± 12.5 kn	53%	74%	63%
Measurements within ± 17.5 kn	70%	88%	78%
Median error (kn)	11.7	3.5	8.6

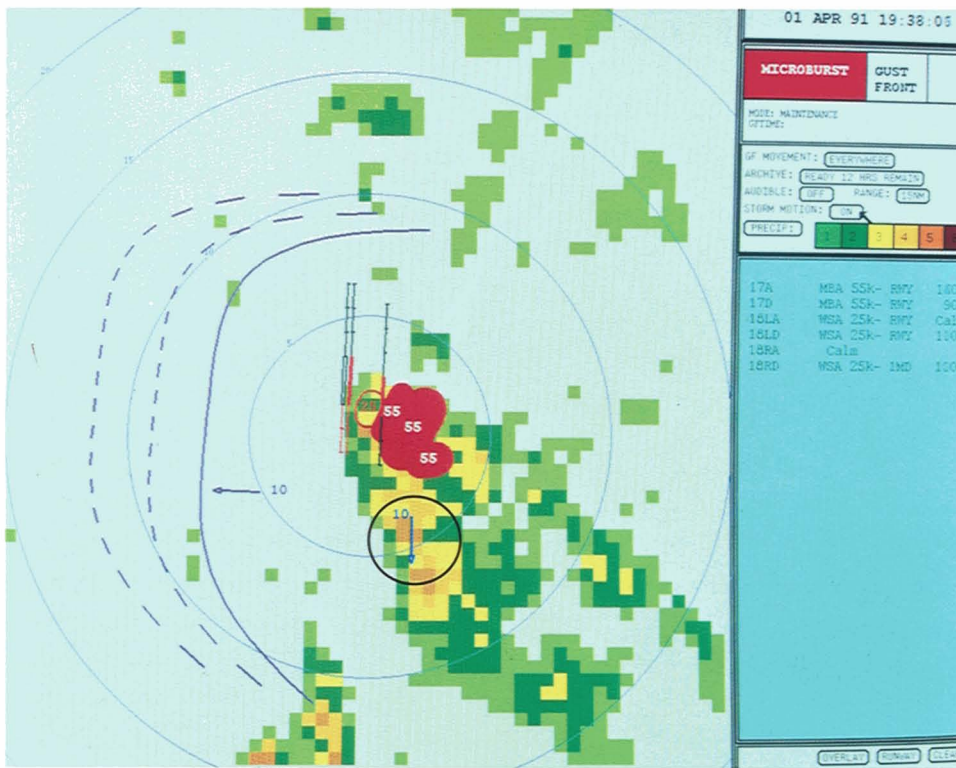


FIGURE 10. Testing of TDWR storm-motion product in Orlando in 1990. Storm motion of 10 kn to the south is indicated by the circled blue vector.

storm impacts on the terminal area. The baseline TDWR product, however, did not provide any information on storm motion to assist in planning.

There were two basic approaches to storm-motion estimation that had been considered in the radar meteorological literature:

1. tracking of storms as three-dimensional (3-D) entities with an emphasis on features that are characteristic of severe storms (e.g., hail and/or tornadic characteristics), and
2. pattern matching of precipitation fields with the spatial correlation function [8–10] to identify best-fit displacement of the weather at one time instant for matching the weather at a later time.

The Next Generation Weather Radar (NEXRAD) system currently uses the first approach. The TDWR scan pattern, however, is such that 3-D information is available only over approximately one-third of the terminal area. Furthermore, the 3-D NEXRAD trackers have to date performed poorly on nonsevere multicell storm-cell clusters because of the merging and splitting of the cells over time. Such storm-cell

clusters are far more common than severe storms and are of concern to aviation even though they generally do not produce significant damage on the ground.

Consequently, Lincoln Laboratory has focused on methods that use the spatial correlation function to estimate storm motion (see the article “Automated Storm Tracking for Terminal Air Traffic Control” in this issue [9]). Figure 10 shows an example of the TDWR storm-motion product for a storm observed in Orlando. The storm motion is indicated by an arrow with the speed in knots shown adjacent to the arrow.

Although there are some basic deficiencies in the storm representation by the TDWR precipitation product (see the article “The Integrated Terminal Weather System (ITWS)” in this issue [11]), terminal air traffic personnel have found the TDWR-based storm-motion product to be very useful. In fact, the FAA has issued a formal requirement for the product’s implementation. As described below, a new approach was taken to assist Raytheon in the software implementation of the storm-motion prod-

uct. The approach was quite successful, and the product is expected to be implemented at operational TDWR sites in 1995.

Data-Link Products

Although pilots are the end recipients of the TDWR wind-shear warnings, the initial TDWR operational concept provided these warnings only via verbal VHF radio messages from tower controllers. To augment the tower-controller messages, data-link transmission of wind-shear information could be used to make a pilot aware of the wind-shear activity (and precipitation) prior to the pilot's tuning to the tower frequency.

The long-standing plan was to provide the ribbon display messages to pilots via the Mode S data link. Delays in the Mode S program and related data-link processing programs, however, prevented any operational evaluation of TDWR data-link products from 1988 through 1992.

In the summer of 1992, Lincoln Laboratory personnel had the opportunity to ride as cockpit passengers on Northwest Airlines flights, during which they observed that pilots were attempting to obtain airport weather information by using the data-link text transmission capability of the Aircraft Communications Addressing and Reporting System (ACARS) to gain access to National Weather Service (NWS) standard airport observations. From this experience, it was clear that automatically generated text messages based on TDWR would be much more timely and germane because the NWS observers had no access to the TDWR information.

Thus an ad hoc users' group was formed with representatives from air traffic controllers and pilots from the major airlines. Because the ACARS text display unit could display only 10 lines of 21 characters each, prioritization of the information was essential. Additionally, air traffic controllers were concerned about the increase in their work load if the text messages were misleading.

During operational tests at Orlando in 1993 [12], pilots found the automatically generated text messages to be operationally useful and effective in reducing pilot-controller discussions regarding the weather. Although using a data link to provide TDWR informa-

tion had been a part of the overall FAA system architecture throughout the TDWR development, the 1993 test provided some very useful insights:

1. Most pilots access the data-link information at least 10 min prior to landing (when their work load is lower). Hence the planning elements of TDWR (e.g., information on wind-shear activity and precipitation location/motion near the airport) are very important.
2. Providing pilots all details of the current wind-shear warnings that are available to the tower controllers may be misleading because the microbursts will evolve and/or move significantly in the time between the receiving of a data-link message and the encountering of a microburst.
3. The greatest benefit in reducing controller work load occurs in inclement weather situations (i.e., when holding patterns are necessary) because pilots can easily get frequent updates from the data link without having to call an overloaded controller.
4. Airline flight operations personnel (e.g., dispatchers) could use the text message to obtain timely information on conditions at the airport.

Currently, the FAA data-link program is working with the TDWR program to implement the data-link products operationally.

Improved Gust-Front Detection/Wind-Shift Prediction

The wind shear that occurs in a typical gust-front encounter tends to increase the lift on a plane; i.e., a plane penetrating a convergence boundary will encounter increased headwinds. Consequently, there was no quantitative performance requirement for gust-front detection/wind-shift prediction for the initial TDWR deployment.

The operational evaluations with the TDWR tested, however, showed that the wind-shift product had operational benefits for airport planning. For example, an aircraft having to change runways at a major airport can result in a 15-to-20-min loss of airport capacity if the runway change occurs unexpectedly. Also, several incidents occurred in which fronts aligned along the flight path produced very turbulent conditions even though there was no appreciable net change in headwind.

The initial TDWR gust-front algorithm, developed by the National Oceanic and Atmospheric Administration (NOAA) National Severe Storms Laboratory (NSSL), used only a convergent radial velocity feature to detect and delineate gust fronts. Such features are very difficult to observe when the leading edge of a gust front is aligned along a radar beam because the velocity change in those situations will be largely transverse to the radar LOS. This condition occurred particularly at Kansas City and Orlando, where the TDWR was located south of the airport with gust fronts moving either to the west or east. The result was that the probability of detecting a given portion of a gust front near the airport could be well below 50%, and the ability to predict wind shifts was also low.

The radar meteorology community had widely recognized that, in addition to the convergent radial velocity feature, a number of other Doppler radar image characteristics could be used, including

1. thin lines in reflectivity images,
2. regions of high variability in the radial velocity within a range gate,
3. azimuthal wind shear, and
4. feature movement.

Additionally, the past motion of a gust front could be used to aid in resolving ambiguous situations. The challenge was to find an effective algorithmic approach for utilizing all of these various clues while minimizing false alarms triggered by erroneous data mimicking a gust front. (For example, range-aliased echoes from a distant storm appear to be elongated in range when represented as first-trip returns. The elongated shape can look like a gust front aligned along the radar beam.)

The first effort in this direction used a rule-based expert system similar to that used for microburst detection [13]. The resulting rule set, however, was quite complex and difficult to optimize.

The use of functional templates as described by R.L. Delanoy and S.W. Troxel [14] provides a more robust algorithm. With this approach, the fusion of multiple features is achieved by the pixel-by-pixel weighted addition of *interest* levels corresponding to the confidence in the existence of each feature. The possibility of various error sources in the data emulat-

ing a feature is addressed by the building of feature detectors for the error sources. The interest associated with these error sources has a negative value; thus the net interest for a pixel is the weighted algebraic sum of various feature-detector outputs. For detecting gust fronts, interest image features can be manipulated, e.g., to look for the motion of line-like features. Overall, the use of functional templates has been very successful in improving the detection of gust fronts. In Orlando, the probability of detection increased from 30% to 71%, with the false-alarm probability decreasing from 8% to 3.5%.

Algorithm Transition to Contractors

Practical operational use of the weather-detection algorithms developed by Lincoln Laboratory requires implementation of the algorithms in the production system software. The production software must meet stringent government requirements for the design, coding, and testing process. Furthermore, the software must also be written in a single high-level language. As discussed earlier, the Lincoln Laboratory software that was used to develop, test, and demonstrate the products operationally met few, if any, of the government acceptance criteria for production software. Thus technology had to be transferred from Lincoln Laboratory to the organization responsible for the production software. (Note: Although Raytheon Corp. was in charge of the production TDWR software, Lincoln Laboratory also worked with the Unisys Corp. on technology transfer in the context of the terminal NEXRAD program, which was eventually canceled, and with NOAA's Forecast Systems Laboratory on the validation of the TDWR algorithm specification.)

Contractor implementation, design, and testing of the initial TDWR product algorithms from a Lincoln Laboratory-generated specification revealed a number of areas for improvement:

1. In some cases, the software designers at the contractor were not familiar with the implementation of data structures as called for in the specification; e.g., the designers had difficulty attempting to implement an algorithm in FORTRAN that had been developed by Lincoln Laboratory in a LISP dialect.

2. Determining the cause of differences between Lincoln Laboratory's and the contractor's results for a given algorithm on a given input data set was difficult because there were insufficient common test points between the two software implementations.
3. The contractor handled the error conditions and resolved the ambiguous situations differently from Lincoln Laboratory because of misunderstandings that arose concerning the rationale for some elements of the specification.

The need to transfer the new and refined algorithms described in the previous section offered Lincoln Laboratory the opportunity to investigate alternative approaches to the technology transfer. The first of these approaches was to work more closely with the contractor personnel in developing the specification. A major element of the government-prescribed procedure for software development is the preparation of a contractor specification of the algorithm that goes beyond the government (i.e., Lincoln Laboratory) specification to include the following tasks: accessing the data required, providing the results to the display systems, and handling various error conditions/contingencies, e.g., processor overload. Lincoln Laboratory worked with the contractor to create a revised Lincoln Laboratory specification that would accomplish the same functional capability as the original specification while facilitating the contractor software validation and testing process.

The second element was to provide the contractor with a copy of the Lincoln Laboratory algorithm software that had been used for product operational validation. This software had been restructured to facilitate its utility for the contractor. For example, Lincoln Laboratory-specific input/output calling routines had been replaced with more generic calls, and comments had been added to indicate the relationship of the code to the specification features. We found that the copy of the working code helped the contractor to reduce the design effort significantly. For example, Lincoln Laboratory code modules could be run on the production system computer to give a rough assessment of the areas of highest computational load that might require greater design effort.

Finally, the Lincoln Laboratory algorithm develop-

ers were available as consultants to the contractor software designers and test team to explain elements of the algorithm and to assist in resolving differences in the results obtained by the Lincoln Laboratory software and the contractor software. When appropriate, the Raytheon designers would work with the algorithm developers and computer system at Lincoln Laboratory to address specific issues. In some cases in which the Lincoln Laboratory code needed to be revised significantly because of its age, the software approach developed by the contractor was incorporated into the Lincoln Laboratory system.

Much of the initial development of the Raytheon specification and preliminary software design was accomplished under a separate contract between Lincoln Laboratory and Raytheon that did not explicitly involve the FAA TDWR contract with Raytheon. This arrangement had two advantages:

1. The contractual arrangement mitigated concerns of government procurement officials that the close working relationship between a research organization and a contractor such as described above could lead to situations in which the research organization provides unapproved technical direction that results in increased costs to the government for the production contract.
2. Industrial firms are generally risk adverse to the incurring of losses on a fixed-price contract. As a consequence, if there is considerable uncertainty as to the worst-case cost that might occur, the firms will tend to bid a higher price than on a similar job that they understand better. When Raytheon and the government met to discuss the full formal implementation of the software, Raytheon had a very good idea as to the amount of effort that would be required for the job. This reduction in Raytheon's uncertainty as to the magnitude of the work enabled the company to have much greater confidence in its ability to estimate the cost. Hence Raytheon could quote a lower cost to the government.

The flexible nature of this contractual relationship has also made it much easier to respond in a coordinated way to other unexpected issues that have arisen with other elements of the program, e.g., the develop-




FIGURE 11. Distribution of microburst strengths from various TDWR testbed sites. The total number of microbursts are as follows: Huntsville (1986), 236 microbursts; Denver (1987), 472 microbursts; Denver (1988), 694 microbursts; Kansas City (1989), 318 microbursts; Orlando (1 January through 15 August 1991), 1243 microbursts; and Orlando (1992), 1663 microbursts.

ment of improved recording systems for the TDWR to accomplish site-specific studies.

Overall, this multifaceted approach to technology transfer has worked very well. The overall transition of the TDWR/LLWAS integration and storm-motion algorithms has gone considerably faster and at lower cost than the transition of the initial TDWR algorithm packages to Raytheon.

Summary and Future Work

Since the formal demonstration of TDWR in Denver in 1988, a number of significant accomplishments have occurred during the transition of the system to full-scale development:

1. The microburst-detection and clutter-suppression algorithms that Lincoln Laboratory devel-

- oped were transferred successfully to Raytheon,
2. Major functional interfaces between the contractor features and the government algorithms were validated experimentally and refined with the Lincoln Laboratory–developed TDWR testbed, and
3. A number of system refinements were made to address the site-specific problems that were identified in the Lincoln Laboratory TDWR testbed experiments at Kansas City and Orlando.

Lincoln Laboratory is continuing to provide technical support to the FAA in determining site locations for the TDWR at different airports. The laboratory is also acting as a consultant on various system issues. Additionally, Lincoln Laboratory has developed new

products that further improve the wind-shear warnings and increase the utility of the TDWR for terminal planning.

The principal objective of the eleven-year TDWR program at Lincoln Laboratory has been to improve the safety and operations at major airports. Since 1988, testbed operations have provided wind-shear protection at various major airports while the production systems were being developed. As shown in Figure 11, several thousand microbursts were observed during this period, from which data are available for future studies. Currently, commissioned production TDWR systems are providing operational TDWR wind-shear warnings at Memphis International Airport (site of the first Lincoln Laboratory testbed experiments), St. Louis Airport, and Houston Intercontinental Airport, and approximately fourteen other sites are testing a production TDWR as a prelude to formal commissioning.


In the next two years, research at Lincoln Laboratory will focus on the recording and analysis of data from a number of key TDWR sites to verify the site-specific system operation and to develop a methodology for ongoing system analysis and optimization. The technology developed from this effort will be transferred to the FAA TDWR Program Support Facility in Oklahoma City, which will be responsible for long-term support of the deployed TDWR systems.

al Test and Evaluation, Orlando 1990," MIT Lincoln Laboratory, *Project Report ATC-179*, DOT/FAA/NR-90/2 (9 Apr. 1991).


6. R.E. Cole and R.F. Todd, "Terminal Doppler Weather Radar (TDWR) Low Level Wind Shear Alert System 3 (LLWAS 3) Integration Studies at Orlando International Airport in 1991 and 1992," MIT Lincoln Laboratory, *Project Report ATC-216*, DOT/FAA/RD-94/12 (20 Apr. 1994).
7. M.M. Wolfson, R.L. Delaney, B.E. Forman, R.G. Hollowell, M.L. Pawlak, and P.D. Smith, "Automated Microburst Wind-Shear Prediction," *Linc. Lab. J.*, in this issue.
8. J.C. Brasunas, "A Comparison of Storm Tracking and Extrapolation Algorithms," MIT Lincoln Laboratory, *Project Report ATC-124*, DOT/FAA/PM 84-1 (31 July 1984).
9. E.S. Chornoboy, A.M. Matlin, and J.P. Morgan, "Automated Storm Tracking for Terminal Air Traffic Control," *Linc. Lab. J.*, in this issue.
10. E.S. Chornoboy, "Storm Tracking for TDWR: A Correlation Algorithm Design and Evaluation," MIT Lincoln Laboratory, *Project Report ATC-182* (14 July 1992).
11. J.E. Evans and E.R. Ducot, "The Integrated Terminal Weather System (ITWS)," *Linc. Lab. J.*, in this issue.
12. S.D. Campbell, "Terminal Weather Message Demonstration at Orlando, FL," MIT Lincoln Laboratory, *Project Report ATC-210*, DOT/FAA/RD-94/3 (4 Mar. 1994).
13. S.D. Smith, A. Witt, M. Eilts, D. Klingle-Wilson, S. Olson, J. Sanford, "Gust Front Detection Algorithm for the Terminal Doppler Weather Radar: Part I, Current Status," *3rd Int. Conf. on Aviation Weather Systems, Anaheim, CA, 30 Jan.-3 Feb. 1989*, American Meteorological Society, p. 31.
14. R.L. Delaney and S.W. Troxel, "Machine Intelligent Gust Front Detection" *Linc. Lab. J.* 6, 187 (1993).

REFERENCES

1. M.W. Merritt, D. Klingle-Wilson, and S.D. Campbell, "Wind Shear Detection with Pencil-Beam Radars," *Linc. Lab. J.* 2, 483 (1989).
2. D. Turnbull, J. McCarthy, J. Evans, and D. Zrnic, "The FAA Terminal Doppler Weather Radar (TDWR) Program," *3rd Int. Conf. on Aviation Weather Systems, Anaheim, CA, 30 Jan.-3 Feb. 1989*, American Meteorological Society, p. 414.
3. J.E. Evans, "Results of the Kansas City 1989 Terminal Doppler Weather Radar (TDWR) Operational Evaluation Testing," MIT Lincoln Laboratory, *Project Report ATC-171*, DOT/FAA/NR-90/1 (17 Aug. 1990).
4. L. Stevenson, "The Stapleton Microburst Advisory Service Project: An Operational Viewpoint," U.S. Dept. of Transportation, Federal Aviation Administration, DOT/FAA/PM-85/21 (Sept. 1985).
5. D.M. Bernella, "Terminal Doppler Weather Radar Operation-



JAMES E. EVANS is leader of the Weather Sensing Group, where his current focus of research is on addressing weather impacts on aviation and everyday life. Jim received an S.B., an S.M., and a Ph.D. degree from MIT in electrical engineering. While at MIT, he received the Compton Award and the Carleton E. Tucker Teaching Award.



DAVID M. BERNELLA is assistant leader of the Weather Sensing Group, where he has been responsible for the development and operation of a weather radar testbed for the FAA Terminal Doppler Weather Radar program. He received a B.S. degree in engineering science from Pennsylvania State University in 1958 and an M.S. degree in electrical engineering from Northeastern University in 1966. While David was at HRB Singer (1958–1961), the company received several device patents for his work on Remote Weather Transponders.