

**Project Report
WXCB-2**

**C-Band Aerosol Release Detector (CBARD)
Algorithms Description**

**S.W. Troxel
J. Cho
R. Frankel
M.P. Matthews
J. Shaw**

19 January 2006

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Lexington, Massachusetts



**Prepared for the Joint Program Manager for Nuclear, Biological, and Chemical
Contamination Avoidance (JPM NBC CA) under Air Force Contract FA8721-05-C-0002.**

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

ABSTRACT

The C-Band Aerosol Release Detector (CBARD) system developed under the Homeland Defense Chemical Biological Umbrella (HDCBU) program utilizes existing FAA airport surveillance and weather radars to detect aerosol releases from crop-duster type aircraft and provide early warning and advisory capabilities for emergency response agencies.

The CBARD system utilizes sophisticated radar image processing and scan control algorithms to maintain a tight track on a possible threat aircraft and to recognize signatures in Terminal Doppler Weather Radar (TDWR) data indicative of the aircraft target and aerosol plume. This report describes the Aerosol Release Detection Algorithm (ARDA) and 3-D Aircraft Tracking and TDWR Scan Strategy Algorithm (3-DAT), and the scientific underpinnings for these.

TABLE OF CONTENTS

	Page
ABSTRACT	iii
List of Illustrations	vii
List of Tables	xi
1. INTRODUCTION	1
2. SUPPORTING EXPERIMENTS AND STUDIES	3
2.1 Particle Size Measurements	3
2.2 Canadian River Tests	6
3. CBARD SYSTEM OVERVIEW	9
4. AIRCRAFT AND PLUME DETECTION ALGORITHM DESCRIPTION	11
4.1 ARDA Data Flow	11
4.2 Radar Signatures	12
4.3 ARDA Processing Overview	15
4.4 Data Preparation	18
4.5 Aircraft Detection	23
4.6 Plume Detection	33
4.7 Status Message Output	49
5. 3-DAT ALGORITHM DESCRIPTION	51
5.1 Search Mode	51
5.2 Track Mode	51
6. RESULTS	55
7. SUMMARY	59
Glossary	61

TABLE OF CONTENTS
(Continued)

	Page
References	63

LIST OF ILLUSTRATIONS

Figure No.		Page
1	Number density versus diameter for simulants measured at Dugway.	4
2	Plume reflectivity distribution for BT and ovalbumin observed by the Enterprise Electronics Corporation C-band radar at Dugway.	4
3	Distribution of plume detectability versus clear air background reflectivity.	7
4	Seasonal and diurnal changes in background reflectivity.	7
5	CBARD functional architecture.	10
6	ARDA data flow.	11
7	Reflectivity peak and differential Doppler velocity associated with aircraft target.	13
8	Cumulative reflectivity difference motion signature associated with an aircraft.	14
9	Crop duster release plume signatures in TDWR data.	15
10	Example functional template for plume reflectivity thin line feature detection.	16
11	ARDA functional block diagram.	17
12	Effects of Cartesian resampling grid resolution.	19
13	Example of background reflectivity normalization.	20
14	Summary of image pre-processing stages for determination of aircraft motion by reflectivity differencing.	22
15	Illustration of Aircraft Reflectivity Peak feature detector.	24
16	Computation of differential velocity image.	25
17	Illustration of Aircraft Tandem reflectivity detector.	26
18	Illustration of Aircraft Reflectivity Difference feature detector.	27

LIST OF ILLUSTRATIONS (Continued)

Figure No.		Page
19	Illustration of Aircraft Motion feature detector.	28
20	Aircraft Anticipation interest images.	29
21	Aircraft interest image fusion.	30
22	Results of feature extraction.	31
23	Summary of Initial Plume detector processing.	35
24	Example results of Thin Line Suppression interest detection.	37
25	Examples of reflectivity images containing large amounts of atmospheric structure and/or ground clutter.	38
26	Illustration of background reflectivity sensitivity bias computation.	39
27	Example of Plume Anticipation interest image.	40
28	Illustration of Plume Reflectivity detector processing.	41
29	Functional template for Frame-Integrated Plume detector.	43
30	Illustration of Frame-Integrated Plume detector processing stages.	43
31	Summary of Plume Reflectivity Difference detector processing.	44
32	Example of multi-path detection.	45
33	Plume Anticipation Averaging Weight Function.	46
34	Illustration of “M-of-N” alert reporting strategy.	48
35	Performance confidence matrix.	49
36	Example of aircraft tracking results.	53

LIST OF ILLUSTRATIONS
(Continued)

Figure No.		Page
37	ARDA diagnostic display showing example plume detection.	55

LIST OF TABLES

Table No.		Page
1	Aircraft Feature Detectors	23
2	Feature Detector Weights Used for Aircraft Interest Fusion	30
3	Plume Feature Detectors	34
4	Feature Detector Weights Used for Plume Interest Fusion	46
5	Plume Detection Performance Results	57

1. INTRODUCTION

Under the Joint Program Executive Office for Chemical Biological Defense, the Homeland Defense Chemical Biological Umbrella (HDCBU) program was established to address the threat of chemical and biological (CB) weapons to the United States through development of an early warning system utilizing the existing network of federal and civilian radar systems. The HDCBU project sponsored a series of field tests and analyses to determine the potential of various radar systems to detect aerosol clouds disseminated from small crop-duster style aircraft and determined that the FAA Terminal Doppler Weather Radar (TDWR) was the best candidate for development of the early warning capability. Under the HDCBU project, Lincoln Laboratory, in conjunction with the FAA, developed the necessary interfaces and algorithm software components for the initial capability known as C-Band Aerosol Release Detector (CBARD). Several additional field tests using the FAA radar test facility in Oklahoma City were conducted during the course of CBARD development and provided the necessary data for algorithm development. The system was successfully validated in an operational demonstration conducted at Oklahoma City in August, 2004.

Within CBARD, the Aerosol Release Detection Algorithm (ARDA) is the software module responsible for the 3-D detection of the threat aircraft and identification of aerosol plume signatures. More specifically, ARDA serves two functions within CBARD:

1. Detect and track the cued threat aircraft in the multi-elevation TDWR data and ascertain the aircraft altitude. This altitude information augments the 2-D aircraft cue position data provided by the CBARD ASR-9 aircraft tracker and is a critical datum needed to provide feedback to the 3-D Aircraft Tracking and TDWR scan strategy algorithm (3-DAT) in order to maintain TDWR antenna positioning on and near the threat aircraft.
2. Process the TDWR data to automatically identify possible aerosol release events and produce a detection message with information necessary to produce an NBC-1 warning message that can be sent to Emergency Operations Centers (EOC).

This report focuses primarily on the scientific underpinnings and implementation of the aerosol plume detection algorithm (ARDA), but also discusses the workings of the 3-DAT algorithm since the two algorithms work synergistically to maintain tracking and detection of the aircraft and associated aerosol plumes with the TDWR.

2. SUPPORTING EXPERIMENTS AND STUDIES

2.1 PARTICLE SIZE MEASUREMENTS

C-band radars such as the TDWR operate at wavelengths of about 5 cm. The particle size of chemical and biological weapons agents are expected to be much smaller than this. Therefore, the scattering of the radar pulse by these small aerosols is in the Rayleigh regime, i.e., the backscattered power is proportional to D^6 , where D is the diameter of the particle. Such a steep dependency of aerosol reflectivity on size means that the size distribution is a key parameter in determining whether it is possible for the radar to detect chemical and biological agent dissemination.

Although a particle sizing experiment was conducted prior to Lincoln Laboratory's involvement in this project (Waxahachie, Texas; December 2002), the instruments used were only capable of measuring up to 0.35 mm in size. However, visual evidence on the instrument table and personnel clothing at this test, and indirect evidence from the reflectivity and inferred fall speeds of simulant plumes in radar data during the Canadian River trials pointed to the existence of larger particles. In an effort to fill in this absence of data in the upper size regime, we leased an optical disdrometer for supplementary particle sizing experiments. Details of the experiment are reported by Cho et al. (2003). Here we briefly summarize the results.

The laser disdrometer (Parsivel M300), normally used to observe raindrop size distributions, had measurement limits of 0.25–25 mm ($\pm 5\%$ accuracy) for size and 0.1–20 m/s ($\pm 10\%$ accuracy) for velocity. For our experiment we used the minimum available integration time of 1 minute.

The measurements were made with the M300 set up in the middle of a clear, flat field, with a crop duster flying 10–15 ft above the ground and releasing the simulant over the instrument. An attendant with a laptop PC connected to the instrument monitored the proceedings from nearby. The experiment was conducted at the Dugway Proving Ground in Utah (28 October 2003) and at an airfield in Waxahachie, Texas (21 November 2003). At the first location, bacillus thuringiensis (BT) and ovalbumin (powdered egg white) were measured. At the second location, a polyethylene glycol/water mixture (PEG) and kaolin were tested.

In Figure 1 we show the particle distribution data for BT and ovalbumin. Clearly, there were particles larger than 0.35 mm and, in the case of ovalbumin, even as large as a few millimeters. The latter measurement is consistent with direct visual observations of pebble-sized aggregates present in the ovalbumin dispersion from the aircraft during the experiment. The clear difference in size distribution between BT and ovalbumin is also manifest in the large radar reflectivity difference (Figure 2). This correlation underscores the importance of the upper-end size distribution tail in determining the radar reflectivity of the dispersant. PEG and kaolin had distributions similar to BT, i.e., diameters of up to ~1 mm were observed.

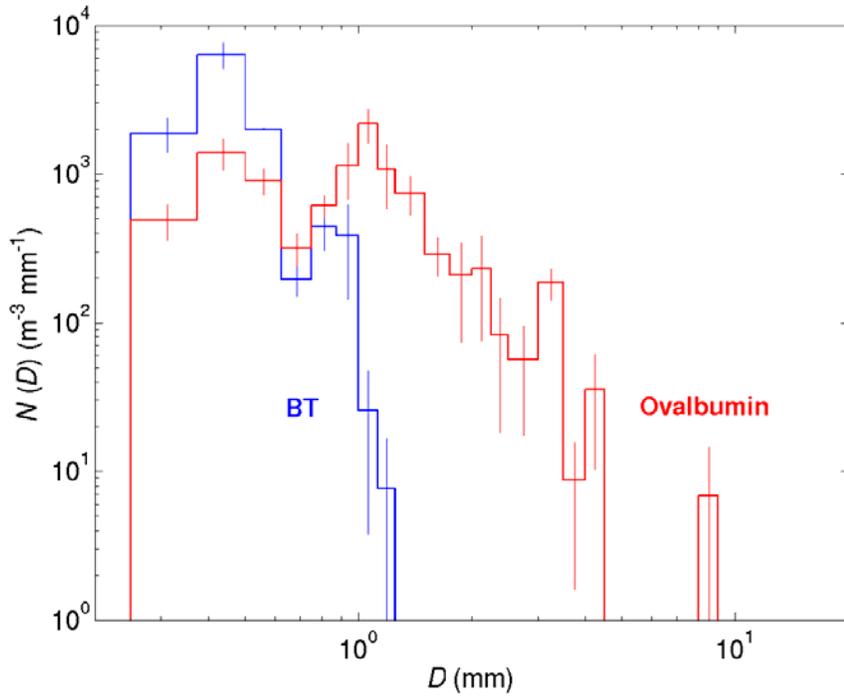


Figure 1. Number density versus diameter for simulants measured at Dugway. The error bars indicate ± 1 standard deviation divided by the square root of $n-1$, where n is the number of integration periods averaged.

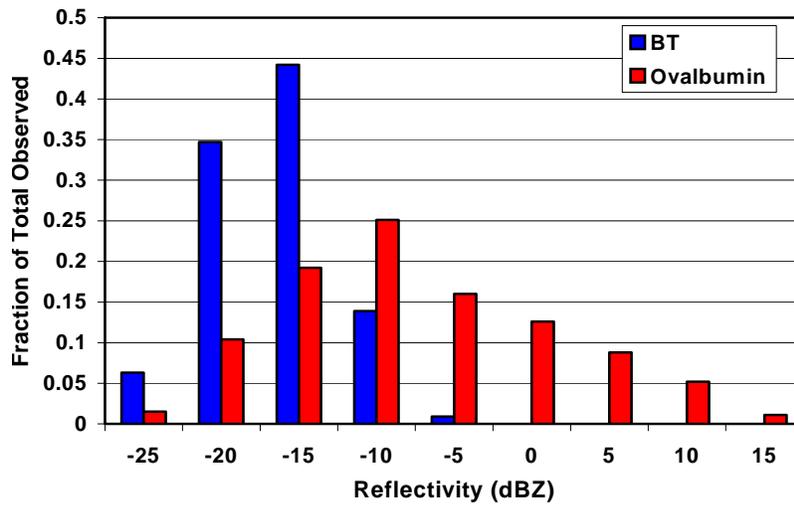


Figure 2. Plume reflectivity distribution for BT and ovalbumin observed by the Enterprise Electronics Corporation C-band radar at Dugway.

The different clumping characteristics of the solid simulants imply that perhaps storage conditions, ambient humidity and temperature, may play a role in determining the detectability of the dispersant. Advanced aerosol preparation technology, such as the incorporation of nanoparticles and surface polymerization for aggregate inhibition, and static charging for mutual repulsion, would lower the likelihood of successful detection by radar. For dry bioweapons agents meant to be inhaled, the real threat is in the low end of the size distribution, so we are relying on a detection of a proxy, i.e., inadvertent aggregates produced by low-tech methodology. For liquid chemical agents such as mustard gas this is not true, since the goal is for the droplets to fall intact to the surface, thus, they are meant to extend to larger sizes. For liquid sprayers, the specific nozzle type may also affect the droplet size distribution, a factor that has not been explored in the experiments conducted during this project.

Because the low-end size distribution was not measurable by the disdrometer and the previous measurements made of the smaller particles by other instruments only provided relative size distributions, we do not have directly measured distributions of the simulants over the entire size range. However, because different size particles fall at different speeds, the time evolution of the radar reflectivity versus beam elevation angle can provide information about the size distribution.

Comparison of radar-observed time evolution of water plume reflectivity with computer simulations incorporating sedimentation and bulk advection effects show a good match in initial reflectivity and transfer rate to lower elevations for a lognormal distribution with count median diameter (CMD) of 0.1 mm and geometric standard deviation (GSD) of 1.6. These size distribution parameters are consistent with the upper-end size measurements made with the Parsivel M300 of PEG. Although the simulated overall decay of reflectivity with time is slower than observed, this is likely due to other physical mechanisms such as particle break-up, evaporation, and turbulent/shear dispersion that were not incorporated into the computer model. See Cho et al. (2003) for further details.

2.2 CANADIAN RIVER TESTS

From April 2003 to July 2004, a series of eight field tests were conducted near the FAA's Mike Monroney Aeronautical Center in Oklahoma City. During these tests, a crop duster disseminated chemical simulants with varying flight paths and release locations near the Canadian River area southwest of the city, and data were collected with the TDWR. The tests were named CR1 through CR8 accordingly.

The objective of the early CR tests (CR1 and CR2) was to determine the ability of Doppler weather radar such as TDWR to detect CB simulants such as Kaolin, PEG, and water dispersed from a crop duster. Drops were made at varying release rates and the TDWR data were examined for radar signatures. Detectability was studied in relation to meteorological variables. An important finding of these early studies was the dependency of detectability on clear-air background returns. Over 50% of the high release rate drops ($> 50 \text{ g/m}$) of PEG and Kaolin were detected when the clear-air background was considered "low" ($\leq -22 \text{ dBZ}$) or "medium" ($-22 < \text{dBZ} \leq -14$), while the detection rate of "high" ($> -14 \text{ dBZ}$) background releases was less than 5%. Figure 3 plots the fraction of total detections against 5 dB increments of background reflectivity.

Variability in the background returns was also found to reduce detectability as the plume signature could be lost amongst the structure in the returns. Additional follow-on studies using existing TDWR data from other locations such as Newark, NJ and Dallas, TX found significant seasonal and diurnal fluctuations in clear-air background levels, with the lowest background levels occurring near sunrise and sunset and during winter months when atmospheric conditions were more stable and biological activity (e.g., birds, insects) was at a relative minimum (Figure 4).

The primary mission of CR3 in November, 2004 was the collection of long release plume signatures in the radar basedata. Many excellent signatures were obtained during the trials at dissemination lengths of 2 km, 6 km, and 10 km. Most of the flights took place at higher altitudes (5000 ft) where the background reflectivity values were considerably lower than at 1200 ft. The plume signatures were useful for the initial detection algorithm development.

The CR4 trials consisted of pre-planned flight scenarios to test the initial algorithms for TDWR scan strategy generation and aircraft tracking, and to understand the limitations. The initial algorithm selected the appropriate scan azimuth sector limits based on the location information provided by the ASR-9, but did not yet determine the aircraft altitude by analyzing the TDWR basedata. The elevation angle scanned by the TDWR during the CR4 trials was determined manually by using the aircraft beacon altitude reports.

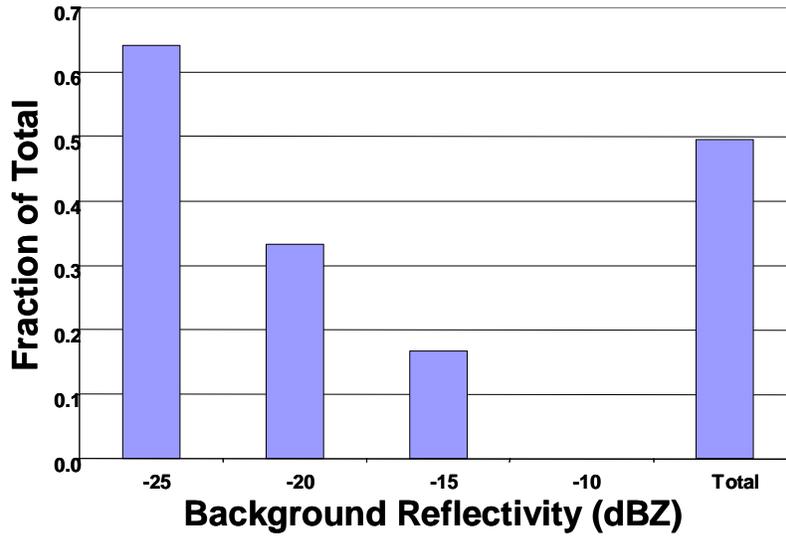


Figure 3. Distribution of plume detectability versus clear air background reflectivity.

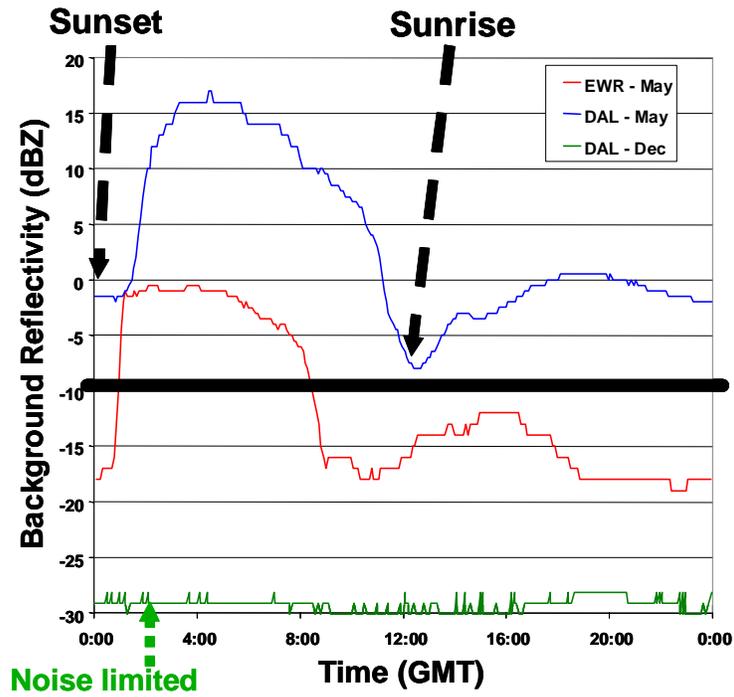


Figure 4. Seasonal and diurnal changes in background reflectivity. Data are from the Newark, NJ (EWR) and Dallas-Love (DAL) TDWRs.

The first end-to-end tests of the CBARD system's major components were performed during CR5 in February, 2004. Tracking algorithm performance was evaluated by flying the same scenarios and flight paths used in CR4. The tracking algorithm calculated the sector azimuth angles using the aircraft location, speed, and heading from the ASR-9, and computed the necessary elevation angles using aircraft detections provided by automated processing of the TDWR basedata (without reliance on aircraft beacon altitude data). Several disseminations were successfully detected and the aircraft tracking was markedly improved compared to CR4. A set of blind trials were conducted at the end of CR5, where no up-front information was given on the location or length of release. A post-test performance analysis revealed that only one out of six "hopper dump" releases during the blind trials was successfully detected, as these produced radar signatures much different in appearance than the low reflectivity thin line echo produced by a more gradual release. Generally, it was felt the additional parameter tuning and algorithm enhancements could further improve performance.

CR6 and CR7 were conducted in March and May of 2004 and were demonstrations of all the major components of the CBARD system. A new method of aircraft tracking using only two elevation tilts (as opposed to four) was employed along with refinements to the aerosol plume detection algorithm. The collection and evaluation of additional data led to the creation of a 3-tiered performance confidence matrix that takes into account the environmental background conditions that exist during a trial. The statistics indicated very good performance in the "green" and "yellow" confidence categories where the mean and standard deviation of the background reflectivity values are low, with decreasing performance as these values increased into the "red" category.

Finally, the CR8 tests conducted in July, 2004 were intended as a sort of "dress rehearsal" for the operational demonstration that would follow the next month. All elements of the CBARD system were exercised. Equipment problems resulted in only ten trials being completed during CR8, and adverse background conditions during the majority tests made detection of the releases difficult. Higher altitude flights between 4000 and 5000 feet were conducted in an attempt to obtain lower background reflectivities, but the resulting high-elevation TDWR scan patterns were sub-optimal for scanning the plume, which further hampered detectability. One of the conclusions from the CR8 tests was that more realistic, lower altitude flight scenarios should be utilized for the upcoming operational demonstration.

Overall, the series of CR tests provided valuable data sets needed to develop, test, and refine the CBARD system.

3. CBARD SYSTEM OVERVIEW

Figure 5 illustrates the functional flow through the various components comprising the CBARD system. A detailed description of the CBARD system design can be found in Elkin et al., 2005.

ASR-9 primary and beacon surveillance data are collected by the CBARD computer, which provides the graphical operator interface to the CBARD system and computes aircraft tracks from the ASR-9 surveillance data using the Lincoln Laboratory Optimized Tracker (LLOT). Once the CBARD operator has selected a possible threat aircraft on the tracker display, the CBARD computer begins outputting 2-D (range, azimuth) cue locations of the threat aircraft. These 2-D aircraft cue locations are used by the 3-D Aircraft Tracking and TDWR Scan Strategy (3-DAT) algorithm to set up the initial volume scan strategy for acquiring the aircraft with the TDWR.

The 2-D aircraft cue locations are also used by Aerosol Release Detection Algorithm (ARDA) to localize the search for the aircraft target and aerosol plume signatures in the TDWR data. If ARDA detects the aircraft signature, it computes a 3-D aircraft location (the TDWR scan elevation angle is used to derive altitude) and outputs this information immediately to 3-DAT so it can begin computing a new TDWR scan strategy with more precise elevation information. ARDA then searches the data regions in the vicinity of the aircraft for aerosol release signatures and outputs any plume detections to the NBCMsgServer output server.

NBCMsgServer converts the ARDA plume detection messages to XML-formatted NBC-1 messages that are posted to a configured list of web URLs associated with one or more emergency operations centers (EOCs).

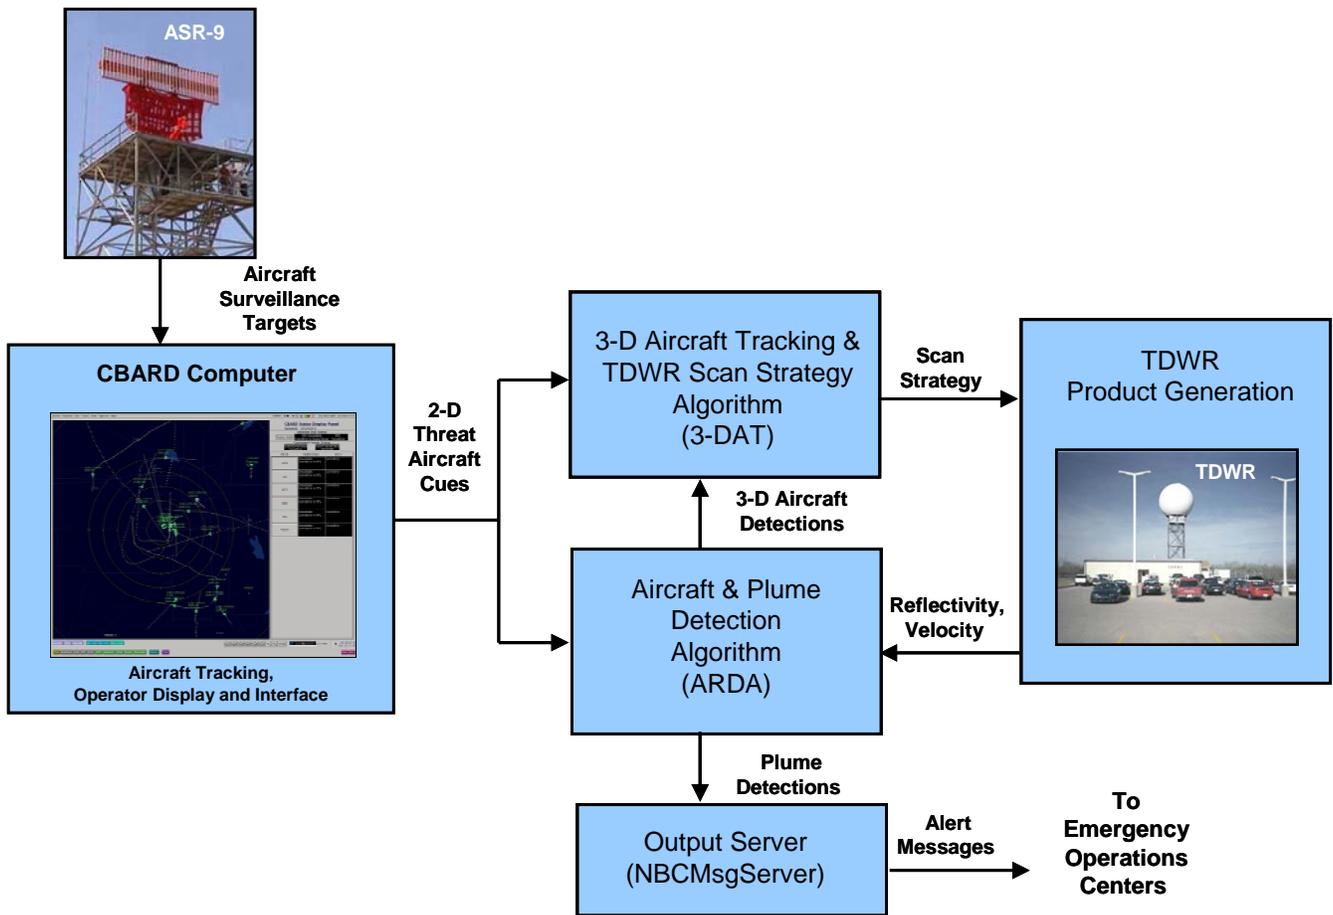


Figure 5. CBARD functional architecture.

4. AIRCRAFT AND PLUME DETECTION ALGORITHM DESCRIPTION

4.1 ARDA DATA FLOW

Figure 6 illustrates the data flow through the Aerosol Release Detection Algorithm. After a CBARD operator has selected a threat aircraft and enabled the CBARD system, 2-D (range, azimuth) cue locations of the threat aircraft and TDWR base data begin flowing to ARDA. The TDWR data are passed through a translator task (tdwrResampler) that converts the TDWR reflectivity and Doppler velocity data messages to a Lincoln Laboratory data format (WxCAR) that can be read by ARDA. ASR-9 based aircraft location messages generated by the CBARD computer are multiplexed with the translated TDWR data messages by the ARDAServer task, which provides a single stream of merged input data to ARDA. This merged data stream provides a convenient input data archival point that ensures the original ordering of the input data will be preserved for future playbacks.

There are two major detection stages in ARDA: aircraft detection and plume detection. The TDWR sector scans are centered over the ASR-9 aircraft cue location and at a series of discrete elevation angles that vary depending on the TDWR scan strategy generated by the CBARD 3-DAT module. ARDA uses the most recent ASR-9 cue location to localize its search for the aircraft target in the TDWR data. If the aircraft is detected on a given TDWR elevation scan, ARDA estimates the aircraft altitude from the TDWR scan geometry and immediately outputs the 3-D aircraft location information to the 3-DAT algorithm where it is used to generate an updated TDWR scan strategy.

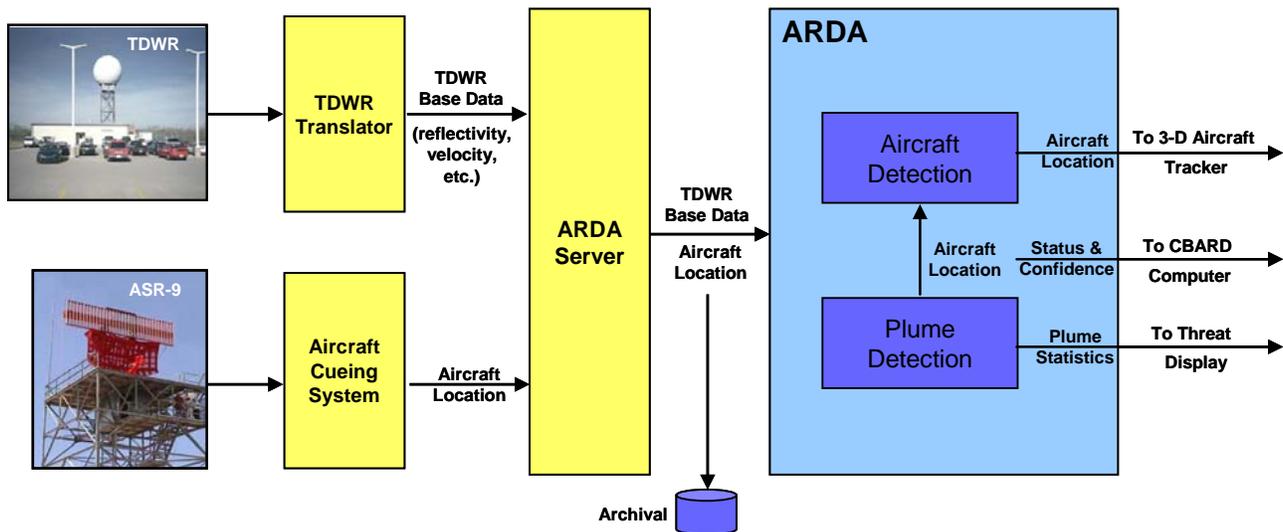


Figure 6. ARDA data flow.

If the aircraft has been detected in the TDWR data, the aircraft location is stored internally and ARDA proceeds with processing of the current radar data, looking for evidence of a plume signature within a region behind the detected aircraft location. If a plume is detected, a report message is generated and output on an Ethernet socket connection, where it is subsequently reformatted and forwarded to one or more EOCs for contamination hazard analysis and display.

Following the end of each processing cycle (approximately once every 9 seconds during CBARD processing or once every 30 seconds if CBARD processing is not engaged), a status message is output to the CBARD computer. The message contains status bits indicating the health of the software, and also contains a 3-level plume detection confidence measure (low, medium, high) based on a real-time assessment of the atmospheric background reflectivity statistics and their effects on detection performance.

4.2 RADAR SIGNATURES

This section describes the radar signatures that ARDA recognizes for aircraft and plume detection.

4.2.1 Aircraft Signatures

Two radar signatures have been identified as indicators of the aircraft in the TDWR reflectivity and velocity data:

1. A small region of substantially enhanced reflectivity (reflectivity peak) co-located with Doppler velocity values that differ substantially from neighboring Doppler velocities (differential velocity peak).
2. Motion in the reflectivity field as revealed by scan-to-scan differencing of the reflectivity images.

4.2.1.1 Reflectivity Peak with Differential Velocity Peak

Figure 7 provides example close-ups of TDWR reflectivity (left) and corresponding velocity data (right) for two different scans in which the aircraft was intercepted. In these images, the open white circle is centered on the expected location of the aircraft as given by the ASR-9 range-azimuth cue location provided by the CBARD tracker. In both of these scans, the aircraft can be clearly seen in the reflectivity image as a 3-4 pixel region of orange or red pixels (> 25 dBZ).

In the associated Doppler velocity images, small regions of corresponding pixels having radial velocities that differ from those of the surrounding environment can be seen. In the left-most example, the relatively strong velocities directed away from the radar associated with the aircraft are indicated by the pink and red pixels. These contrast with the ambient radial velocities which are generally 2 - 4 m/s

directed away from the radar. In the right-most example, which is only 9 seconds earlier, the aircraft can be seen in the velocity image as a small region of green pixels corresponding to a velocity of approximately -2 m/s directed toward the radar. Because the unambiguous velocity range for the TDWR is approximately ± 26 m/s, and because the crop duster aircraft used for tests was flying at nearly 50 m/s, velocities associated with the aircraft are folded and may appear practically anywhere in the Doppler spectrum depending on the heading of the aircraft with respect to the radar viewing angle. It may even fold in such a way as to blend with the ambient velocity returns. Therefore, the absence of an associated differential velocity signature does not necessarily discount a reflectivity peak signature. However, the appearance of the differential velocity signature increases the probability that the reflectivity peak signature is truly associated with the aircraft.

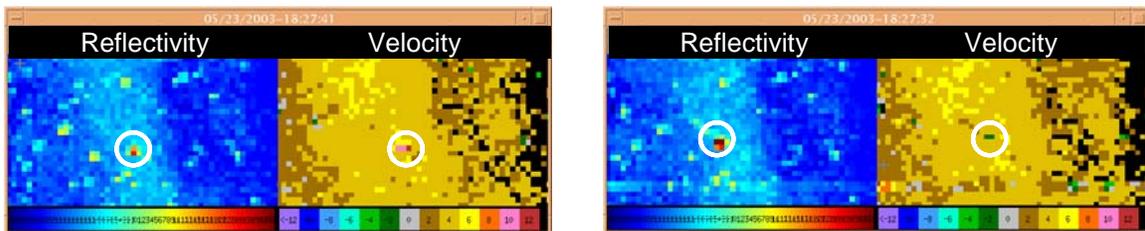


Figure 7. Reflectivity peak and differential Doppler velocity associated with aircraft target.

4.2.1.2 Reflectivity Motion

Provided the TDWR is able to scan the aircraft frequently enough, the cohesive motion of the aircraft can be recognized through reflectivity image subtraction. By pixel-wise subtracting a prior reflectivity image from the current one, moving features such as the aircraft will show up as a small region of positive pixel values followed by a region of negative pixel values in the resulting reflectivity difference image. By cumulatively storing the positive differences and applying a recursive temporal decay factor to the cumulative differences, an image containing a streak with a fading tail along the recent path of the aircraft is generated for which a suitable image pattern recognition template can be applied. Figure 8 shows an example cumulative reflectivity difference image. The path of the aircraft is readily recognized in this example. A detailed description of the preparation and computation of the cumulative motion reflectivity difference image is presented in Section 4.4.3.

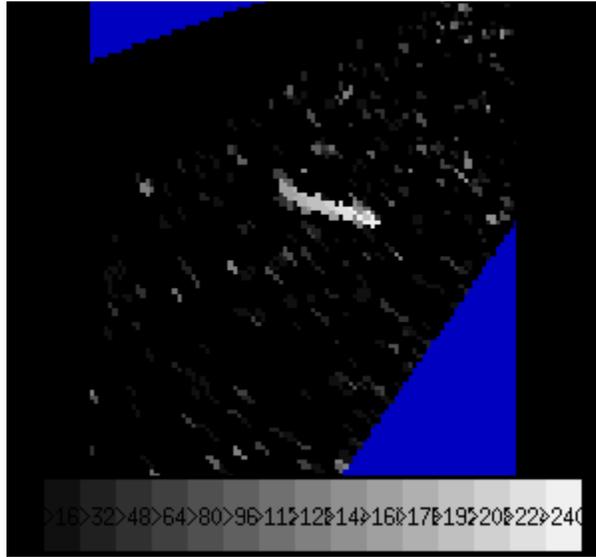


Figure 8. Cumulative reflectivity difference motion signature associated with an aircraft.

4.2.2 Plume Signatures

4.2.2.1 Reflectivity Thin Line

Figure 9 shows TDWR reflectivity images from two different scans that intercepted the plume of water being released from a crop duster during the CR8 field tests. The red circle is centered on the ASR-9 cue location of the aircraft provided by the CBARD tracker. In both of the examples, the reflectivity peak associated with the aircraft can be seen near the center of the red circle. Extending downward from the reflectivity peak is a thin line of enhanced reflectivity associated with the aerosol plume (in this case, water). The reflectivity of the plume is quite low – in these examples it ranges from about -5 to -2 dBZ. If background reflectivity levels are too high, e.g., from insect returns or clutter sidelobes, then it can be difficult to distinguish the plume from the cluttered background (note the considerable background clutter in the left-most example). Spatial proximity of the thin line to the reflectivity peak from the aircraft and the orientation of the thin line with respect to the aircraft heading are important contextual links for interpreting the echo region as belonging to a plume.

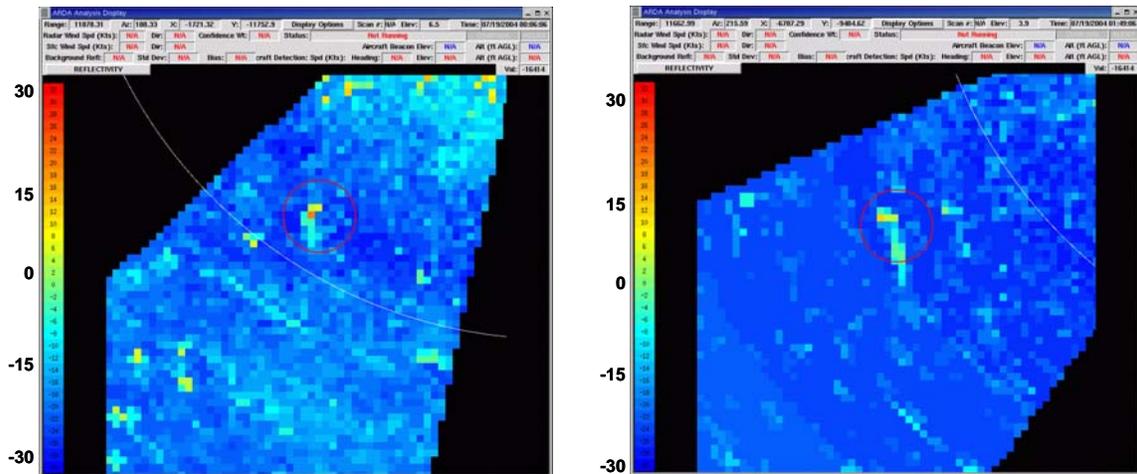


Figure 9. Crop duster release plume signatures in TDWR reflectivity data. Data are from CR8.

4.3 ARDA PROCESSING OVERVIEW

4.3.1 Feature Detection Approach

In ARDA, feature detection is performed utilizing a pattern matching image processing approach originally developed by Lincoln Laboratory for automatic target recognition (ATR) called Functional Template Correlation (FTC) (Delaney et al., 1992). FTC utilizes a pattern template together with associated scoring functions to generate a map of evidence from an input image called an interest image, which indicates the degree to which the feature being sought is present in the imagery (e.g., an aircraft reflectivity peak or a reflectivity thin line echo associated with an aerosol plume). Evidence probabilities are expressed in the interest images as gray scale pixel values ranging from 0 to 255, with values less than 128 indicating disconfirming evidence and values greater than 128 indicating confirming evidence. Interest images provide a mechanism for data fusion and consensus building. Consensus is achieved through interest image combination using a weighted average of spatially co-registered interest image pixels.

As an illustration, Figure 10 shows the functional template used to detect the reflectivity thin line echo signature associated with a plume. The functional template in this example is a 7x7 grid (kernel) of scoring function indices. The kernel resolution is presumed to be spatially matched to that of the input image to be probed (in this case 150 meters). The green stripe region in the center of the template is assigned to the scoring function that returns maximal scores for image reflectivity values that fall within the range expected for the thin line echo region, while the red template pixels are assigned to the scoring

function that returns optimal scores for low reflectivity background values in the flanking regions on both sides of the thin line echo. The white pixels in the template illustration represent guard regions in which the input image (i.e., reflectivity) values are ignored and have no effect on match scores. The template is centered over each reflectivity image pixel and, at each image pixel location, is rotated at 10 degree increments through 180 degrees. At each rotation of the template kernel, the individual pixel scores are averaged. The maximum of the average kernel scores over all rotations of the kernel is clipped to the interval [0, 255] and assigned to the corresponding output interest image pixel location.

Through input image masking, the FTC operation can be restricted to appropriate areas in the input imagery. For example, the functional template shown in Figure 10 has been allowed to operate only within an oval-shaped region of the input image corresponding to an area of prior plume detections (the solid blue area in the output interest image at right is the masked region). Such contextual processing helps reduce false alarms and allows maximum detection sensitivity to be encoded in the FTC scoring functions.

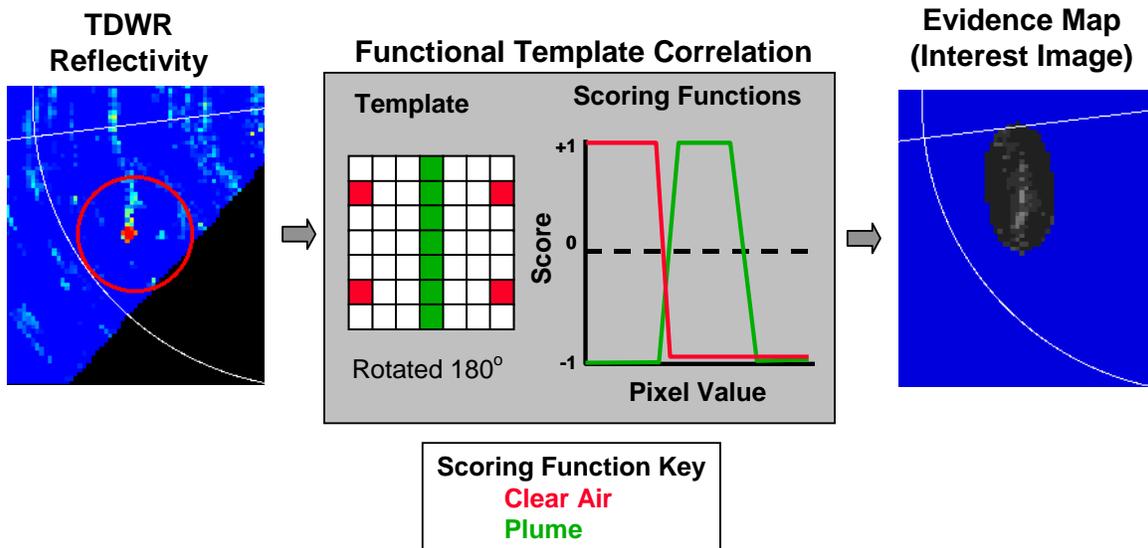


Figure 10. Example functional template for plume reflectivity thin line feature detection.

4.3.2 ARDA Processing Flow

Figure 11 is a block diagram illustrating the ARDA processing flow. Aircraft cue location information are provided by the CBARD ASR-9 tracker and are received by ARDA at the nominal ASR-9 data rate of approximately once every 5 seconds. ARDA also receives a constant elevation azimuthal sector scan of TDWR basedata products (reflectivity, Doppler velocity, spectrum width, and signal-to-noise) roughly every 9 seconds when the TDWR is operating in CBARD mode. Although the elevation angle is constant for a given sector scan, it typically changes with each scan as the TDWR follows the aircraft. Scan elevation angles may change appreciably from scan to scan if the aircraft is flying directly toward or away from the radar, even if the aircraft is flying level.

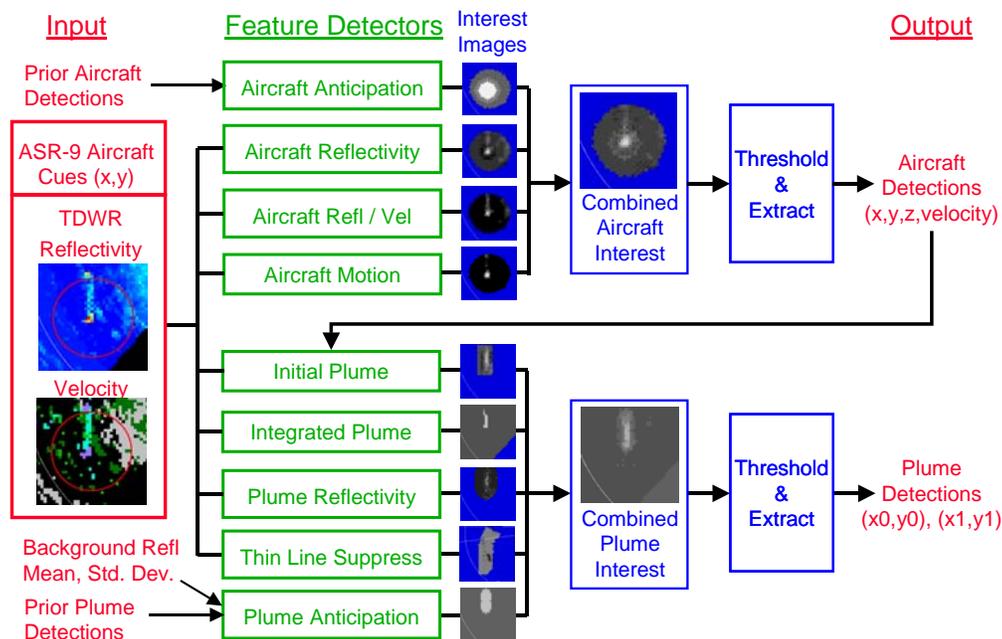


Figure 11. ARDA functional block diagram.

The TDWR and ASR-9 aircraft cue location data are supplied to a series of feature detectors, each of which generates an interest image indicative of the degree to which the feature being sought is present in the data. The aircraft feature detectors are run first, and their resulting interest images are combined using a weighted average to form a combined interest image representing a consensus of evidence. A

threshold is applied to extract one or more interest regions having sufficiently strong evidence to warrant a possible detection. The extracted interest regions are analyzed and subjected to various attribute tests, and if the candidate region is sufficiently “interesting” the aircraft detection is output and the information is internally forwarded to the plume detection stage. Before output, the current aircraft detection is associated with prior aircraft detections to form a track vector for the aircraft. The track vector together with the prior aircraft detection location are used in the processing of the next TDWR scan to form an aircraft anticipation interest image that selectively sensitizes the detection algorithm to locations where the aircraft is expected to be.

Following the aircraft detection stage, the plume detectors process the radar data together with the detected aircraft location to generate interest images indicative of possible plume features. The interest images are averaged to form a combined plume interest image upon which a threshold is applied to extract the sufficiently interesting plume regions. Detected plumes are output and are stored internally where they are used in formulating regions of localized anticipation during the next processing iteration.

4.4 DATA PREPARATION

Prior to feature detection, the received TDWR reflectivity and velocity data are first converted to 2-D Cartesian spatial images (spectrum width and signal-to-noise products are not currently utilized). The Cartesian reflectivity images are then pre-processed to normalize the reflectivity levels with respect to the environmental background.

4.4.1 Cartesian Resampling

To facilitate image processing and pattern recognition of straight-line features (such as aerosol plumes), the input polar coordinate data (range, azimuth) are mapped onto a Cartesian image grid. The choice of Cartesian grid resolution is important. A coarser grid allows image processing operations to proceed more quickly, but the grid cannot be too coarse or details of the radar feature being sought will not be adequately resolved. Different Cartesian grid resolutions were investigated during development. Figure 12 shows an input polar reflectivity image (as a B-scan display with range increasing from left to right and azimuth increasing from bottom to top) and the results of mapping the data onto Cartesian grids with resolutions of 250 meters and 150 meters. A resolution of 150 meters was chosen as it better resolves the aircraft and plume signatures, and algorithmic processing is still fast enough to keep up with the input radar data rate.

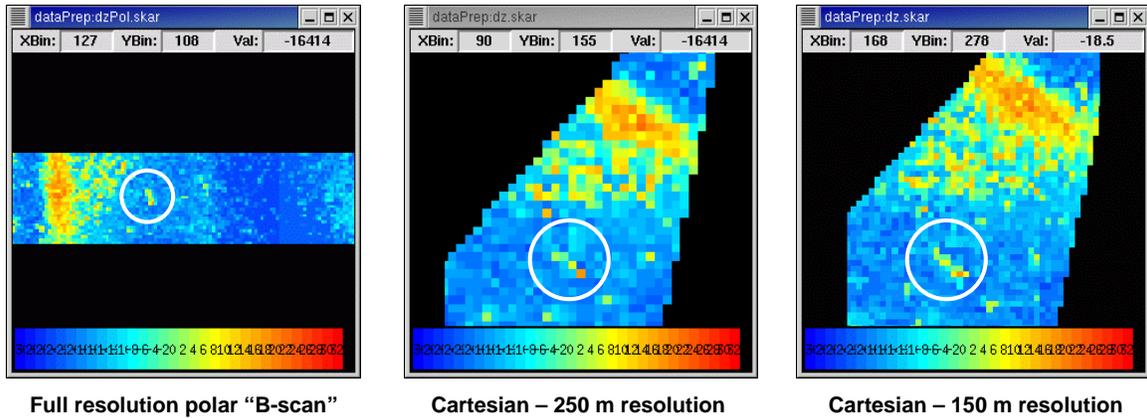


Figure 12. Effects of Cartesian resampling grid resolution. Location of aircraft and plume feature is indicated by the white circle.

4.4.2 Reflectivity Background Normalization

Even in optically “clear” environmental conditions, background reflectivity levels can vary diurnally or seasonally owing to the presence of environmental aerosols, especially biological targets (e.g., insects, birds). The contrast between an aerosol plume and the background is often considerably reduced when the background reflectivity is high. During FTC pattern-matching in ARDA, neighboring background reflectivity pixels are probed by the scoring functions which return optimal scores over a narrow range of expected clear air values. Extending the range of optimal background image values to encompass the broad range of possible background levels is undesirable as it increases the likelihood of false alarms. While one could dynamically adjust the scoring functions based on an overall assessment of the background, there are often smaller-scale regional fluctuations in the background that would not be optimally matched as the aircraft flies through them.

The approach used by ARDA is to apply a local averaging window at each reflectivity pixel to normalize the reflectivity levels prior to image processing. The normalization is performed on the Cartesian reflectivity image using a 20x20 (3 km x 3 km) averaging window. The mean of the reflectivity values in the window is subtracted from the central pixel value in the window to produce the background-normalized estimate. The window is sufficiently large that the relatively small number of pixels in the aircraft and plume regions do not substantially bias the mean that is subtracted. Figure 13 shows the results of the background normalization. All of the subsequent reflectivity-based feature detection is performed using the normalized image.

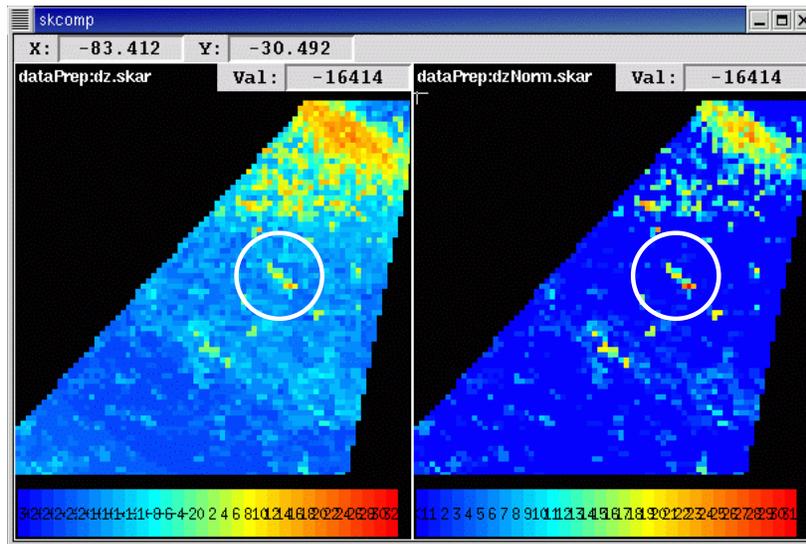


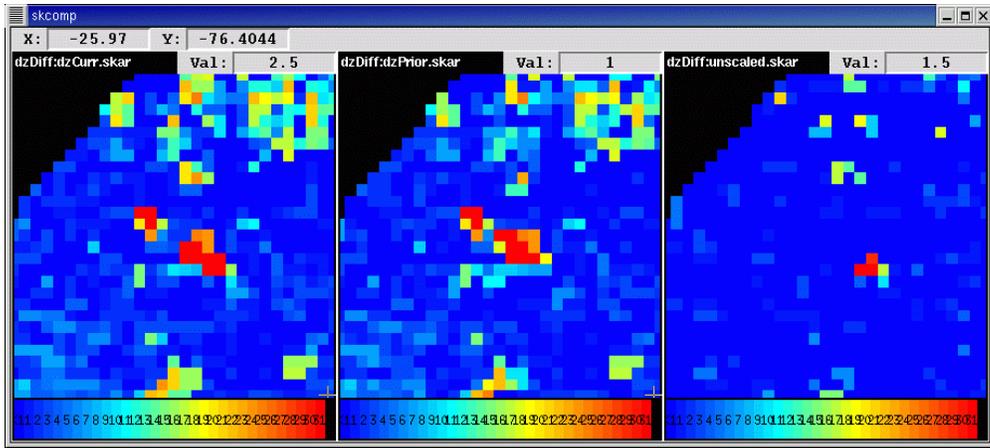
Figure 13. Example of background reflectivity normalization. (Left) Original reflectivity image. Aircraft and plume signature is indicated with white circle. (Right) Results of background normalization using a 20x20 averaging window.

4.4.3 Cumulative Motion Difference Image

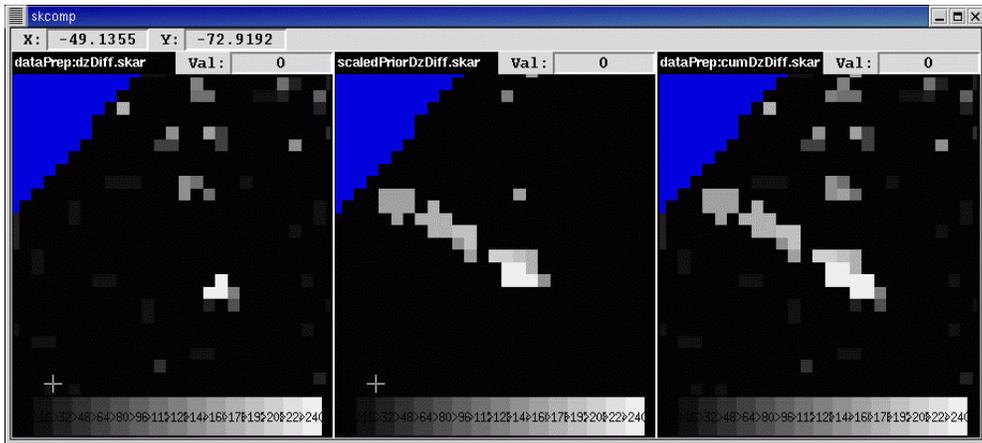
Motion detection is an important aspect of aircraft feature detection that allows ARDA to discriminate the small reflectivity peak associated with an aircraft from other peaks in the image that may be associated with stationary ground clutter residue or relatively slow moving atmospheric clutter. In ARDA, this is accomplished through construction of a cumulative reflectivity motion difference image for which a moving target is seen as a line of positive reflectivity difference pixels. FTC is used to recognize the line feature in the cumulative reflectivity difference image. The computation of the cumulative motion difference image is accomplished in several steps:

First, a composite maximum reflectivity image associated with the prior radar scan (approximately 9 seconds ago) is retrieved from the ARDA image history, and is subtracted pixel-wise from the current composite maximum reflectivity image (see Figure 14: “Current Reflectivity” and “Prior Reflectivity”). These composite maximum reflectivity images represent the maximum reflectivity at each pixel location taken over the preceding 90 seconds (a parameter corresponding to approximately ten TDWR sector scans). The composite maximum reflectivity image is used for motion differencing because the bracketing scan pattern utilized for the TDWR during CBARD mode doesn’t always intercept the aircraft (or only partially intercepts the aircraft) on each scan. If antenna tracking is going well, the aircraft will likely be intercepted at least a few times during the compositing interval, so the aircraft echo will be repeated or smeared along the aircraft path. When the two images are subtracted, the positive difference associated with the most recent interception of the aircraft is revealed as seen in Figure 14: “Difference”.

Next, the reflectivity difference image is re-scaled such that differences of 0-30 dB linearly map to the interval 0-255 (Figure 14: “Scaled Difference”). Values greater than 30 dB are clipped to a value of 255. Negative differences are clipped to zero. Finally, the cumulative motion difference image is updated by taking the pixel-wise maximum of the current scaled difference image and 0.95 times the prior cumulative motion difference image value (Figure 14: “Prior Cumulative Difference”). The recursive decay factor of 0.95 creates a fading tail of differences in the cumulative reflectivity difference image (Figure 14: “Cumulative Difference”).



Current Reflectivity Prior Reflectivity Difference



Scaled Difference Prior Cumulative Cumulative Difference
[0,30] -> [0,255] Difference x 0.95

Figure 14. Summary of image pre-processing stages for determination of aircraft motion by reflectivity differencing. See text for description.

4.5 AIRCRAFT DETECTION

4.5.1 Aircraft Feature Detectors

Following receipt of TDWR sector scan and image preparation, aircraft detection is performed. The prepared reflectivity and Doppler velocity images are processed by a series of aircraft feature detectors to generate interest images. To reduce the likelihood of false aircraft detections, only a 2 km radius circle centered on the ASR-9 cue location is processed by the aircraft feature detectors. No opinions are expressed outside of the 2 km radius. Table 1 provides a brief description of each of the aircraft feature detectors. Detailed descriptions of each of the aircraft feature detectors follow.

TABLE 1
Aircraft Feature Detectors

Feature Detector	Description
Aircraft Reflectivity Peak	Identifies small reflectivity peak associated with aircraft.
Aircraft Reflectivity/Velocity	Tandem detector that looks for reflectivity peak co-located with Doppler velocity that is different from surrounding background
Aircraft Reflectivity Difference	Detects instantaneous aircraft motion through differencing of the current and prior reflectivity images.
Aircraft Motion	Identifies streak in cumulative reflectivity difference image.
Aircraft Anticipation	Localized anticipation interest based on extrapolation of prior aircraft detection.

4.5.1.1 Aircraft Reflectivity Peak Detector

The Aircraft Reflectivity Peak detector returns maximal confirming interest scores for small peaks of high reflectivity surrounded by clear air returns. Prior to FTC, a copy of the background normalized reflectivity image is linearly rescaled such that reflectivity values of 0-40 dBZ map to the interval [0,255]. The scaled reflectivity image is then input to FTC with a processing mask that excludes FTC output for image regions outside of a 2 km radius centered on the aircraft cue position.

The scoring functions for this feature detector were inherited unchanged from an earlier developmental version of ARDA that allowed negative interest scores and used zero as the confirming/disconfirming interest ambiguity point instead of the value of 128 used in the current version of ARDA. In order to properly combine the interest values from this detector with other “128-based” interest images, the initial interest values that are output from FTC are first scaled by 0.5 and then linearly rescaled such that values in the interval [-255, 255] map to the interval [0, 255].

Figure 15 shows the template and scoring functions along with an example of interest image output.

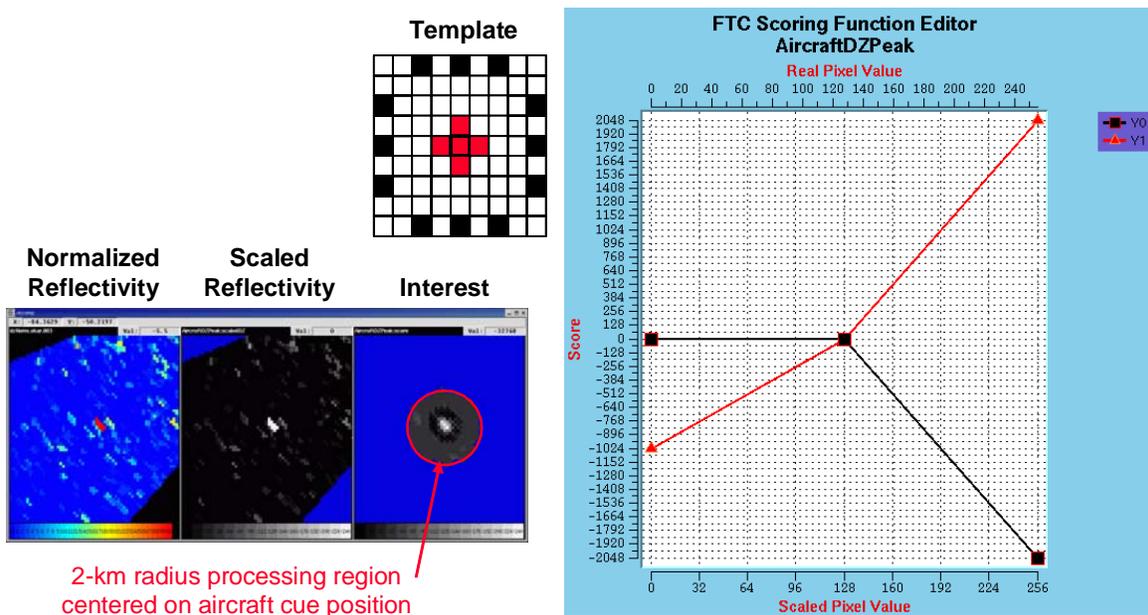


Figure 15. Illustration of aircraft reflectivity peak feature detector.

4.5.1.2 Aircraft Tandem Reflectivity/Velocity Detector

The Aircraft Reflectivity/Velocity detector is an example of a tandem FTC filtering operation. It returns maximal confirming interest scores for small reflectivity peaks that are co-located with velocity values that differ from the surrounding background velocity values (i.e., “differential velocity”). For this detector, the original resolution polar data are used. Prior to FTC, the differential velocity image is computed by subtracting a 3x3 median-filtered copy of the polar velocity image from the original velocity image and taking the absolute value of the differences (Figure 16).

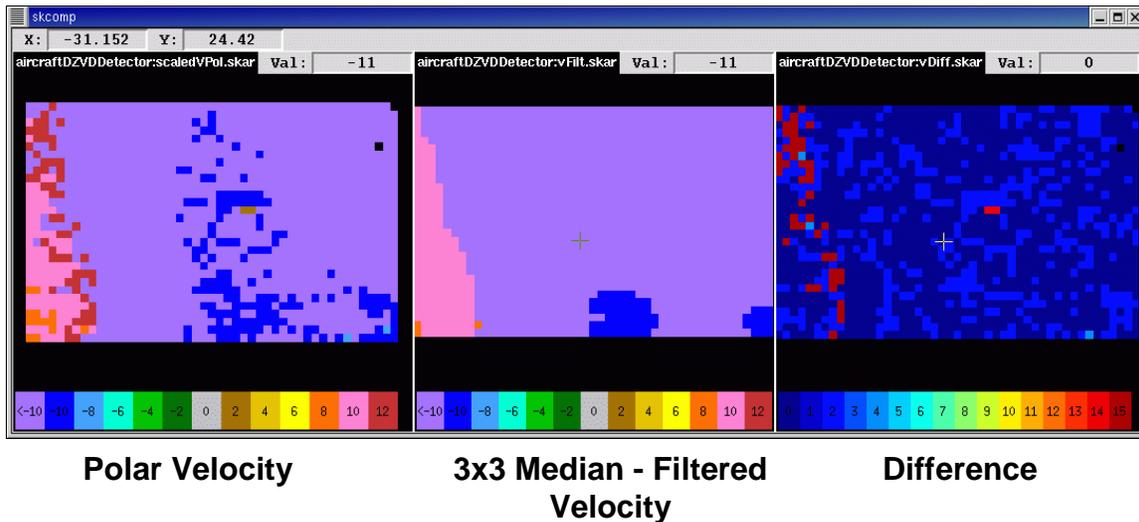


Figure 16. Computation of differential velocity image.

In tandem FTC operation, two separate template kernels and associated scoring functions are applied to the two input images. The templates are passed over the two images in lock-step and the scores at each template location and orientation are simultaneously averaged. Thus, the highest match scores are returned when good pattern matches are obtained in both input images. For this detector, the polar reflectivity and the pre-computed differential velocity images are processed with the FTC templates and scoring functions shown in Figure 17. Again, the FTC operation is restricted from operating outside of a 2 km radius centered on the aircraft cue position. The resulting polar-coordinate interest image is then resampled to the common Cartesian grid used by the other feature detectors so that the evidence can spatially register and fused later.

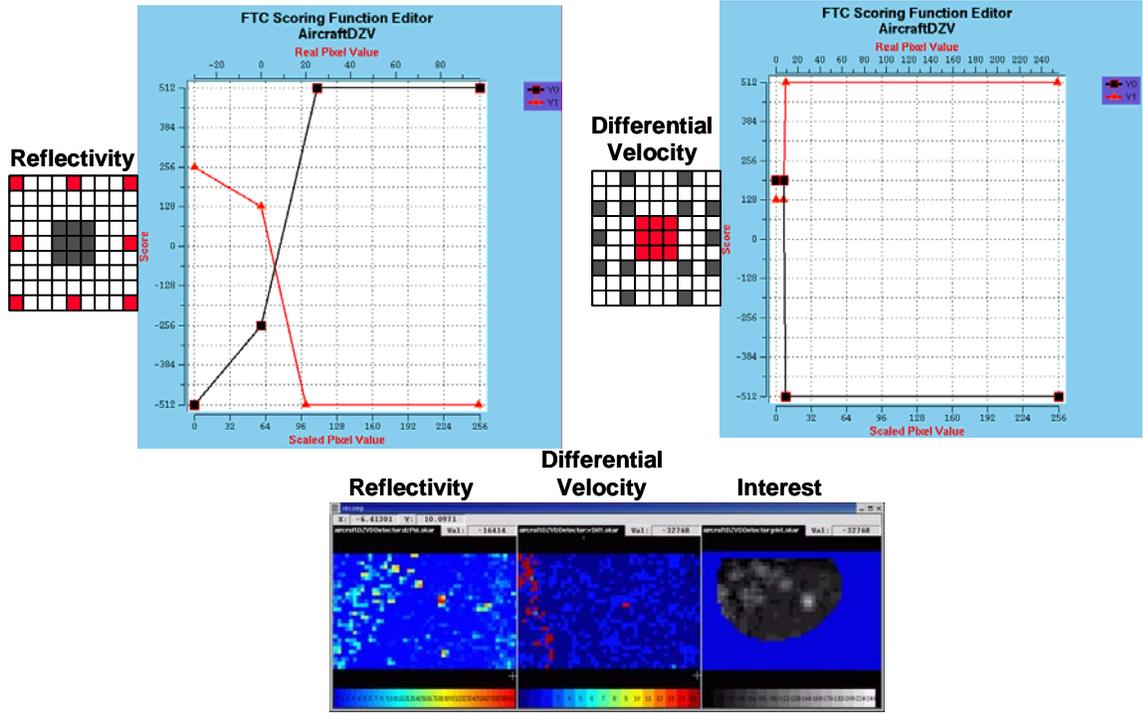


Figure 17. Illustration of Aircraft Tandem Reflectivity detector.

4.5.1.3 Aircraft Reflectivity Difference Detector

The Aircraft Reflectivity Difference detector responds to instantaneous aircraft motion by identifying the positive reflectivity difference peak in a reflectivity difference image. The difference image is created by subtracting the prior normalized reflectivity image from the current one and linearly rescaling the differences such that values of -20 to 30 dB map to the interval [0, 255]. Like the Aircraft Reflectivity Peak detector described earlier, the scoring functions for this feature detector were inherited unchanged from an earlier developmental version of ARDA that allowed negative interest scores and used zero as the confirming/disconfirming interest ambiguity point instead of the value of 128 used in the current version of ARDA. In order to properly combine the interest values from this detector with other “128-based” interest images, the initial interest values that are output from FTC are first scaled by 0.5 and then linearly rescaled such that values in the interval [-255, 255] map to the interval [0, 255]. Figure 18 illustrates the Aircraft Reflectivity Difference feature detector processing.

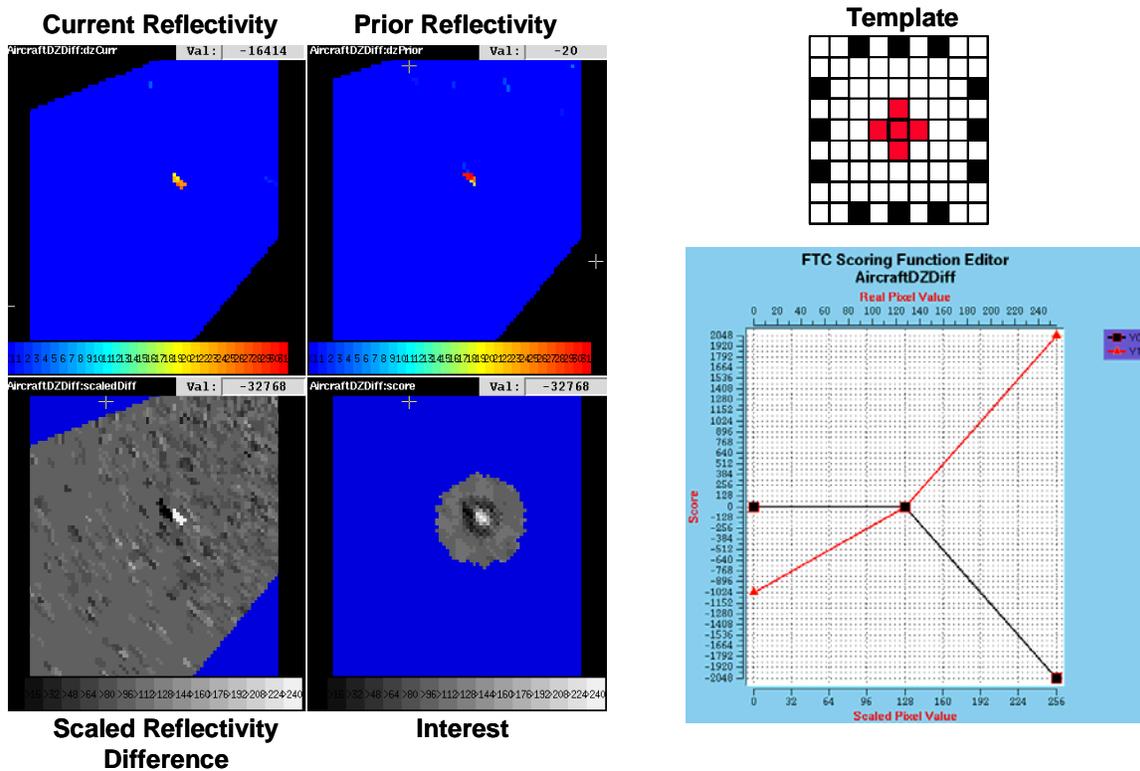


Figure 18. Illustration of Aircraft Reflectivity Difference feature detector.

4.5.1.4 Aircraft Motion Detector

The Aircraft Motion detector is another tandem feature detector. It simultaneously probes the current reflectivity difference image (the current reflectivity image minus the prior reflectivity image), and the cumulative reflectivity motion difference image that was pre-computed during data preparation (see Section 4.4.3). The cumulative motion template has a band of scoring functions that respond to the trail of cumulative reflectivity differences that fall behind the current aircraft location. The current motion template looks for the differential reflectivity peak associated with the leading edge of the aircraft target motion. Figure 19 shows the FTC templates and scoring functions along with example results.

Since the aircraft is not necessarily intercepted on every scan, the current motion difference image may not reveal the current location of the aircraft. Performing tandem detection with the cumulative motion image helps ensure that at least a modest confirming interest score is returned if there is a trail of motion-induced cumulative reflectivity differences. The score is higher, of course, if the aircraft is present in the current motion image.

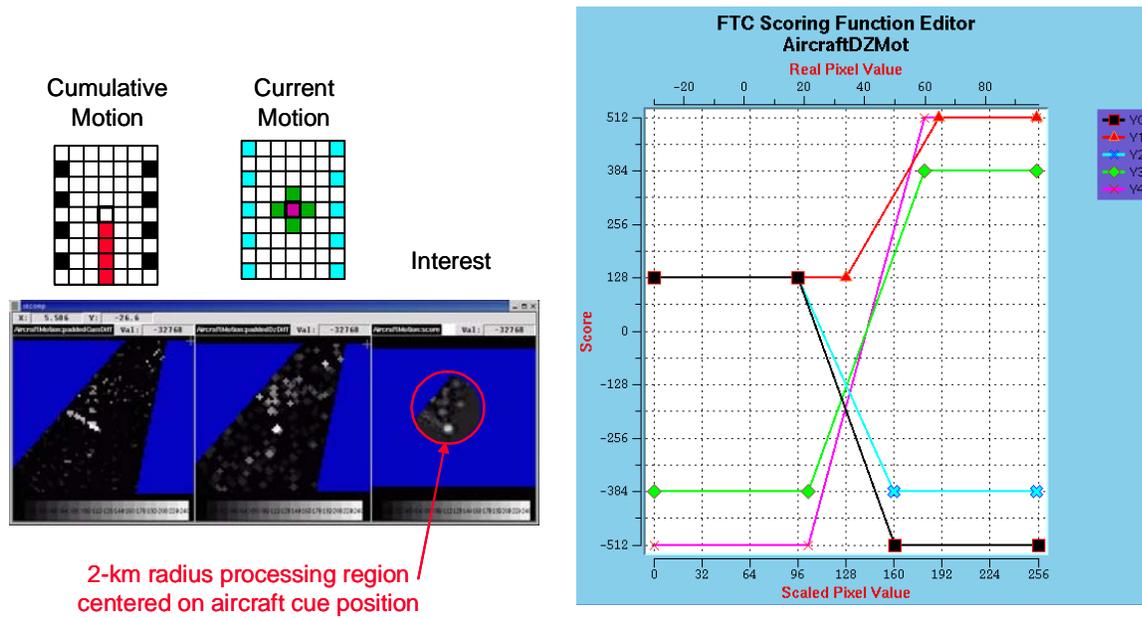


Figure 19. Illustration of Aircraft Motion feature detector.

4.5.1.5 Aircraft Anticipation

The Anticipation interest image provides a mechanism to spatially adjust the detection sensitivity of ARDA, based on situational or environmental context. When high localized anticipation interest values (based on extrapolation of prior detections) are averaged with interest values from the feature detectors, weak evidence that might otherwise be too low to produce a detection is boosted, allowing detection to be maintained. Conversely, low anticipation values can suppress the likelihood of detection.

By default, the Aircraft Anticipation image has a background global interest value of 64 everywhere (this represents a suppressive value within the 0-255 interest scale having 128 as the ambiguity point between confirming and disconfirming evidence). A site-specific parameter file (targZones.param) contains definitions of sensitive target areas where ARDA is to maintain higher detection sensitivity. Each entry in the parameter file describes a central location (latitude/longitude) and a radius of influence in meters. Within each circular target region, the background interest value is boosted to 96 (still suppressive, but not so much so). Figure 20 shows an example of three target zones represented in an Aircraft Anticipation image (two of them overlap). For the OKC tests, a single target zone centered at the radar and extending to the full 50-km ARDA processing range was used, so effectively the global background anticipation for aircraft detection was 96.

In addition to the regionalized boosting from the target zones, a more localized anticipation boosting is accomplished through extrapolation of the currently tracked aircraft position. At the extrapolated location, the anticipation level is boosted to a value of 128 within a 1350 meter radius and to 192 within a 900 meter radius centered on the aircraft cue location. This provides important localized detection support in locations where the aircraft is expected to be (see Figure 20). Again, the anticipation interest is limited to the 2-km region surrounding the current aircraft cue location, and provides no opinions outside of that radius.

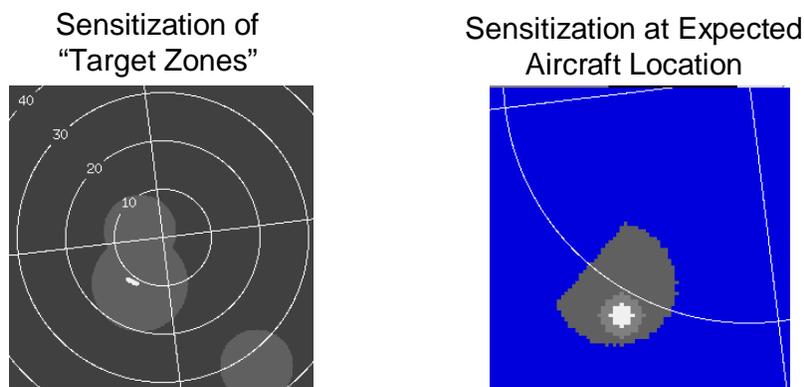


Figure 20. Aircraft Anticipation interest images.

4.5.2 Combination of Aircraft Feature Detector Evidence

The individual aircraft feature detector interest images are combined using a pixel-weighted average to form a pre-combined feature detector interest image. Table 2 lists the confirming and disconfirming weights that are applied to each aircraft feature interest pixel during the averaging. The weights are tuning parameters that are readily adjustable. Interest values less than 128 are assigned the disconfirming weight, while interest values greater than or equal to 128 are assigned the confirming weight. Since the various feature detectors may have different degrees of aptitude for correctly expressing confirming versus disconfirming opinions, the differential weighting scheme allows this skill level to be accounted for when forming the consensus.

TABLE 2
Feature Detector Weights Used for Aircraft Interest Fusion

Feature Detector	Confirming Weight	Disconfirming Weight
Aircraft Reflectivity Peak	1.00	1.00
Aircraft Reflectivity Difference	1.00	0.50
Aircraft Reflectivity/Velocity	1.00	0.50
Aircraft Motion	1.00	0.50
Aircraft Anticipation	0.33	0.33

The combined feature detector evidence is assigned confirming and disconfirming weights of 1.0 and then averaged against the aircraft anticipation interest (with its weights as indicated in Table 2) to form the total combined aircraft interest image (Figure 21).

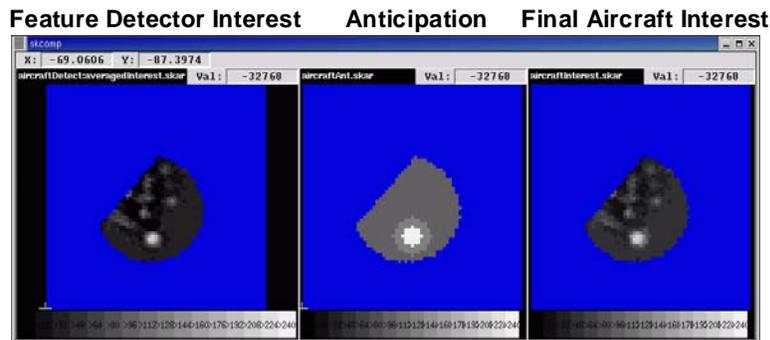


Figure 21. Aircraft interest image fusion.

4.5.3 Extraction of Aircraft Features

After combining the aircraft feature detector interest with the aircraft anticipation interest, an interest threshold of 127 is applied to extract the candidate regions of evidence that may be associated with the aircraft target (Figure 22). Region analysis is then conducted on each connected region of above-threshold interest to compute statistics, and the regions are subjected to an area threshold parameter currently set to 0.05 sq km.

For each candidate region that exceeds the area threshold, the location of the maximum reflectivity within the region is identified and is assigned as the location of the aircraft target. A measure of the total aircraft target power is computed by summing the nine highest reflectivity values lying within the extracted aircraft target region.

An initial estimate of the target altitude is then computed from the target range and the elevation angle of the current TDWR scan. Since the aircraft's true location may fall in between the elevation scans of the TDWR, the initial altitude estimate will be adjusted later based on comparisons with prior aircraft target reflectivities on adjacent elevation scans.

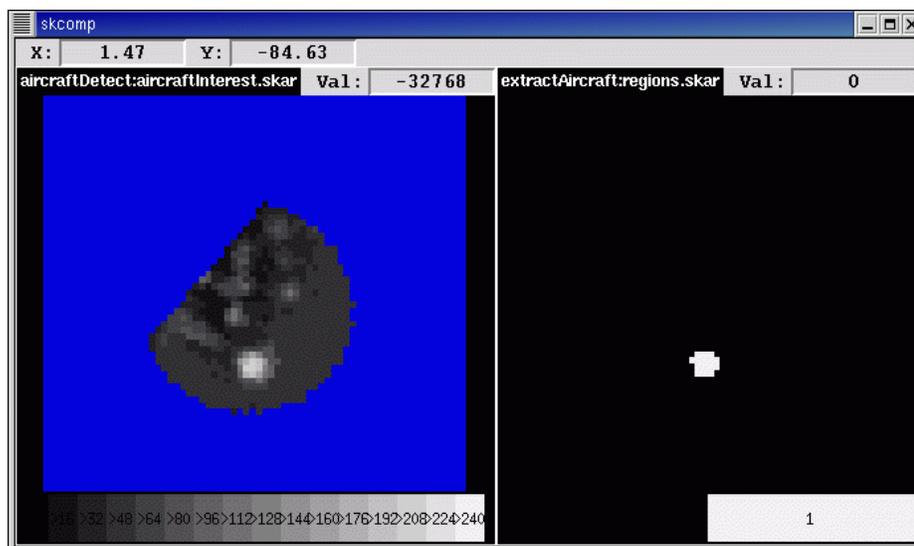


Figure 22. Results of feature extraction. (Left) Combined interest image. (Right) Extracted region after applying interest threshold.

4.5.3.1 Correspondence and Tracking

Next, an attempt is made to establish correspondence between each of the candidate aircraft targets and the prior aircraft detection. Correspondence is established if the candidate target is within a parameter distance threshold (currently 1500 meters) of a prior aircraft detection. If more than one prior detection is within the distance threshold, then the closest one is chosen for the correspondence association. Once correspondence is established, the unique ID number of the prior detection is linked to the current candidate, and the aircraft velocity (speed and heading) is computed based on the target location displacement over time between the prior and current scan. The aircraft velocities are temporally smoothed by performing vector addition of the current and prior velocity vectors and multiplying the vector result by 0.5.

If no candidate regions are found on the current scan, then a check is made to see if a prior aircraft detection should be coasted. This is accomplished by checking the aircraft detections in the prior event (there should be only one) and seeing if the detection ID is linked to a current detection candidate. If the prior aircraft detection is not linked to a current one, and if it has not been coasted too many times (a parameter currently set to 3 scans), then the prior detection is coasted by synthesizing a new aircraft detection object based on attributes copied from the prior detection, except that the synthesized detection has an updated location based on extrapolation of the prior detection's location and trajectory.

Finally, the list of candidate aircraft targets is pruned to retain the single detection having a location closest to the current aircraft cue location provided by the CBARD ASR-9 tracker.

4.5.3.2 Aircraft Altitude Adjustment

Accurate estimation of the aircraft altitude is important because the 3-DAT algorithm uses it to dynamically adjust the TDWR scan pattern so that the aircraft and possible trailing aerosol plume are optimally scanned by the radar. An initial aircraft altitude estimate is computed using the range to the target and the current TDWR elevation angle. However, it is possible that the true aircraft altitude falls between the discrete elevation scans of the TDWR.

To account for this, the initial altitude estimate is refined by comparing the elevation angle and total reflectivity power of the current aircraft detection against the elevation angle and power of the prior aircraft detection. If the elevation angles are less than a TDWR beam width (0.5 degrees) apart, then the two aircraft signal powers are fit to a Gaussian-shaped TDWR antenna pattern model. The location of the fitted Gaussian peak is presumed to be the elevation angle that would have returned the maximum aircraft power, and from this elevation angle, a revised altitude is derived. If the prior aircraft detection elevation angle is more than a beam width away from the current one, then the altitude of the detection having the maximum aircraft power is assigned to the current detection.

4.5.4 Aircraft Detection Output

Once the final estimation of the aircraft altitude has been completed, the 3-D aircraft location is immediately output so that the 3-DAT scan strategy generation module can compute the next sequence of TDWR scans required to maintain track of the aircraft. Having now identified the location of the threat aircraft in the TDWR data, ARDA can begin the process of looking for aerosol plumes in the regions trailing the aircraft.

4.6 PLUME DETECTION

4.6.1 Plume Feature Detectors

Plume detection follows immediately after the aircraft detection stage. This processing sequence is necessary because several of the plume detectors require knowledge of the location and movement of the aircraft so that processing can be focused on the area in the vicinity of the aircraft. Plume detection is accomplished using the same general processing sequence used for aircraft detection: feature detection using FTC, interest combination, thresholding, and extraction. In general, the plume feature detectors use various approaches for identifying the weak reflectivity thin line echo characteristic of an aerosol release plume from a crop duster. However, two of the detectors (Thin Line Suppression and Multipath) are not looking for plumes but are instead looking for other thin line echo features that can mimic plume signatures and trigger false alarms.

Table 3 provides brief descriptions of each of the aerosol plume feature detectors and detailed descriptions of each of the detectors follow this section. They are presented in the order that they are run. The order is significant as several of the detectors require information that is computed in prior detectors.

TABLE 3
Plume Feature Detectors

Feature Detector	Description
Initial Plume	Identifies thin line echo in the region immediately behind the aircraft.
Thin Line Suppression	Identifies pre-existing reflectivity thin line features ahead of the aircraft and generates disconfirming interest in those regions.
Plume Anticipation	Global anticipation interest value fluctuates according to computed background reflectivity statistics in order to regulate overall detection sensitivity. Provides localized anticipation interest boosting based on locations of prior plume detections.
Plume Reflectivity	Looks for persistent thin line echo only within area of prior plume detection. This helps maintain detection length as aircraft moves away from initial release.
Frame-Integrated Plume	Temporally integrates reflectivity in a rectangular frame region behind the aircraft to enhance the signal from the plume with respect to the noise background. Pattern template then identifies the resulting thin line of enhanced signal.
Reflectivity Difference	Looks for thin lines of positive reflectivity difference after subtracting a prior reflectivity image from the current one.
Multi-path / Sidelobe	Identifies radial and cross-radial streaks of enhanced reflectivity associated with multi-path and sidelobe returns from the aircraft.

4.6.1.1 Initial Plume Detector

The Initial Plume detector identifies reflectivity thin lines in the initial plume release region behind the aircraft. FTC processing is performed on the normalized reflectivity image and is restricted to a 10x20 pixel (1.5 km x 3.0 km) rectangular frame region behind the aircraft. The sub-image frame is rotated to an aircraft coordinate system such that the image Y-axis is aligned with the aircraft heading. The FTC template is rotated from 155 to 205 degrees at 5 degree increments to seek the best match over plume orientations that may differ slightly from the aircraft heading (a cross-wind may skew the plume orientation, for example). The resulting interest image, which is in the aircraft coordinate system, is rotated back to the original Cartesian image coordinate system for later fusion with the other feature detector interest images. Figure 23 summarizes the Initial Plume detector processing.

Initial Plume Template
(Rotated 155 to 205 deg)

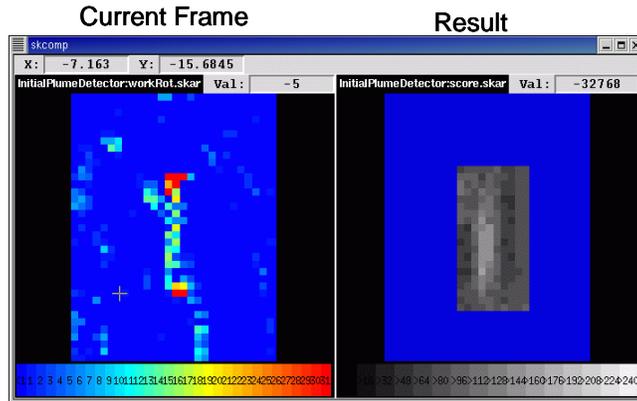
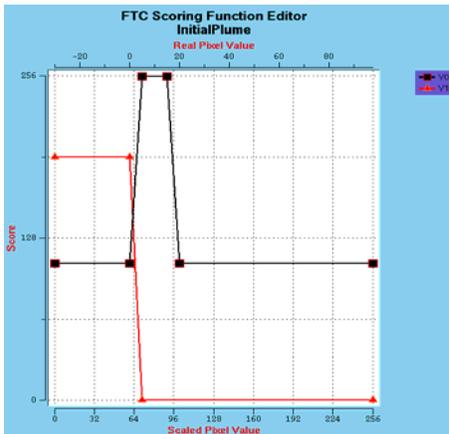
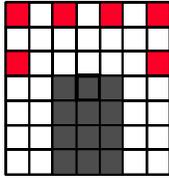


Figure 23. Summary of Initial Plume detector processing. The image labeled “Current Frame” is a 10x20 sub-image of the normalized reflectivity image corresponding to the aircraft location and its trailing region, and is oriented along the aircraft heading. The resulting interest image from application of the scoring functions and template is shown in the image labeled “Result”.

4.6.1.2 Thin Line Suppression Detector

The Thin Line Suppression detector applies FTC pattern-matching to identify pre-existing thin line echo features in regions ahead of the aircraft based on its current trajectory. Conceptually, this is done by taking the rectangular frame processing region as defined in the Initial Plume detector (see Section 4.6.1.1) and “flipping” that analysis region to lie ahead of the aircraft instead of behind it. The functional template that is used is essentially the same one that is used for the Plume Reflectivity thin line detector (described later), except that the confirming interest scores that are reversed to disconfirming interest scores. The resulting interest image has dark gray and black (disconfirming interest) regions instead of white ones wherever thin line features are found.

The localized suppressive interest generated by the FTC operation is then superimposed on a suppression map image that continually accumulates the regions of suppressive interest from scan to scan. If these were accumulated indefinitely and without modification, then eventually more and more of the suppression map would contain dark areas of disconfirming interest as the aircraft moves. Eventually, the aircraft could turn and cross one of these older regions, where the feature that caused the negative interest may no longer be truly present. Plume detection capability in the region behind the aircraft might therefore be inappropriately suppressed. To prevent this from occurring, the prior suppressive interest values in the suppression map are increased by a value of 5 before superimposing the latest suppressive interest regions. This causes the prior disconfirming interest regions to fade toward neutral (128) as time progresses.

Figure 24 shows the template and scoring functions used by the Thin Line Suppression feature detector and example results. In the figure, the current normalized reflectivity image on the lower left shows an aircraft with trailing plume signature. There are also several other echo streaks not associated with the plume. Dark areas in the interest image on the lower right are areas where pre-existing thin lines have been identified. In this example, the location of the real plume signature behind the airplane lies very close to some of the suppressive interest produced by this detector.

Template
(Rotated 0 to 180 deg)

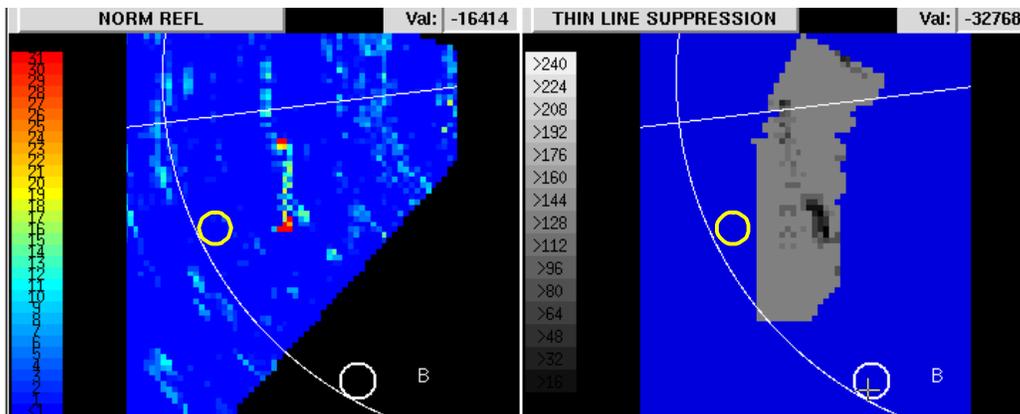
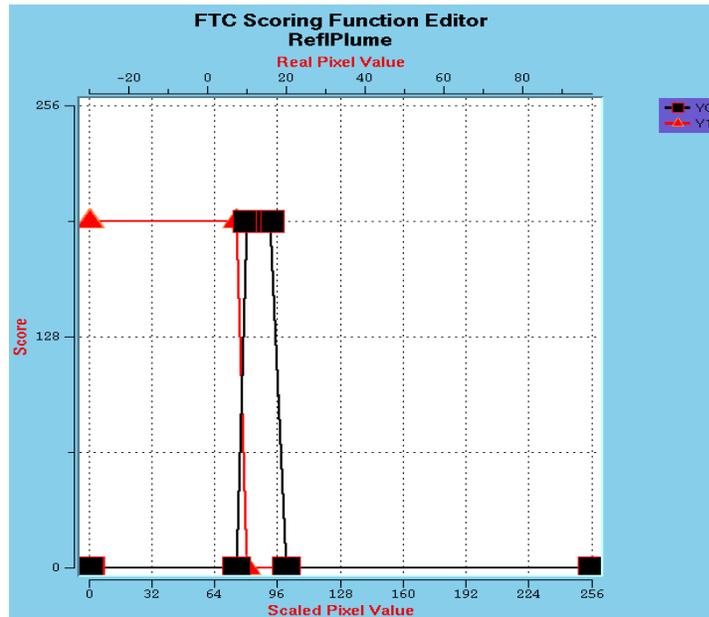
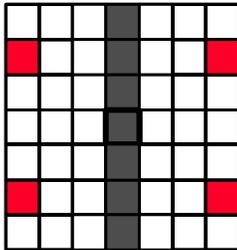


Figure 24. Example results of Thin Line Suppression interest detection. Functional template and associated scoring functions are shown at top. The current normalized reflectivity image is shown on the lower left. Dark areas in the interest image on the lower right are areas where pre-existing thin lines have been identified.

4.6.1.3 Computation of Background Reflectivity Statistics

The characteristics of the reflectivity background play a significant role in plume detection performance. When background reflectivity is high (e.g., due to insect or bird activity), the probability of detection is lower. When there is a lot of background noise or structure (e.g. thermally-induced atmospheric “popcorn” structures, or enhanced ground clutter residue during ducting conditions), there is increased likelihood of false plume detections (see Figure 25). In order to minimize false alarms, the overall plume detection sensitivity is modulated by the Plume Anticipation detector. The Plume Anticipation detector automatically adjusts its global background anticipation interest level based on the statistics of the background (mean and standard deviation), and their corresponding sensitivity bias factors. The background reflectivity statistics need to be computed prior to running the Plume Anticipation detector, but after the Initial Plume detector is run.

The local reflectivity mean and standard deviation are computed from the normalized reflectivity image over a 4-km diameter area centered on the current aircraft location. Prior to computing the statistics, the reflectivity pixels associated with the aircraft are excluded by editing out the pixels in a 9x9 cell window centered on the detected aircraft location. Additionally, in order to avoid possible biases from a released plume, the rectangular frame region used in the Initial Plume detector to define the release area behind the aircraft is used as a mask to edit out the corresponding reflectivity pixels in the working copy of the reflectivity image. The mean and standard deviation are then computed from the remaining pixels in the 4-km diameter circle.

The mean and standard deviation are input to linear falling ramp functions that return bias values between 0.0 and 1.0. The bias values for the two quantities are multiplied together to form the total bias factor. The initial background anticipation interest value will be multiplied by this bias factor to reduce the background anticipation accordingly. Figure 26 shows the ramp functions used to compute the bias factors, and illustrates their application to an example reflectivity image shown in the lower left. In this example, the global anticipation interest is reduced (multiplied) by a bias factor of 0.7.

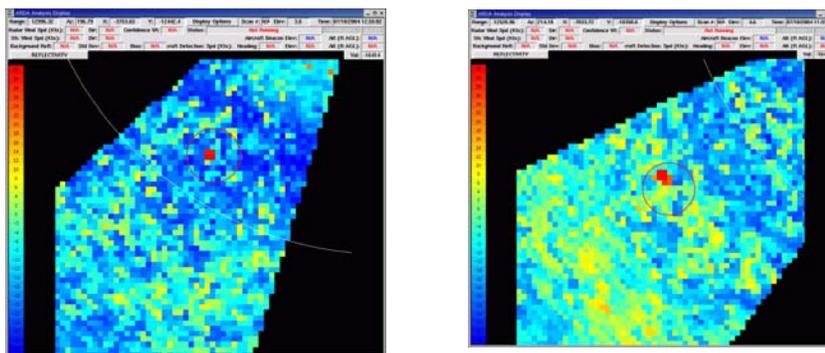


Figure 25. Examples of reflectivity images containing large amounts of atmospheric structure and/or ground clutter. The aircraft can be seen as a bright red peak in the two images.

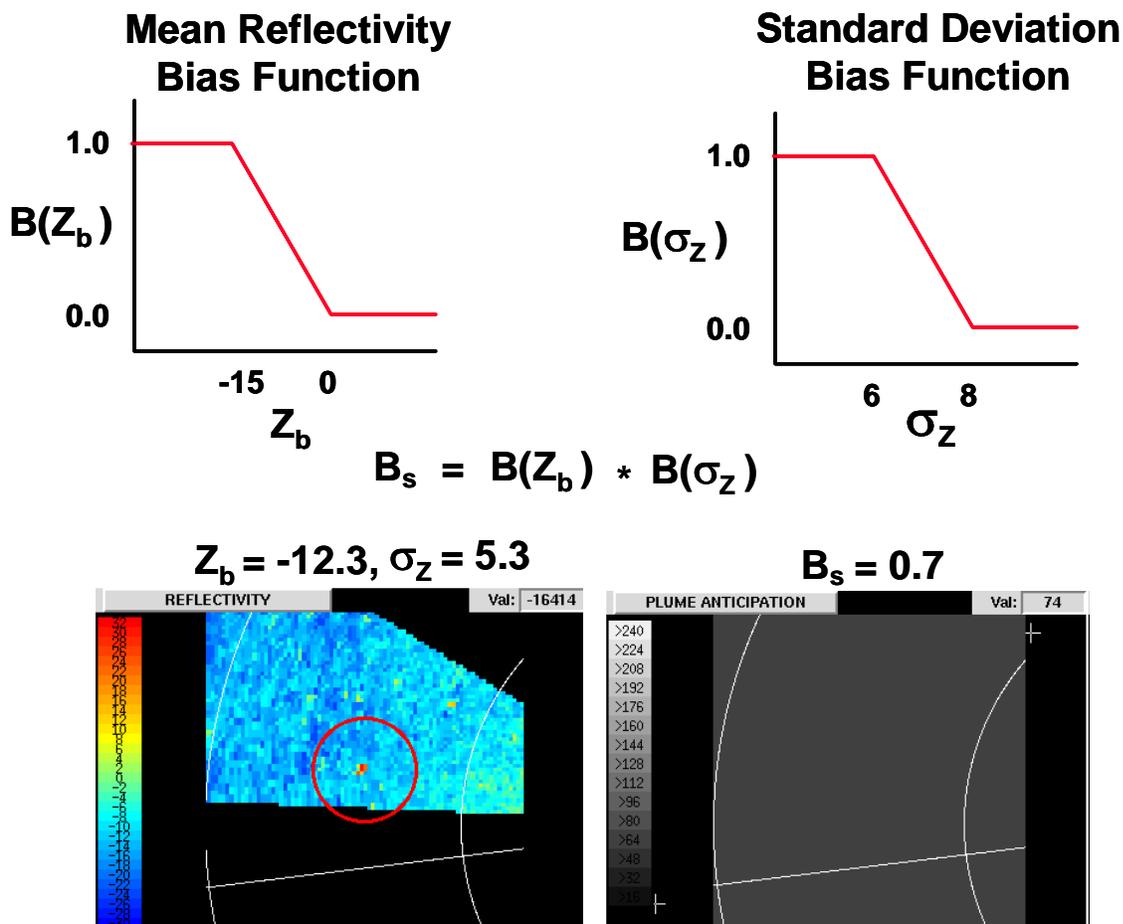


Figure 26. Illustration of background reflectivity sensitivity bias computation.

4.6.1.4 Plume Anticipation

In similar fashion to the aircraft anticipation interest image, the plume anticipation interest image provides the mechanism for dynamic sensitization of plume feature detection. When the plume anticipation interest is averaged against the pre-combined plume feature detector interest, the feature detector evidence can be supported or suppressed based on situational and environmental context.

As is the case with the aircraft anticipation interest image, an interest value of 64 is defined as the initial base value to use for the global plume anticipation interest level, but this value can be overridden in circular target zone regions as defined in the ARDA targZones.param file. Again, for the OKC tests, the entire 50 km ARDA processing range was defined as a single target zone, so the global plume anticipation interest value was effectively set to 96 everywhere.

Once the initial global background interest value has been painted onto the anticipation interest image, the initial values are adjusted by the environmental background bias factor previously computed as described in Section 4.6.1.3. The adjusted values are computed by performing pixel-wise multiplication of the initial interest values by the constant bias factor. Next the plume detection history is checked to see if any plumes were detected during the previous two scans. Plume anticipation interest values are locally boosted to values of 128 in regions where the plume was detected on only the most recent scan, and to values of 192 in regions where the plume has been detected on two prior scans. In this way, evidence from the plume feature detectors is supported by the averaging of boosted anticipation interest values in regions of high expectation based on prior detections. Figure 27 is an example of a plume anticipation interest image showing a region of localized boosting based on prior plume detections.

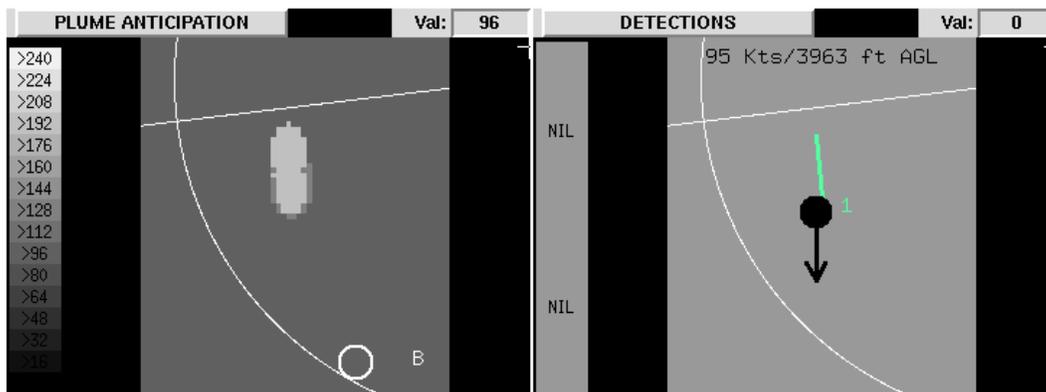


Figure 27. Example of Plume Anticipation interest image. The mid-level gray and white areas in the plume anticipation image (left) are regions corresponding to prior detections where expectation is high. Current plume detection that was supported by anticipation is indicated by the green segment in the image on the right.

4.6.1.5 Plume Reflectivity Detector

The Plume Reflectivity detector looks for reflectivity thin line signatures in areas where a plume has been detected previously. Whereas the other feature detectors operate only within the region immediately behind the moving aircraft, this detector examines regions where prior plume detections were made (not necessarily directly behind the current aircraft position) in order to check for continued existence of plume signals and to help maintain the fullest overall possible detection length.

The previously run Plume Anticipation detector provides the needed map of prior plume detection regions. A copy of the Plume Anticipation interest image is obtained and the image copy is thresholded at 128 (pixel regions of prior plume detections which have values of 128 or 192 will survive the thresholding). The region of above-threshold prior plume interest is then expanded through morphological dilation with a 9x9 pixel expansion kernel. This expanded region of prior plume detections is applied as a mask to the normalized reflectivity image that is input to FTC. Pixels falling outside of the mask region are ignored during FTC processing. The FTC template encodes a thin line echo pattern consistent with that of a plume signature (see Figure 28). The template is rotated through all possible orientations (180 degrees of rotation with the symmetric template) to produce the initial interest image. Following the initial pattern match, the interest values are adjusted based on the template rotation angle at which the best match interest score was achieved. Interest values whose best-match template orientation differs by more than 30 degrees from the orientation of the prior detected plume are reduced by 0.5.

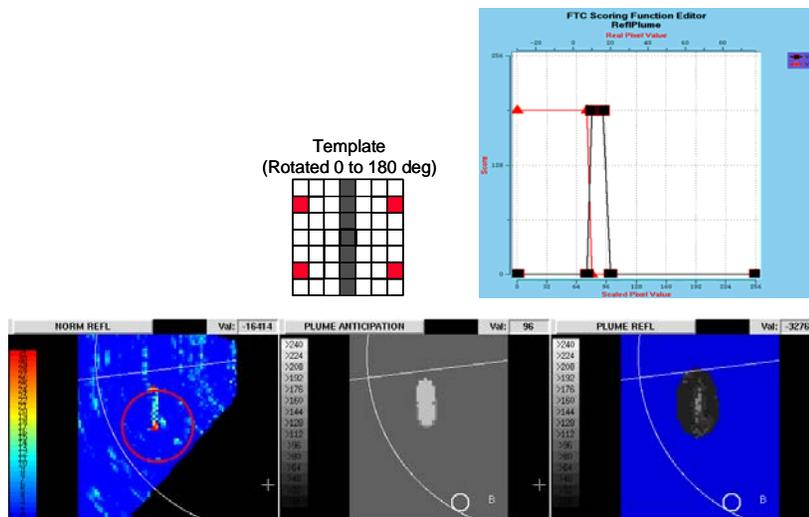


Figure 28. Illustration of Plume Reflectivity detector processing. FTC thin line pattern template and scoring functions are shown at top. Images at bottom from left to right are: (a) Input normalized reflectivity, (b) Plume Anticipation image showing region of prior plume detection where FTC will be applied, and (c) Resulting interest image.

4.6.1.6 Frame-Integrated Plume Detector

The Frame-Integrated plume detector employs time-integration of the reflectivity values in a rectangular “frame region” following the aircraft in order to increase the signal-to-noise contrast of the aerosol plume with respect to the background. A thin line template identifies the ridge of enhanced integrated signal associated with the plume (see Figure 29).

A key element is the construction, placement, and orientation of the frame region used to accumulate the reflectivity values for the temporal integration. The frame is an 11x21 pixel (1.6 km x 3.2 km) box that is positioned in the image over and behind the detected position of the aircraft, and is oriented with its major axis aligned along the current aircraft heading.

To integrate the values, a sub-image of the current normalized reflectivity image encompassing the aircraft and its trailing region is extracted and the pixels are rotated to the aircraft coordinate system such that the aircraft is located top and center. The 11x21 integration frame is then placed on the rotated sub-image and the underlying reflectivity values are retrieved and stored in a temporary array (“Current Frame” of Figure 30). Next, the temporally integrated frame values associated with the prior aircraft detection are retrieved from the history (“Prior Frame” of Figure 30), and the integrated reflectivity average is updated at each pixel in the current frame region using the following weighted average formula:

$$\textit{Integrated value} = ((2 \times \textit{Prior Integrated Value}) + \textit{Current Value}) / 3$$

The integrated sub-image values (“Integrated Frame” of Figure 30) are then used as input to the FTC operation, which identifies the temporally enhanced line of reflectivity associated with the aerosol plume (“Results” of Figure 30). The resulting interest image pixels are then rotated back to the original image orientation so that the values can be fused with other feature detector interest values later during interest combination.

Before exiting the detector, the freshly integrated reflectivity values in the current frame region are stored as a “plume frame” attribute of the currently detected aircraft so that it can be retrieved for temporal integration the next time the Frame-Integrated Plume detector is run.

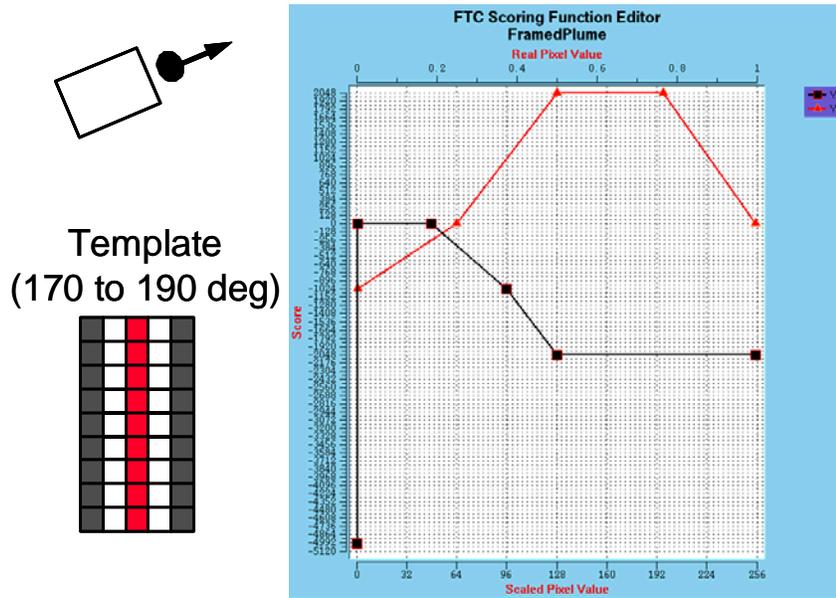


Figure 29. Functional template for Frame-Integrated Plume detector. The template is rotated from 170 to 190 degrees at 5 degree increments with respect to a rotated coordinate system aligned with the aircraft heading.

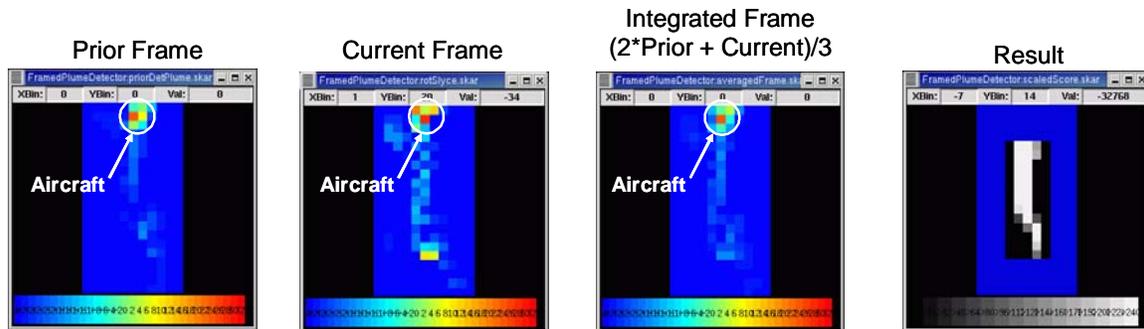


Figure 30. Illustration of Frame-Integrated Plume detector processing stages.

4.6.1.7 Plume Reflectivity Difference Detector

The Plume Reflectivity Difference Detector applies FTC to a reflectivity difference image to identify newly appearing reflectivity thin lines having orientations that are closely aligned with the aircraft trajectory. Processing begins by attempting to find a suitable prior reflectivity image that will be subtracted from the current reflectivity image. There are two selection criteria for obtaining a prior image: age and elevation angle difference. The selection algorithm searches the reflectivity image history for an image that is between 30 and 90 seconds old and has an associated radar elevation angle no greater than 2.0 degrees from the current scan elevation angle. The most recent image in the history that falls within the age window and meets the elevation angle difference tolerance is selected as the prior image to be subtracted. If no suitable image is found, then the detector returns with an interest image that is set to “nil” everywhere (no opinion).

Once a prior image has been obtained, it is pixel-wise subtracted from the current reflectivity image to create a reflectivity difference image (see Figure 31 for an example). In similar fashion to the Frame-Integrated plume detector described previously, a rectangular sub-image region of the image is taken from a region immediately behind the aircraft and is rotated to an aircraft coordinate system with the image Y-axis aligned with the current aircraft heading. FTC is then performed on this rotated sub-region using the template and scoring functions shown in Figure 31. Only a limited range of template rotations (165 to 195 degrees) are used in order to restrict the pattern matching to line features having orientation similar to that of the aircraft. The resulting interest image is then rotated back to the original image orientation for the final interest image result.

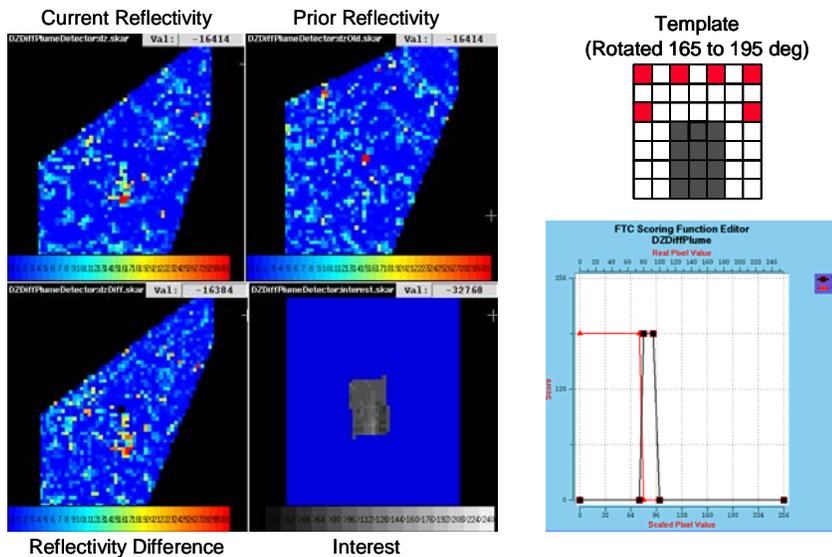


Figure 31. Summary of Plume Reflectivity Difference detector processing.

4.6.1.8 Multipath Detector

Multipath returns from the aircraft can appear as low reflectivity streaks that mimic plume thin line signatures, causing false plume detections. The Multipath detector is a special feature detector that generates disconfirming interest for low reflectivity thin line features that are radially aligned and located at range intervals beyond the range of the aircraft (consistent with multipath returns). The disconfirming evidence generated by this detector can counteract the confirming (but false) interest that might be generated by the plume detectors on multipath returns.

Since multipath returns are always radially aligned, FTC is performed on the original polar coordinate reflectivity image where each row is an azimuth radial and range increases column-wise from left to right. The pattern template then needs to only be passed over the image at a single orientation (90 degrees) in order to find optimal matches. Since the multipath returns only occur at ranges greater than the aircraft range, all polar grid cells with range less than or equal to the aircraft are masked prior to FTC to prevent opinions from being generated in those areas. The resulting polar-coordinate interest image is then resampled to the standard Cartesian grid so that it can be fused with the other Cartesian-based feature detector interest images. Figure 32 shows an example of multipath detection.

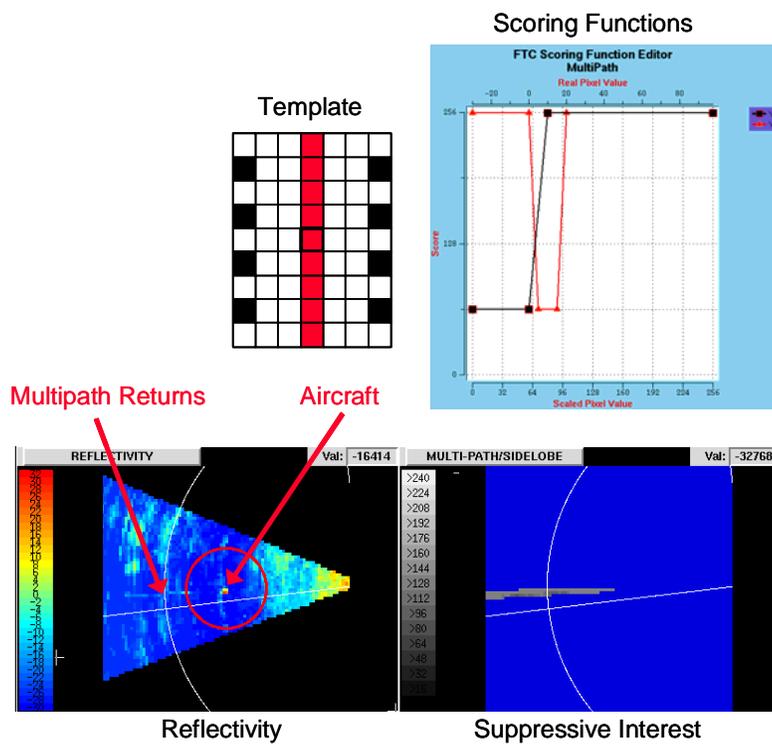


Figure 32. Example of multi-path detection.

4.6.2 Combination of Plume Feature Detector Evidence

The individual plume feature detector interest images are combined using a pixel-weighted average to form a pre-combined feature detector interest image. Table 4 lists the confirming and disconfirming weights that are applied to each interest pixel during the averaging. As with the aircraft feature detectors, the averaging weights are tuning parameters that are readily adjustable. Interest values less than 128 are assigned the disconfirming weight, while interest values greater than or equal to 128 are assigned the confirming weight.

TABLE 4
Feature Detector Weights Used for Plume Interest Fusion

Feature Detector	Confirming Weight	Disconfirming Weight
Initial Plume	1.00	1.00
Thin Line Suppression	0.00	0.75
Plume Reflectivity	0.33	0.20
Frame-Integrated Plume	0.85	1.00
Plume Reflectivity Difference	0.60	0.55
Multipath	0.00	4.00

The combined feature detector evidence is assigned confirming and disconfirming weights of 1.0 and then averaged against the plume anticipation interest. To allow more fine control over the effects of disconfirming and confirming anticipation, the plume anticipation averaging weights are determined by a function rather than the typical bi-polar weighting scheme (see Figure 33).

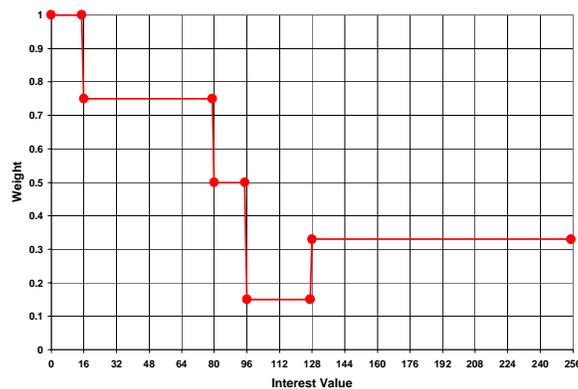


Figure 33. Plume Anticipation Averaging Weight Function.

4.6.3 Plume Extraction

After combining the plume feature detector interest with the plume anticipation interest, an interest threshold of 127 is applied to extract the candidate regions of evidence that may be associated with an aerosol plume. Each connected region of above-threshold interest is analyzed to compute statistics, and the regions are subjected to a length threshold of 0.7 km. The above-threshold candidate plume regions are then converted to plume line segments based on the region major axis length, center of gravity, and orientation.

An attempt is then made to associate each candidate plume region with prior plume detections based on spatial proximity and similar orientation. Correspondence is established if the distance between any two endpoints of the candidate and prior plume segments are within 2500 m (parameter `maxPlumeCorrDistM`) of each other and if the angular difference of the segment orientation differs by less than 45 degrees (parameter `maxPlumeAngCorrDiff`). When correspondence is established, the detection counter (i.e., the number of times the same plume has been detected) for the new detection is set to one plus the previous detection counter value. The current and prior corresponding plume segments are joined in order to “grow” the plume detection length as it is being detected, with the revised segment endpoints assigned to the new detection.

Next, a check is made to see if any prior plume detections should be coasted. For each prior plume detection, the list of new plume detections is checked to see if any of them have been associated with the prior detection under consideration. If a prior detection has no associated new detection, and if the number of coasts associated with the prior detection has not exceeded a `maxPlumeCoasts` parameter (3), then a new coasted detection is fabricated from the attributes of the prior detection. The number of coasts associated with the coasted detection is then incremented.

4.6.4 Detection Reporting Strategy

To reduce false alarms from transient reflectivity features, a delayed “M-of-N” reporting strategy was implemented. The M-of-N output filter requires that a plume be detected on at least $M=2$ of the last $N=4$ scans before an alert will be reported. In order to ensure that the entire detectable duration of the release event has been processed, a wait interval of three consecutive scans without plume detections must transpire following satisfaction of the M-of-N test. The final reported alert contains the two endpoint coordinates of the entire length of the plume trail (as described previously, the intermediate plume segments are internally joined as successive detections are made). Figure 34 is an illustration of the M-of-N reporting strategy for a sample sequence of scans.

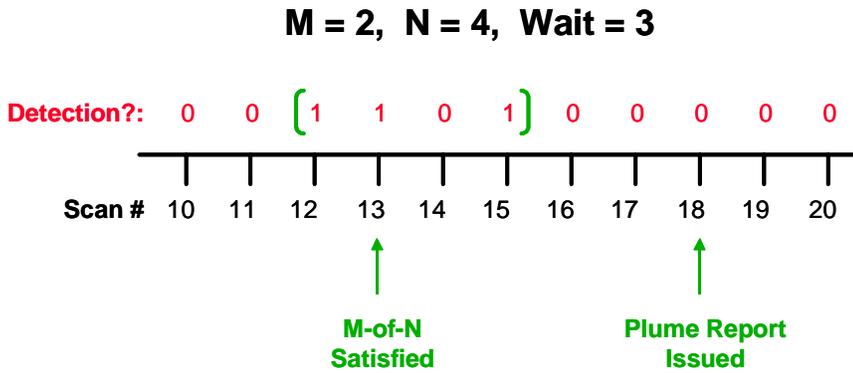


Figure 34. Illustration of “M-of-N” alert reporting strategy.

4.6.5 Plume Detection Output

Plume detection messages are sent via TCP/IP socket communication where they are received by the NBCMsgServer task that reformats the plume detection data into a standard NBC message format, and posts it to one or more Emergency Operations Centers (EOC).

4.7 STATUS MESSAGE OUTPUT

Following the end of each processing cycle (or once every 30 seconds if CBARD processing is not engaged), a status message is output to the CBARD computer. The message contains status bits indicating the health of the software, and also contains a 3-level plume detection confidence measure (low, medium, high). The confidence measure is obtained by inputting the background reflectivity mean and standard computed as described in Section 4.6.1.3 to a 2-D look-up table that returns the confidence value (Figure 35). The table partitioning values were empirically derived based on performance during the Canadian River tests preceding the CBARD operational demonstration.

		Mean Background Reflectivity		
		≤ -20 dBZ	> -20 dBZ & ≤ -15 dBZ	> -15 dBZ
Background Reflectivity Standard Deviation	≤ 4 dBZ	High	Medium	Low
	> 4 dBZ & ≤ 7 dBZ	Medium	Medium	Low
	> 7 dBZ	Low	Low	Low

Figure 35. Performance confidence matrix.

5. 3-DAT ALGORITHM DESCRIPTION

3-DAT is responsible for selecting the scan strategy used by the TDWR when CBARD is activated. The algorithm receives 2-D aircraft location data from two sources; 2-D aircraft cue position from the CBARD ASR-9 aircraft tracker, and the aircraft altitude estimated by ARDA. 3-DAT utilizes two different modes of operation that govern the type of scan strategy selected. Each mode is described in detail in subsequent sections.

5.1 SEARCH MODE

This mode is used during the initial acquisition of the aircraft's vertical position and continues until a valid aircraft detection is received from ARDA. Once the detection is made, the algorithm switches into the Track Mode. In Search Mode the algorithm develops a scan strategy that provides radar coverage from the surface to one kilometer using a total of eight scans. Each scan is centered at a different altitude corresponding to the aircraft's horizontal position. If data from the aircraft beacon transponder is provided, three additional tilts are provided at the beginning of the scan strategy. The first is at the reported aircraft altitude, with the remaining two tilts 100 meters above and below the beacon altitude.

5.2 TRACK MODE

In Track Mode the algorithm attempts to maintain a tight "lock" on the aircraft in the vertical. During this mode the algorithm develops a scan strategy that consists of only two radar scans. The two scans are 0.2 degrees above and below the elevation angle that corresponds to the estimated altitude of the aircraft. In the majority of cases this approach provides the smallest difference in elevation angles and, thus, the best tracking of the aircraft during level, straight flight. The method of bracketing the aircraft between two radar scans also allows the algorithm to maintain a track during aircraft turns, climbs and descents. The strategy has proved effective in maintaining a track of the aircraft in all but the most extreme maneuvers.

In some instances, the algorithm may have difficulty tracking the aircraft using only two tilts. Two techniques are applied to try and identify these scenarios and provide a modified scan strategy to help mitigate this problem. First, the closer the aircraft is to the radar, the faster the elevation angle needed to maintain a track can change over a very short period of time. 3-DAT determines the aircraft location relative to the TDWR and if it is in close proximity to the radar (nominally five kilometers) the algorithm switches to a "cone of silence" mode. This mode produces two tilts at elevation angles that intersect the aircraft's forecasted position based upon its location and speed assuming level straight flight. In this mode it is likely that the two tilts are greater than one beam width apart. The second scenario occurs when the aircraft is greater than five kilometers away from the radar but at an altitude that causes rapid changes to the elevation angles needed to maintain track. In this scenario the forecast position is computed and the

radar tilt that intersects this location is modified to center at a future aircraft position while leaving the initial tilt elevation unchanged. This approach provides the best tracking during level, straight flight and during aircraft turns, climbs, and descents.

In addition to selecting the appropriate elevation angles needed to track the aircraft, the algorithm also selects sector azimuth angle values in an attempt to keep the aircraft within in a wedge of data. The ability to track the aircraft within a small sector allows for sampling the aircraft position and possible aerosol releases at a greater temporal resolution. The selected sector azimuth values are based on the aircraft's position relative to the TDWR, as well as the location of a potential aerosol plume behind the aircraft. The future positions are computed by projecting the plane position forward at its current speed and rotating the aircraft's heading in 10 degree increments. After the position information is calculated, the scan sector is chosen that contains all the computed azimuth values. The scan sector is then limited to a minimum of 30 degrees. If the sector is larger than 180 degrees, the algorithm switches to a 360 degree scan mode. The rotation rate for all the radar scans is 15 degrees/second.

The algorithm continually checks to make sure the 2-D aircraft location information supplied by ARDA is spatially correlated with the ASR-9 aircraft cues. If the distance between the two positions is greater than five kilometers horizontally, the algorithm assumes that the ARDA detection is invalid and returns to the Search Mode. In addition, the algorithm checks the computed elevation angle and limits the radar from scanning below 0.3 degrees and above 20 degrees. Below 0.3 degrees ground clutter returns severely impact the data quality, while above 20 degrees it is very difficult to maintain a track and at that height the aircraft poses little threat to ground-based targets. In the event that the ASR-9 cue data become unavailable while in Track Mode, the algorithm immediately switches to an ARDA "detection only" mode. In this mode, the aircraft's horizontal position (normally supplied by the ASR-9) is obtained from the valid ARDA detection. As long as ARDA maintains a successful track, 3-DAT continuously generates new scan strategies.

Figure 36 plots aircraft positions detected by ARDA from one of the crop duster trials during the operational demonstration as it processed TDWR scan sequences that were issued by 3-DAT. The ARDA detected positions are plotted with purple square symbols. "Truth" positions of the aircraft are based on the ASR-9 cue location and are plotted for reference using blue squares. As can be seen from the plot, ARDA was able to maintain accurate aircraft tracking through the majority of the trial period, including a turn from an easterly to a southerly heading. However, tracking was dropped briefly after the aircraft executed a sharp reverse in direction, but was re-established after the aircraft trajectory straightened out and before the trial ended.

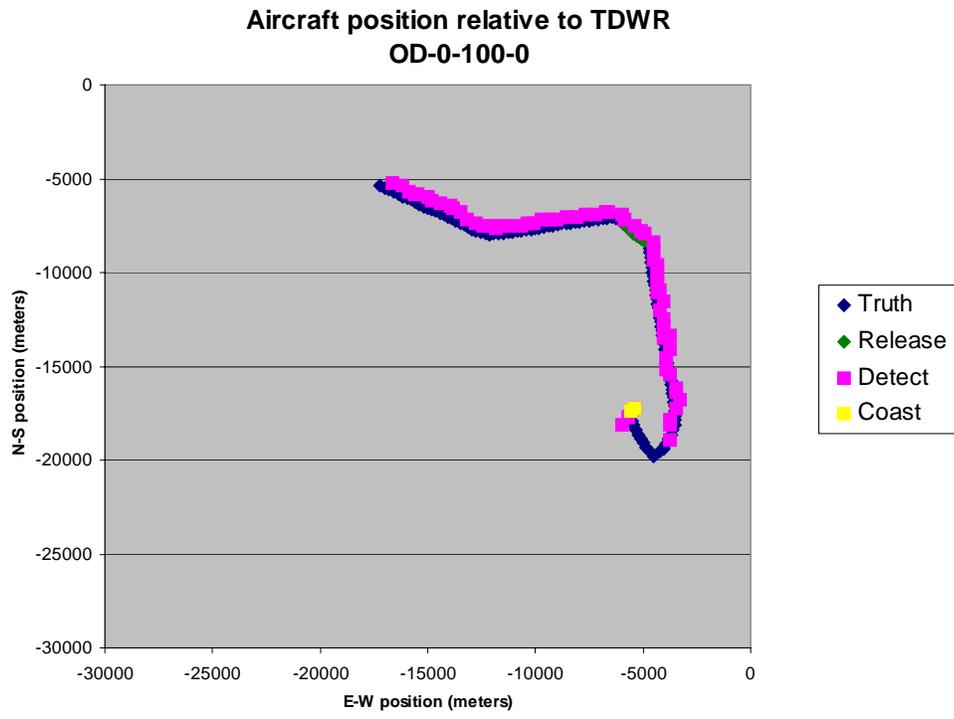


Figure 36. Example of aircraft tracking results.

6. RESULTS

Figure 37 shows an example of the ARDA diagnostic display following processing of a scan of TDWR data during a crop duster release of water. The top row shows the input reflectivity image on the left, followed by interest images generated by three of the aircraft feature detectors. The relatively high intensity peak and trailing low reflectivity thin line associated with the aircraft and plume are readily seen in the reflectivity image. The aircraft feature detectors are all showing strong interest on the aircraft target. The second row shows the Doppler velocity image that corresponds to the reflectivity image on the left. The aircraft and plume can be seen quite easily in this image as well. Interest images from three of the plume feature detectors are shown to the right of the velocity image, and they are each registering different levels of interest on the plume feature. On the bottom left of the third row, the output of the look-ahead thin line suppression detector can be seen. Following that on the bottom row are the total combined aircraft and plume interest images respectively.

The image at the lower right displays the results of the aircraft and plume detection stages. The aircraft is plotted using a filled black circle to indicate the current aircraft position, and an arrow indicating the aircraft heading. The location and extent of the detected plume is indicated by the green line segment. A green segment indicates an intermediate plume detection that has not yet met the reporting criteria. The segment turns red when the alert is actually issued.

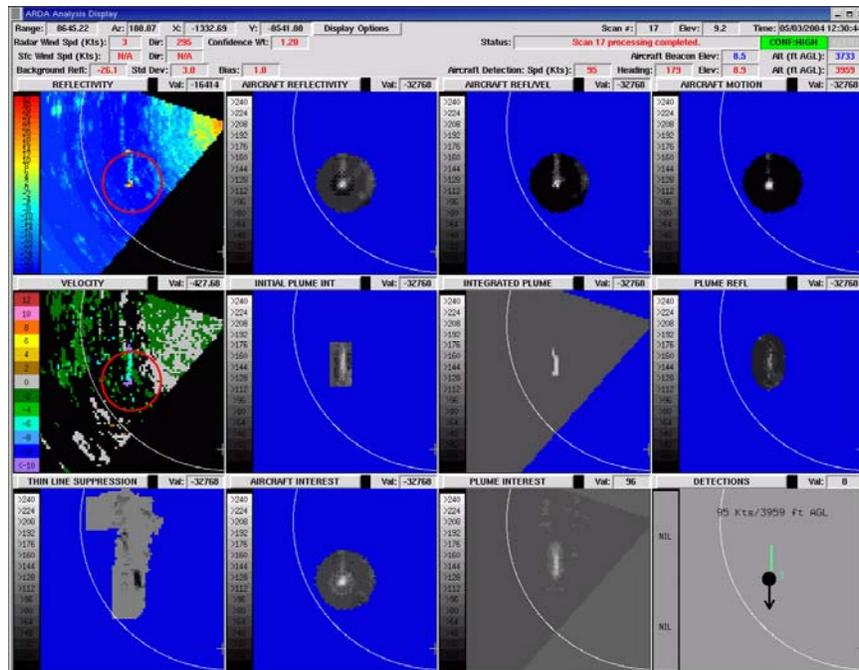


Figure 37. ARDA diagnostic display showing example plume detection.

Table 5 shows results of the current versions of ARDA and 3-DAT operating on data sets collected during the CR7 field test and the Operational Demonstration conducted in August, 2004. During both of these field tests, the aircraft flew at altitudes varying from 500 ft. to 4000 ft., under varying background conditions, and along different flight paths and release locations. Trials were conducted with and without simulatant releases of Kaolin and water. Simulant releases were made at a rate of 100 g/m over a release length of 2 km.

Probability of Warning (P_w) and Probability of False Warning (P_{wf}) were computed by comparing the time of each ARDA plume detection alert against the dissemination periods (for trials in which there was a dissemination). An alert was declared valid if it occurred after the start time of the release and before 3 minutes following the end time of the release. The additional end time allowance was provided to account for continued presence of the aerosols after dissemination and to include tolerance for the 3-scan wait period used in ARDA's delayed alert reporting scheme. An alert was counted as false if it fell outside of the time window during a dissemination trial or if it occurred any time during a non-dissemination trial. Ensemble P_w and P_{wf} statistics were then computed for the CR7 and OD data sets using the following event-based scoring metric:

$$P_w = \frac{\text{Number of trials with at least one valid alert}}{\text{Number of trials with releases}},$$

$$P_{wf} = \frac{\text{Number of trials with at least one false warning}}{\text{Total number of trials (with and without releases)}}$$

The results show that the system correctly reported 80% or more of the events having disseminations that occurred during high confidence background conditions having low mean reflectivity and standard deviation. Although 3 false alarms occurred in supposedly "green" conditions during the Operational Demonstration, subsequent review indicated that one or more of the false detections occurred when the aircraft briefly encountered "yellow" conditions along the flight path. The broad classification of each trial into one of the three confidence categories does not always capture the variations in background that the aircraft may encounter during the trial. As expected, performance declined for trials conducted in the "yellow" and "red" background conditions. In fact, reduced detection capability is practically guaranteed during "red" conditions as ARDA substantially reduces its sensitivity in these conditions.

TABLE 5
Plume Detection Performance Results

		Environmental Confidence Category		
		High $Z_b \leq -20$ $\sigma_z \leq 4$	Medium $-20 < Z_b \leq -15$ $4 < \sigma_z \leq 7$	Low $Z_b > -15$ $\sigma_z > 7$
CR7 Database	Pw	9/ 11 = 0.82	10/ 20 = 0.50	1/ 6 = 0.17
	Pwf	0/ 31 = 0.00	4/ 47 = 0.11	1/ 18 = 0.16
Operational Demo Database	Pw	8/ 10 = 0.80	4/ 12 = 0.33	0/ 2 = 0.00
	Pwf	3/ 17 = 0.18	5/ 26 = 0.19	1/ 6 = 0.17

7. SUMMARY

Within the C-Band Aerosol Release Detector (CBARD) system, the Aerosol Release Detection Algorithm (ARDA) and 3-D Aircraft Tracking and TDWR Scan Strategy Algorithm (3-DAT) are the two principal processing algorithms that work together to track a threat aircraft with the TDWR and detect aerosol plumes disseminated by an aircraft. The two algorithms are described in detail in this report. Key technological achievements include:

1. Use of 2-D aircraft location data from an airport surveillance radar (ASR-9) to cue a TDWR to follow a suspicious aircraft in a special CBARD sector scan mode.
2. Development of a TDWR scan strategy generation algorithm for continuously computing optimal scan sectors and elevations needed to maintain a tight track of the aircraft and nearby regions where aerosol plumes may exist.
3. Development of knowledge-based, radar image processing methods to recognize the aircraft and plume signals in the TDWR radar data.
4. An understanding of the effects of the background reflectivity environment on detectability and relating of the background reflectivity statistics to system confidence.

The CBARD algorithms and system interfaces were continually developed and refined as crop duster release data were collected in a series of supporting field tests in the Canadian River area near Oklahoma City from April, 2003 to July, 2004. The tests culminated in a successful operational demonstration of the complete CBARD prototype in August, 2004.

GLOSSARY

ARDA	Aerosol Release Detection Algorithm
ASR	Airport Surveillance Radar
BT	Bacillus Thuringiensis
CB	Chemical/Biological
CBARD	C-Band Aerosol Release Detector
CR	Canadian River
EOC	Emergency Operations Center
FAA	Federal Aviation Administration
FTC	Functional Template Correlation
HDCBU	Homeland Defense Chemical Biological Umbrella
LLOT	Lincoln Laboratory Optimized Tracker
OKC	Oklahoma City
PEG	PolyEthylene Glycol
Pw	Probability of Correct Warning
Pwf	Probability of False Warning
TDWR	Terminal Doppler Weather Radar
3-DAT	3-D Aircraft Tracking and TDWR Scan Strategy Algorithm

REFERENCES

1. Cho, J., B. Collins, and J. Shaw, 2003: "Simulant Particle Size Distribution: Optical Disdrometer Measurement and Time Evolution Comparison of Radar-Observed Plume Reflectivity with Computer Simulation", MIT Lincoln Laboratory, Report No. 43PM Wx-0092.
2. Delanoy, R.L. , J.G. Verly, and D.E. Dudgeon, 1992: "Functional templates and their application to 3-D object recognition", *Proceedings of the 1992 International Conference of Acoustics, Speech, and Signal Processing*, San Francisco, CA, pp 23-26.
3. Elkin, G., M. Matthews, S. Maloney, and S. Troxel, 2005: "Terminal Doppler Weather Radar (TDWR) C-Band Aerosol Release Detector (CBARD) Prototype System Design", MIT Lincoln Laboratory, Project Report No. WX CB-1.