Enhanced Regional Situation Awareness

Curtis W. Davis III, James M. Flavin, Robert E. Boisvert, Kyle D. Cochran, Kevin P. Cohen, Timothy D. Hall, Louis M. Hebert, and Ann-Marie T. Lind

Airspace protection in the capital area is provided by an Integrated Air Defense System (IADS) created through the coordinated response of U.S. government and local law-enforcement agencies, including the Department of Defense, the Department of Homeland Security, the Federal Aviation Administration, and the Capitol Police. The IADS includes U.S. Coast Guard helicopters, fighter aircraft, and airborne early-warning aircraft cued by surveillance radars. Under Operation Noble Eagle, the response to a threat includes warning flares deployed from fighter aircraft and, ultimately, the use of surface and air-launched missiles. Selecting the appropriate response requires a means for rapidly assessing the aircraft threat. New and existing sensors must be simultaneously cued to the target of interest and integrated with existing sources of information to display a common-air-picture display to support the decision makers. This article describes the development of an Enhanced Regional Situation Awareness system, an integrated sensing and decision support system developed for the complex and busy airspace surrounding the National Capital Region.

To most, the national capital region (NCR) symbolizes the strength and leadership of a nation. In this compact location, a world-class transportation system that includes multiple international airports routinely provides air travel for the President, Vice President, and employees of vital government agencies, as well as throngs of enthusiastic tourists visiting the monuments and museums that capture our nation's spirit and history. However, the tragic events of 11 September 2001 remind us that this location represents a prime target for terrorists, and the transportation system is a means of disguising and delivering a devastating attack.

Historically, defense of the NCR has been the mission of North American Aerospace Defense Command (NORAD), specifically, the Northeast Air Defense Sector (NEADS) in Rome, New York. The role of NEADS has been to maintain surveillance of the airspace over the northeastern United States, give early warning of attack, and provide command and control of defensive forces. NORAD's mission focused on outward detection of Soviet bombers, not on the identification of an aircraft launching an attack from within the United States in a sufficient time period to allow an effective response. Figure 1 highlights key events that motivated changes to the NORAD mission. On 12 September 1994, a stolen Cessna 150 crashed onto the South Lawn of the White House. The impact of this light aircraft inflicted minimal damage but exposed the vulnerability of the White House to an air incident. In response, the Air Surveillance Center was established to protect the White House and to provide an early warning sufficient for the President and other VIPs to be moved to a safe location. The Air Surveillance Center uses existing Federal Aviation Administration (FAA) radars and cameras to observe aircraft in and around the NCR. However, the warning time was reliant on the extent of surveillance coverage, cueing of the cameras, and quality of the

FIGURE 1. Evolution of the NCR air surveillance architecture. Recent events have shaped air defense of the NCR.

radar tracks that were obtained from Federal Aviation Administration radars designed primarily to maintain safe separation between aircraft cooperating with Air Traffic Control (ATC).

The most significant air attack in the NCR occurred on 11 September 2001, when Al Qaeda terrorists hijacked American Airlines Flight 77, a Boeing 757, and crashed it into the Pentagon. This flight had departed from Dulles International Airport and was en route to Los Angeles when it was hijacked. To improve the time available to detect and defend against future attacks from commercial aircraft, important changes were made under Operation Noble Eagle for the NCR. Airspace restrictions were extended to allow greater warning time. NORAD deployed military assets, including Short-Range Air Defense systems comprised of Sentinel radars and Stinger/Avenger missile batteries. Fighter aircraft and helicopters were placed on alert in the region to intercept aircraft violating restricted airspace. However, the NEADS, Noble Eagle, and Air Surveillance Center systems supporting the NCR were not fully integrated to provide a common air picture, nor were the sensors sufficient to provide low-altitude surveillance and target identification. Other than an aircraft intercept, response options were limited to using voice communication to get the attention of the pilot. Overall improvements were needed to enhance the ability of military and civilian decision makers to rapidly detect, identify, and coordinate a response to an atypical aircraft track. This article outlines the development of elements added to the NORAD architecture to increase the situational awareness and expand the palette of decision support tools available to the NCR Operation Noble Eagle console operators and NORAD decision makers.

It is important that these tools allow any possibly relevant information to be sifted through and presented quickly and concisely. Clearly an important role for automated decision support is to share some of the burden of maintaining a high state of readiness. Due to the critical nature and high cost of missed threats and false alarms, the final decisions for interdiction and setting of state-of-alert levels must remain with the human operator. A principal function of decision support is to alert the operator to the subset of cases in which further investigation is required.

ERSA Overview

The United States Air Force Rapid Capability Office tasked Lincoln Laboratory with developing an integrated sensing and decision support system to enhance the situation awareness of decision makers responsible for protecting and responding to attacks from the complex airspace surrounding the NCR. Improved aircraft surveillance and a means for integrating the information available from government and non-government databases were to be combined into a common air picture that also automatically detects and highlights aircraft behaving in a suspicious manner. A means was also to be provided for quickly and effectively sending a visual warning to pilots of the aircraft violating airspace restrictions, thus reducing the delays and confusion sometimes incurred during radio communication.

The Enhanced Regional Situation Awareness (ERSA) system improved four major air defense areas: airspace surveillance, threat assessment and decision support, distribution of a common air picture to multiple agencies, and new ways to respond to aircraft violating the NCR. The key components of ERSA are shown in Figure 2.

Table 1. Glossary of Acronyms

ADIZ	Air Defense Identification Zone
AIMM	Altitude Inference using Multilateration and Multiangulation
ARSR	Air route surveillance radar
ASR-9	Airport surveillance radar
CD2	Common digitizer
COTS	Commercial off-the-shelf
DARPA	Defense Advanced Research Projects Agency
EO/IR	Electro-optical/Infrared
ERSA	Enhanced Regional Situation Awareness
FAA	Federal Aviation Administration
FRZ	Flight Restricted Zone
NCR	National Capital Region
NORAD	North American Aerospace Defense Command
RRDL	Remote radar data link
TADIL-J	Tactical Digital Information Link
VWS	Visual Warning System

A Complex Airspace

The NCR comprises a significant number of high-value assets that would sustain significant damage if struck by an aircraft. Identifying aircraft threats requires a thorough understanding of airspace restrictions and acceptable practices. Figure 3 shows the three separate levels of restricted airspace that are currently in effect. These are known as Prohibited Area 56 (P-56), the Flight Restricted Zone (FRZ), and the Washington Air Defense Identification Zone (ADIZ).

Before the attack on the Pentagon, civilian aircraft were not permitted to fly within Prohibited Area 56. P-56 consists of two small regions, an approximately three-square-nautical-mile area that encompasses the National Mall, the White House, the Capitol building, and the Supreme Court, and a one-nautical-mile-diameter region around the National Observatory. Both portions of P-56 extend from the surface to 18,000 feet above mean sea level.

After 9/11, a Notice to Airmen (NOTAM) was issued that added the FRZ and the ADIZ. The FRZ includes the area around the heart of Washington, varies in radius from 12 to 15 nmi, and encompasses the Reagan National Airport (DCA) and Andrews Air Force Base (ADW). The ADIZ surrounds the FRZ and extends to a distance varying from approximately 30 to 45 nmi from the center of Washington. Dulles International (IAD), Baltimore-Washington International (BWI), and numerous smaller airports lie within the ADIZ. Both the ADIZ and FRZ also extend to 18,000 feet in altitude. To enter the ADIZ or FRZ, aircraft must first obtain authorization from ATC and must meet certain operating requirements.

Even with these airspace restrictions in place, there is still a high density of air traffic in the NCR. Figure 4 shows one hour's worth of beacon tracks in this region on a typical day. During the period shown, there were 1300 beacon tracks within the ADIZ, with some of the traffic passing near the heart of Washington. Many of these close-proximity tracks are unavoidable, since departures and arrivals to Reagan National Airport, which is situated less that 2 nmi south of P-56, pass very close to numerous high-value targets.

As noted above, one of the requirements to enter the ADIZ is to have a working transponder, but transpon-

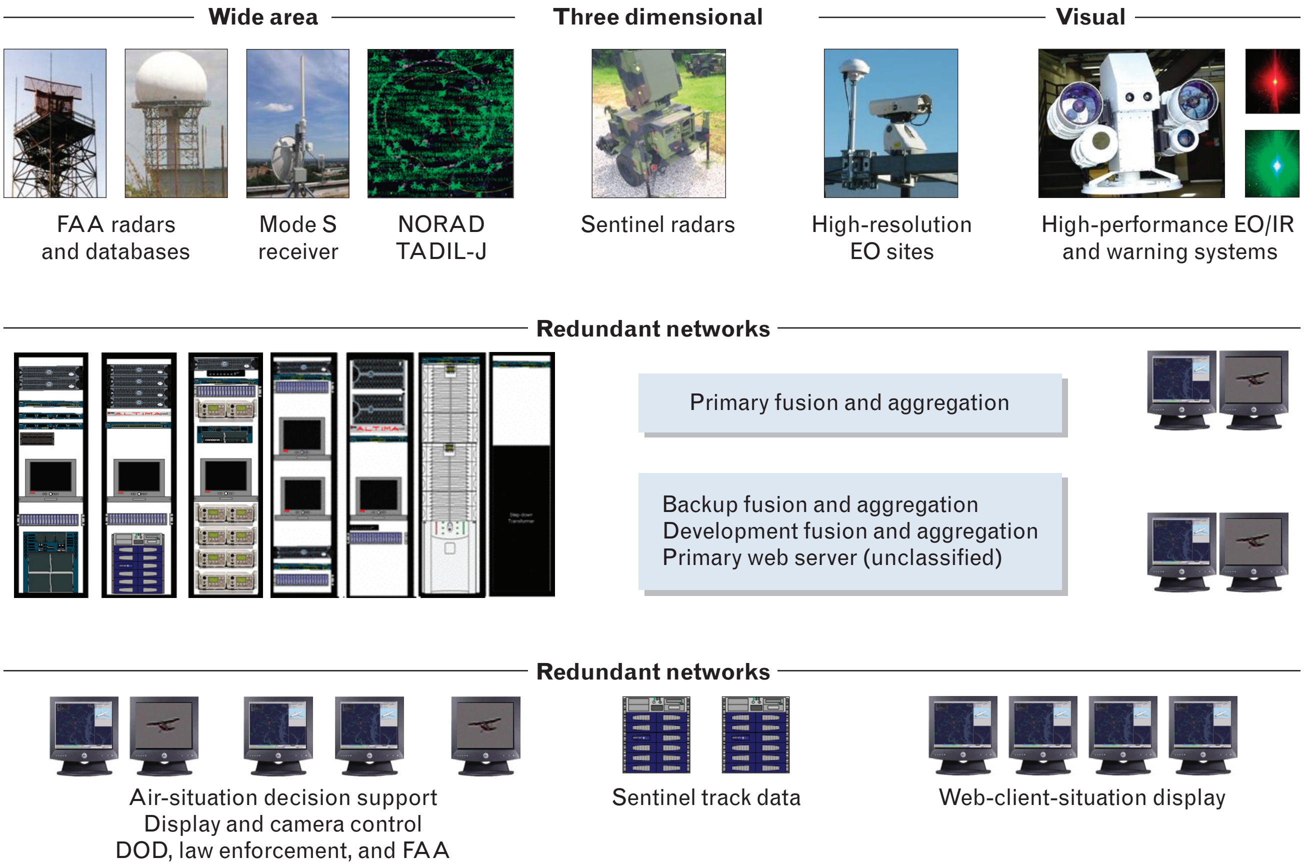

FIGURE 2. Enhanced Regional Situation Awareness (ERSA) system elements. Data are collected, compiled, and distributed in the first two layers. The users then evaluate the data and respond.

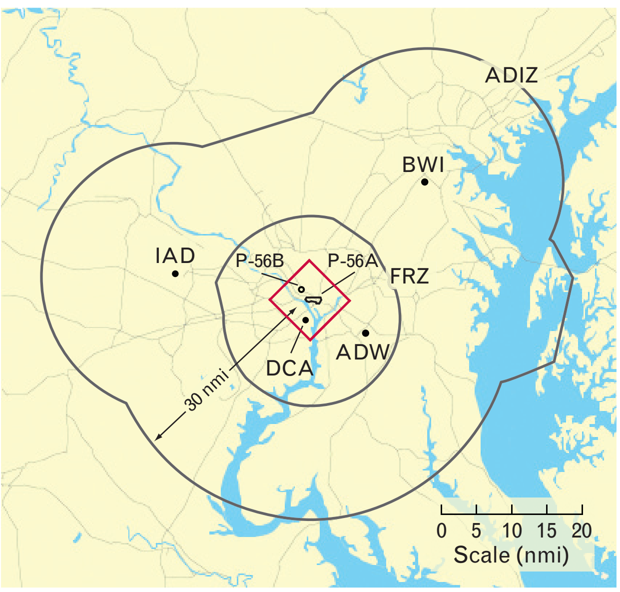

FIGURE 3. National Capital Region. The outer ADIZ and inner FRZ boundaries surround the two parts of the Prohibited Area 56 (P-56).

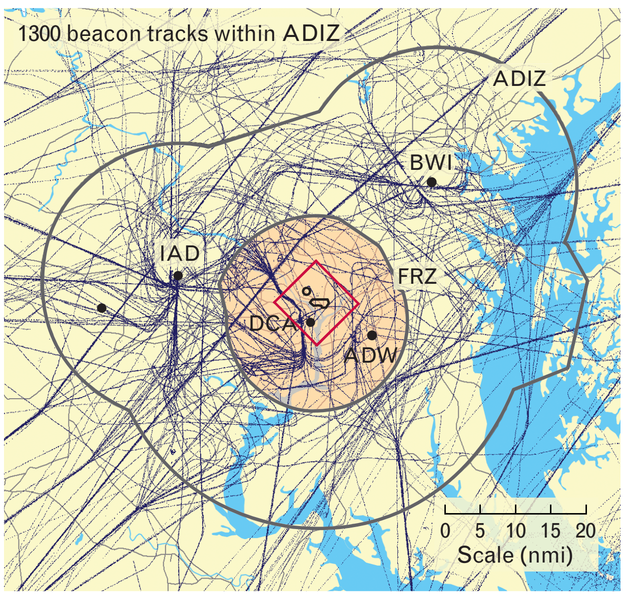

FIGURE 4. One hour's worth of beacon tracks in the NCR. Despite flight restrictions, there is a high density of aircraft traffic within this region.

ders occasionally malfunction. An example that made headlines occurred on 9 June 2004, the day of the funeral for President Reagan. On that day, a King Air 200 aircraft entered the ADIZ with a malfunctioning transponder. The aircraft flew along the trajectory depicted in Figure 5 and landed at National Airport. The pilot of the aircraft was in contact with ATC and received permission to proceed to National Airport despite transponder problems. Unfortunately, other agencies monitoring air traffic in the area were unaware that permission had been granted to allow the aircraft to enter the ADIZ. The aircraft appeared as a radar-only target on a suspicious ground track during a period of heightened alert (many dignitaries were gathered in Washington for President Reagan's funeral). Because of the suspicious nature of the track, the Capitol and Supreme Court buildings were evacuated.

Poor coordination between ATC and other agencies was the primary cause of the 9 June 2004 incident. The lack of a common operating picture, the lack of identification systems (e.g., cameras), and the inability to warn the aircraft (other than by ATC radio) that it had entered restricted airspace were contributing factors. Since that incident, coordination amongst agencies has improved and ERSA has implemented many situation-awareness and intent-assessment tools. The ERSA tools are described in subsequent sections.

Although the 9 June 2004 incident was unusually severe, there have been numerous other instances during which U.S. Coast Guard (USCG) helicopters or NORAD fighter jets have been vectored to intercept suspicious-looking tracks. There are a variety of reasons why an innocent aircraft may appear threatening. In addition to an equipment malfunction, pilot error (e.g., failure to check or understand NOTAMs, incorrect Mode 3/A code entry, and failure to monitor the correct air-ground communications frequency), ATC error, and weather avoidance can all lead to an unexpected ground track. To avoid innocent loss, a layered defense system for the NCR must reliably sort innocent aircraft from true threats.

ERSA Layered-Design Approach

The ERSA architecture has four infrastructure layers connected through a redundant network, which are the sensors, data aggregation and processing, the common air picture, and response. The use of the products from ERSA can be applied to current and future NCR needs.

Sensor Layer

ERSA combines existing FAA ATC radars with new military radars to detect and track aircraft in the region. Sentinels provide accurate 3-D cues for pointing the camera systems, and target-identification capabil-

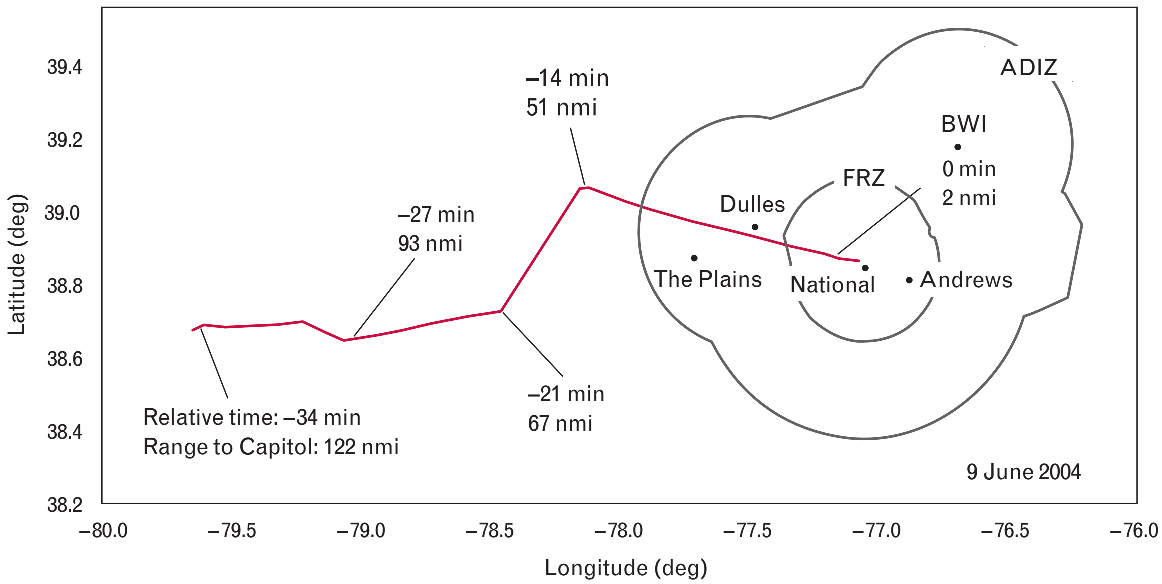

FIGURE 5. Flight path of King Air 200 aircraft that entered the ADIZ on 9 June 2004 with a malfunctioning transponder. Because of the suspicious nature of this track, the U.S. Capitol and Supreme Court buildings were evacuated.

ity through the use of both Identification Friend-or-Foe equipment and radar-signature measurements. An extensive network of electro-optical (EO) and infrared (IR) cameras was deployed across the region. These cameras provide the operator with a capability for visual identification throughout the ADIZ. The control of these cameras includes an autonomous video-tracking capability.

Data Aggregation and Processing Layer

ERSA performs track processing to fuse the radar data to form high-quality metric reports of aircraft positions in the region. The track data are aggregated with weather and other data, then processed through software threat-conditioning logic to assess aircraft compliance with airspace restrictions and to detect unusual behavior. Aircraft that are flagged as possible threats are highlighted on the display, as described below. The aircraft tracks are also used to cue the cameras to the aircraft of interest.

Common Air-Picture Display

The display of aircraft tracks is distributed to civilian and military agencies to aid decision making during the identification and coordinating of the response. An information drill-down capability is provided to allow operators to perform passive intent assessment by viewing detailed information on tracks of interest to look for inconsistencies and reasons for concern. Aircraft approaching high-value assets are highlighted on the display and may be accommodated by video tracking with cameras or a visual warning.

Response Layer

The ERSA Visual Warning System (VWS) gives NORAD another means to respond to airspace violators. The VWS is a set of blinking lights (described in the published NOTAM [2]) directed at and visible only to airspace violators. The lights can be used to warn pilots who have blundered into the restricted airspace and who are not responding to radio calls from ATC. Prior to the deployment of the lights, interceptor aircraft were the only means to warn pilots of a violation and the potential use of lethal force. The sidebar "The ERSA Visual Warning System" gives further details on VWS.

ERSA Layered Decision Support

Since there is a very high volume of innocent tracks and a very low volume of hostile traffic in the NCR, the overall intent-assessment process must be extremely reliable. To responsibly infer intent, decision makers must rapidly collaborate across government agencies and assimilate all information available on the target of interest. Innocent loss must be avoided while false alerts and the overuse of military aircraft need to be minimized. To meet these goals, ERSA implemented a layered decision support architecture that includes passive monitoring of the airspace as well as active interaction with potential threats in order to reliably determine intent.

Passive Decision Support

In this context, passive decision support refers to actions taken to infer intent that do not require direct interaction with pilots (beyond interactions normally required for standard air traffic management operations). Examples of passive intent assessment include monitoring pilot-to-ATC communications, identifying aircraft through the use of EO or IR camera systems, and using special logic that analyzes radar tracks to identify unusual and potentially threatening behavior.

Monitoring pilot-to-ATC communications allows an operator supporting air defense to determine which aircraft are operating abnormally in some manner. Examples include aircraft that are off course, fail to respond to ATC instructions, or report problems such as mechanical failure or turbulence. The use of EO and IR cameras is also a valuable passive intent-assessment tool for daytime and nighttime monitoring. A potential use would be to assess whether the aircraft appears to match the description available in a flight database. Should a decision maker elect to vector a fighter or helicopter to intercept an aircraft, knowledge of the aircraft type facilitates the interceptor's search for the track of interest. In many cases the identification may allow the operator to decide that the track is not a threat.

For ERSA, special intent-assessment logic was designed and implemented to analyze radar-track and other associated data to identify potential threats. The ERSA system receives radar feeds from FAA as well as military radars. Detections received by the network of radars associated with ERSA are processed through a track fusion engine to form composite tracks. Whenever a track update is received from the track fusion engine, the intent-assessment logic decides whether an alert should be issued and, if so, what information should be provided to an operator to help comprehend the situation.

There are myriad checks performed by the intent-assessment logic to detect unusual and threatening behavior. The logic begins by determining whether the track falls into one of four categories: (1) aircraft with valid beacon returns (valid-beacon category), (2) aircraft that previously responded to ATC beacon interrogations but are no longer doing so (beacon loss), (3) tracks without associated beacon returns that exhibit unusual dynamics for an aircraft (suspected clutter), and (4) radar-only tracks with reasonable dynamics (radar only).

Different techniques are used to identify potential threats, depending on the track category. Most tracks fall into the valid-beacon category. The beacon identification (Mode 3/A or Mode S code) associated with these tracks provides information on where and how aircraft should be operating. For example, general aviation aircraft flying under visual flight rules (VFR) generally use a Mode 3/A code of 1200—such aircraft are not permitted to fly within the ADIZ. In addition, there are codes recognized internationally that pilots can set to represent problems such as radio failure, hijacking, or emergency.

A series of decision trees were developed to process tracks with valid-beacon information with each decision tree tuned to the pertinent beacon identification. The individual decision trees check to see if a track is about to enter a zone within which it is not allowed. If the track is about to enter and if the violation poses a threat, an alert is issued (in a few cases, alerts are suppressed for false-alarm mitigation; e.g., very slow-moving aircraft on the ground within the ADIZ that squawk Mode 3/A code 1200). In addition to checking to see whether a track has violated a boundary, the logic performs a series of additional checks to identify potentially suspicious tracks. These checks include unusual ground speed, heading, altitude rate, Mode C errors, and beacon identification history.

Turning off an aircraft's transponder is a suspicious action and one that was used by the terrorists on 9/11. A special track category, beacon loss, was included to rapidly detect and highlight aircraft that appear to have turned off their transponder and have a ground trajectory that could be threatening. Logic for this beacon-loss category uses a combination of decision trees and threat-scoring logic to infer intent.

The suspected-clutter category is included to reduce false alerts. In many cases, the dynamics associated with a track are not aircraft-like. If such a track has no associated beacon returns, then it is placed into the suspected-clutter category. The logic for tracks in this category will produce alerts only under limited circumstances.

The final category is for radar-only tracks that do not fall into the suspected-clutter category (i.e., radar-only tracks with dynamics that are reasonable for aircraft). Although aircraft are required to have working transponders to enter the Washington ADIZ, some radar-only tracks sometimes occur. Cameras are a valuable passive intent-assessment tool for radar-only tracks in identifying a track inside the ADIZ. A threat-scoring approach is implemented to passively assess intent.

With this approach, individual features associated with a radar track (e.g., speed, heading, proximity to NCR) are each assigned a threat score. A total score is obtained by summing up the individual threat scores and adjusted on the basis of contextual information. Thresholds are then applied to assign an alert level.

Active Decision Support

There are numerous reasons that an innocent aircraft may appear threatening solely on the basis of passive intent assessment. Examples include transponder, navigation, or communication system failures; pilot errors such as not reading the latest NOTAM; weather avoidance; and mechanical problems. In most cases, actively interacting with an aircraft will help determine the cause of the suspicious track.

Active intent-assessment measures include air-to-ground communications, fly-bys of the potential threat with fighter aircraft, flares, and VWS. Communicating directly to the pilot by radio is probably the easiest and fastest method to help explain an unusual ground track. However, the pilot, especially if lost, may not have the on-board radio tuned to the correct communications frequency. Also, some aircraft are not equipped with an air-to-ground radio. Even if the pilot responds with a reasonable explanation, continued monitoring of the track is warranted because a pilot who had commandeered an aircraft may try to delay response by making up an excuse for erratic behavior. The sidebar "Pilot Voice Authentication for ERSA Decision Support" provides details on the process of pilot identification and terrorist discrimination.

Intercepting a track with a fighter or helicopter is an effective method to capture a pilot's attention. However, a sufficient timeline must be available to allow the intercepting aircraft to take off, reach the target, and as-

THE ERSA VISUAL WARNING SYSTEM

Instrument failure or pilot inexperience may cause a pilot to inadvertently enter the NCR ADIZ or even the FRZ without proper authorization. Because this behavior may appear threatening, a method is needed to communicate to the crew of the aircraft that they are violating the airspace and that they need to take immediate corrective action.

The first option is radio communication. However, many smaller aircraft operating under visual flight rules (VFR) are not required to maintain continuous contact with air traffic control (ATC). In fact, radios are not required equipment on aircraft operating under VFR. In these instances, an alternative approach is needed to communicate a warning to aircraft. The Visual Warning System (VWS) was developed to provide this alternate means of warning pilots to contact ATC and exit the ADIZ.

The VWS consists of a networked node, shown in Figure A, that contains a red and green laser, as well as cameras to point at targets of interest. The lasers provide a conspicuous signal with a narrow beam to ensure that only the intruding aircraft is warned, thereby minimizing the impact on ATC.

The laser system is designed so that illumination levels are eye-safe and non-hazardous at all ranges. Large 11 inch aperture telescopes spread the beam and hence reduce the laser intensity at the aperture to insure that the intensity is below the American National Standards Institute maximum permissible exposure level for these wavelengths and pulse sequence.

FIGURE A. Visual Warning System (VWS) node. Cameras are used to aim the VWS, and red and green lasers illuminate the target through 11 inch telescopes.

To maximize attention getting and warning, the lasers are modulated (turned on and off) in a regular pattern shown in Figure B. The pattern consists of a sequence of three pulses every second: red/off/red/off/green/off. This pattern and rate of flashing have been determined by human factors experiments to be effective as a warning signal. This pattern is also distinct from other light signals currently used by ATC, yet similar to the red-green flashing pattern used by ATC for a general warning signal. Figure C shows a typical view of the VWS from the pilot's perspective.

Only aircraft that are unauthorized or unidentified (e.g. no flight plan, no transponder signal) and unresponsive to ATC voice contact will be visually warned. The ERSA fused radar data provide a cue to point the electro-optical/infrared video cameras of the laser warning device at the intruding aircraft. The cameras are used to track the aircraft to precisely direct the red and green laser beams.

The procedures to be used by pilots when they observe the visual warning are described in a Notice to Airmen (NOTAM) published by the FAA [1]. Pilots are instructed to contact the ATC and turn away from the center of the FRZ when signalled by the warning system.

NORAD operators monitoring the radar and camera displays operate the VWS. If an intruder aircraft is detected (e.g., no beacon or an unauthorized beacon code) within the ADIZ, NORAD coordinates with the FAA (via existing channels) to determine if the pilot is in contact with FAA. If not, the NORAD operator establishes an optical track on the aircraft with the video cameras mounted on the laser warning device and then illuminates the aircraft with the visual warning signal described above. The visual warning signal will continue until the aircraft responds to the visual signal (see pilot procedures described in the NOTAM text). The initiation and termination of the visual warning signal are coordinated with the FAA.

—*Ronald J. Legere*

Reference

1. At https://pilotweb.nas.faa.gov/distribution/home.html, key in FDC for accountability and the NOTAM number 5/4122, and click View NOTAMs.

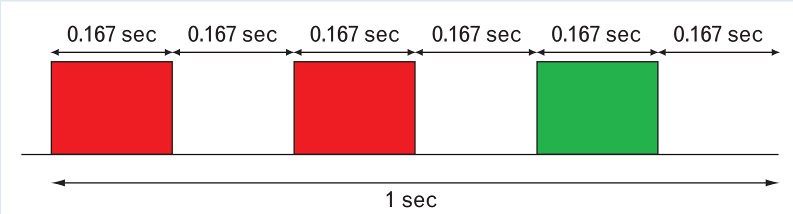

FIGURE B. Pulsed-signal waveform for VWS. This signal ordering is distinct from all other airport signals that pilots receive during routine operations.

FIGURE C. Visual images of pilots' observations of VWS of a single set of laser pulses, taken from a videotape filmed from within an aircraft cockpit.

sess the situation. Interceptor aircraft can be placed on Combat Air Patrol (CAP) to reduce the intercept timeline, but CAPs may not be maintained at all times.

For ERSA, the VWS was developed, tested, and deployed as an additional means of warning pilots that they have entered restricted airspace. VWS consists of eye-safe red and green laser lights that can be directed toward the cockpit of an aircraft. The beamwidth of the laser light is sufficiently small that only the intended aircraft is able to see the red/green flashing lights. VWS is a valuable intent-assessment tool for two major reasons: it does not require any special equipage on the target aircraft and immediately reaches the aircraft of interest.

Example of ERSA Intent Assessment

The 9 June 2004 incident depicted in Figure 5 provides a good example of how ERSA may be used to support the peaceful resolution of an ADIZ incursion. If ERSA had been fully operational at the time of the radar-only aircraft track depicted in that figure, it would have helped resolve that incident in the following ways.

Common Operating Picture. ERSA provides fused radar tracks and a common air traffic display to several users in the NCR. If all participants on 9 June 2004 had the same air picture, the operator who had given the aircraft permission to enter the ADIZ would have known immediately that it was that track that was deemed suspicious by other agencies.

Early Alerting. The passive intent-assessment logic that examines radar tracks would have started to issue alerts when the aircraft was still outside the ADIZ. Although aircraft are allowed to fly in some regions outside the ADIZ without a transponder, the combination of aircraft speed and heading would have made this track sufficiently unusual to automatically trigger an early alert.

Precision Track. Since the aircraft in question did not have a working transponder, the initial track on the aircraft was only a two-dimensional track (FAA radars rely on Mode C transponder replies to obtain target altitude). As the track approached the NCR, it would have flown into coverage of the military radars that have been deployed to provide precision three-dimensional tracking.

Camera Identification. The three-dimensional precision track would have allowed for the pointing of EO and IR cameras to identify the track of interest.

Visual Warning System. If the track continued to approach the center of Washington and was still considered a potential threat, the VWS could have been employed to warn off the pilot. A pilot cognizant of the NOTAMs would have observed the warning, turned, and contacted air traffic control.

Aircraft Intercepts. In the unlikely event that the above steps did not resolve the incident, helicopters and/or fighter jets could have been scrambled to intercept the target. ERSA would have facilitated these intercepts by identifying the type of aircraft that had penetrated the restricted airspace.

PILOT VOICE AUTHENTICATION FOR ERSA DECISION SUPPORT

Positive identification and authentication of pilots in control of aircraft has significant importance for the security of an air traffic system. Since pilots are in radio communication with the tower, using speaker recognition technology [1] to authenticate pilots by voice is a natural approach to addressing this problem. An automatic voice-authentication system that indicates whether or not a particular pilot is the one who is expected to be operating the aircraft can provide valuable information in the context of an air-situation-awareness system. An indication that an unexpected or unauthorized person is speaking on the radio would significantly raise the probability of hostile intent.

The general concept is shown in Figure A. As a pilot is communicating with ATC during normal operations, the radio communications are monitored by a remote system for voice authentication. The remote system will have access to a database associating flight identifications with cockpit crew members (pilot, copilot) and will perform authentication by comparing the incoming speech utterances with speaker models of the pilot and cockpit crew. An indication of whether or not these utterances match the stored models would be provided to a decision maker as part of the overall air-situation-awareness system. The principal utility of the voice-authentication system would be to provide evidence that the aircraft was not being operated by an authorized individual. This might be the case, for example, if a hostile agent were to commandeer an aircraft or attempt to pose as an authorized aircraft operator in order to gain access to restricted airspace. A combination of passive and active authentication could be used. In passive authentication, the system would perform voice authen-

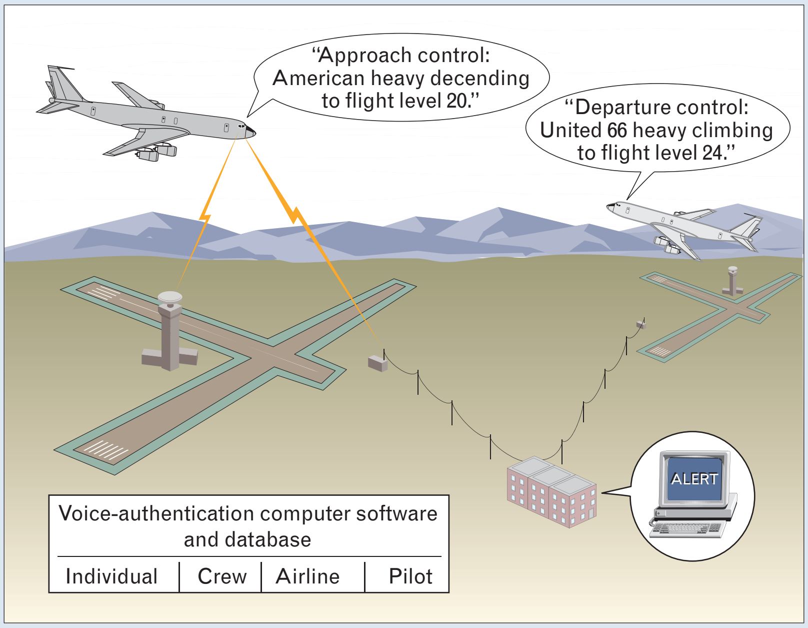

FIGURE A. General concept of pilot authentication by voice. Pilots and crew members voices are compared to a database for authentication.

tication on normal radio communications. It is also possible to equip the cockpit with far-field microphones that would provide a second channel for passive authentication of cockpit crew voices. Such a secondary channel would require either downlinking this audio for remote authentication or performing authentication locally and reporting any anomalous results. In active authentication, a ground operator could request a voice response from the pilot, and voice authentication would be performed on the response. Of course, a failure to respond could be an indication of potential hostile intent.

Potential Voice-Authentication System Configurations

There are a number of ways that a speaker-authentication system could be configured and applied in an air-situation-awareness system. One way is to attempt to match all pilot utterances with a database of models for all pilots. This approach would provide the most comprehensive evidence but, because of the large number of models, would increase the occurrence of missed detections (authentication of an utterance as belonging to an authorized pilot when the speaker is not an authorized pilot). Another approach would be to track the utterances from a particular flight and attempt to detect changes in the speaker. This approach would require the formulation of only a single model, but would be able to detect a change only while the aircraft was in coverage. An intermediate approach would be to track utterances from a particular airline or flight. For example, a model could be generated from the utterances of a flight crew departing on an international flight. When that flight returned from overseas, the authentication system could determine if the same crew were aboard, if another expected crew were aboard (provided that a similar model was stored for that crew), or if a completely new speaker was operating the aircraft. It should be noted here that the data from a voice-authentication system would be used in the context of a larger decision support system. It is very likely not practical that voice authentication be used alone as an indication of hostile intent; rather, the system would be used to provide another piece of evidence. For example, a flight returning from overseas could be exhibiting suspicious behavior (erratic course or lack of transponder replies). If the voice-authentication system also indicated that the speaker was not one of those expected on that flight, it would provide another piece of evidence of the potential for hostile intent.

The issues of how and when to obtain the speaker models for an authentication system, and of how to select the model to be used for a particular authentication, are central to the application of pilot voice authentication. For example, in selection of the authentication model, it may be important to automatically determine the flight identification by a combination of speech recognition on the content of the communications with other information such as radar or beacon tracks, and knowledge of what aircraft are in the airspace under observation. These speaker models will need to be generated and updated from prior speech collected from these crew members. The preferred concept of operations (CONOPS) is for this prior speech to be collected over the

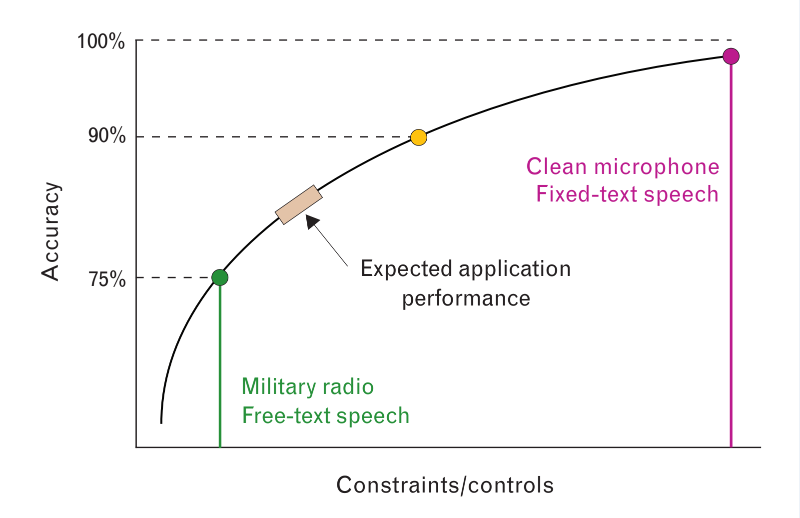

FIGURE B. Nominal trade-off of speaker-authentication accuracy versus data and application constraints and controls.

air, with no special requirements on the pilots for enrollment procedures, which may also require flight identification recognition from speech content. If the match score to the models is below some specified threshold, an alert will be issued, triggering some secondary testing procedure (e.g., further tower questions).

Speech Corpora and Performance Trade-Offs

As with all speech processing applications, the quality and amount of the speech data used for speaker enrollment and authentication are the main dimensions affecting performance. For this application of voice authentication, challenges include variable noise and distortion in the speech transmission from the aircraft to the remote system; the desirability to do training and updating of speaker models on the basis of over-the-air communications; and the effects of misrecognition of flight identification's in pilot speech. On the other hand, the limited population of pilots and the limited vocabulary generally used for pilot communications may be advantages for this application. Currently there are no data collected for this particular application, but there are some relevant speech corpora available from which we can estimate likely performance. These include the Greenflag corpus [2, 3] of military ATC speech; a corpus collected under an earlier DARPA program [4] of pilot-tower communications collected from the Reagan National, Dallas–Fort Worth and Boston airports; and another ATC corpus utilized in a recent Lincoln Laboratory project [5] aimed at recognizing the content of ATC speech in order to associate radio transmissions with radar tracks.

In Figure B we show a nominal trade-off of speaker-authentication accuracy versus application constraints and controls. Constraints and controls generally refer to the how much the application can be designed to expect cooperation of users (e.g., give speaker a fixed-text phrase when prompted), collect high-quality speech (e.g., demand a particular microphone be used) or limit the scope of the problem (e.g., recognize only a small set of speakers with a large quantity of enrollment data). A few representative points from existing speech corpora are shown. The best accuracy is obtained when we have cooperative users and high-quality data. The lowest accuracy happens when we have little to no control over the users (as in passive authentication) and data quality (we must use existing communications channels and equipment). It is expected that the current application of pilot authentication would provide accuracies well above the worst case, but not near the best case.

System Issues

Voice authentication has not yet been applied in the air traffic security system context described here. In order to proceed, we need a system analysis to provide a framework to analyze the feasibility of the concept. This analysis would be done by developing a model of the air traffic flow and tower-pilot communications for a typical airport of interest and analyzing the effects of varying the accuracy (probability of missed detection and probability of false alarms) of a speaker-authentication system. This system analysis would include gathering realistic information on air traffic operations counts, the pace and content of pilot-ATC communications, and the costs associated with handling false alarms from the speaker-authentication system. The key objective of this system analysis and feasibility study would be to develop and compare the potential effectiveness of a variety of system designs and CONOPs for pilot voice authentication, and to recommend a preferred CONOP or set of CONOPs.

—*Clifford J. Weinstein and Douglas A. Reynolds*

References

1. D.A. Reynolds, "Speaker Identification and Verification Using Gaussian Mixture Speaker Models," *Speech Commun.* **17** (1–2), 1995, pp. 91–108.
2. J.P. Campbell, Jr., and D.A. Reynolds, "Corpora for the Evaluation of Speaker Recognition Systems," *Proc. Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP 1999)* **2**, *Phoenix, 15–19 Mar. 1999*, pp. 829–832.
3. R. Ricart, E.J. Cupples, and L. Fenstermacher, "Speaker Recognition in Tactical Communications," *Proc. Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP 1994)* **2**, *Adelaide, Australia, 19–22 Apr. 1994*, pp. 1-329–1-332.
4. J.R. Rohlicek, D. Ayuso, M. Bates, R. Robrow, A. Boulanger, H. Gish, P. Jeanrenaud, M. Meteer, and M. Siu, "Gisting Conversational Speech," *Proc. Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP 1992)* **2**, *San Francisco, 23–26 Mar. 1992*, pp. 113–116.
5. C. Quillen and W. Shen, private communication

ERSA Display-System Overview

Figure 2 reveals the major components in each layer of the ERSA system. In review, operators obtain an understanding of the current airspace by viewing aircraft position symbols formed by a fusion tracker that combines the output of the ATC surveillance radars as well as military radars. The military radars are selectively deployed to improve coverage at lower altitudes and provide the altitudes of aircraft not equipped with transponders. Threat-conditioning logic automatically assesses fusion track and aircraft beacon-code emissions to identify behavior not conforming to normal or allowed operational procedures.

The users themselves must make the time-critical and final decisions in assessing threats and taking action. Behind the displays, the decision support algorithms are sifting through the large volume of air traffic information, generating information for advisories and alerts. The display system is a key component of the decision support system as it interprets the internal products of the decision support algorithms and communicates them to the operator. The display tools provide both a broad view and the facility for the operator to quickly drill down and glean a concise picture from a complex field of information.

The ERSA display uses a combination of symbology, colors, and text messages to assist an operator in determining intent. The display highlights the level of concern by color-coding the aircraft symbol, yellow for moderate and red for severe. Each track has a symbol associated with it with different symbols utilized to indicate the major track types, i.e., discrete beacon, nondiscrete beacon, foreign aircraft (if known), helicopters (if known), and radar-only. If a track is determined to be unusual and suspicious, the color of the icon (and in some cases the icon itself) changes to cue an operator that the track warrants further investigation. Text messages are provided to inform the operator why the logic has determined that a track appears suspicious. The operator can also select the track to obtain additional information on the aircraft, pilot, and route of flight.

Selecting a track also allows EO and IR cameras to be automatically pointed at the target of interest. Aircraft posing a significant concern may be warned by using the VWS described publicly to them in a NOTAM. Pilots complying with the NOTAMs will turn and contact ATC. Pilots ignoring the visual warning by continuing on a threatening route of flight are implied to be hostile and will be intercepted by military and law-enforcement aircraft.

ERSA User Interface

The ERSA user interface was designed to support the time-critical mission of the users tasked with the protection of the NCR. The display of information to support evidence accrual and intent assessment is the product of spiral development. Spiral builds of the system were designed, implemented, and tested to enable the additional functionality of each spiral to be delivered to and evaluated by the users. The functionality for each spiral was documented in a user guide, then training was provided. Training was also a means for obtaining user feedback and suggestions for enhancements to include in future spirals.

A user-centric approach was followed throughout the development of ERSA. Developers visited user sites, discussed user needs with operators, consulted documentation on user tasks and mission, and became familiar with the resources (displays, controls, and communication devices) used by operators. Through an understanding of the current system, appropriate features could be maintained and integrated with new functions.

An ERSA New England demonstration system was developed as a prototype for the ERSA NCR system. The following workstation images illustrate ERSA display functionality. The ERSA workstation shown in Figure 6 consists of a dual-head display and mosaic display provided in a total of four monitors. The common-air-picture radar display is shown in the lower left quadrant and the main video display is shown in the lower right quadrant. There is a single keyboard and mouse to control the displays. Above the radar display and the main video display is the mosaic, which consists of two tile displays showing the output of up to four cameras each. While this is the display configuration used by the NORAD operators, other agencies in the NCR have all or parts of this configuration, depending on their information needs.

The Tactical Display Framework (TDF) developed by Raytheon Solipsys was used as the basis for the radar display. The decision to use TDF was based on three factors. First, it is a commercial-off-the shelf (COTS) product that provides the data fusion capability needed. Second, it provides a flexible framework for data visualization that can be customized to meet user needs.

FIGURE 6. ERSA workstation. Four monitors show an array of images, including radar tracking and camera images. A single keyboard and mouse control the displays.

Third, and very importantly, NORAD already uses the TDF. This common framework made training more efficient.

Through development licenses obtained from Raytheon Solipsys, Lincoln Laboratory produced a highly customized version of the TDF. ERSA customizations include aircraft registration data, three-dimensional models and two-dimensional pictures of aircraft, camera availability and status, and aircraft arrival and departure information for major airports. Custom display filters allow operators to quickly drill down and access information quickly. Other customizations include estimated altitude for targets with no active transponder, custom symbology for international flights, a six-level weather overlay showing convective activity in the area, and the ability for an operator to assess tracks by type, such as birds or small airplanes.

Display Design Considerations

ERSA must facilitate situation awareness and enable users to make timely and correct decisions. A widely accepted definition of situation awareness is “the perception of the elements in the environment within a volume of time and space [level 1], the comprehension of their meaning [level 2], and the projection of their status in the near future [level 3] [3].”

ERSA provides many sources of information to support each of these situation-awareness levels. Intent assessment is ultimately achieved by integrating the essential data available to aid the user in making the best possible determination on whether an aircraft poses a threat. The user interface is the means for conveying this information. For developing the user interface, the following display design considerations were addressed:

Provide a Comprehensive Picture. ERSA provides operators with information from the multiple sources needed for evidence accrual and intent assessment to facilitate accurate and timely decision making. As described above, information is provided from FAA and military radars, EO and IR cameras, databases, side-by-side comparisons of aircraft models to the camera video, and results from the threat-conditioning algorithms. The ERSA user interface integrates this information into a comprehensive view, simultaneously available to various decision makers, without overloading operators with the task of acquiring, sorting, and assimilating information from diverse sources.

Focus User Attention. Customized symbology and color coding of information are used to draw the user's attention to objects of greatest interest. For example, symbols are used to designate type of aircraft (commercial passenger-carrying aircraft, international flights, general aviation aircraft, and helicopter), as well as facilities (e.g., airports, runways, power plants, and sensors). Color coding alerts the user to the threat level. When there are multiple alerts, the color-coded threat levels not only aid in focusing attention, but they also help the user in prioritizing which track should be investigated first. Examples of symbols and the meaning of color-coded alerts are shown in Figure 7.

The display provides selection buttons (the colored on-screen buttons seen at the bottom of the radar display in Figure 6) that enable the user to customize what is shown on the radar display. A user can choose to see only aircraft of specific types (e.g., international flights) and within a certain altitude level of airspace by clicking on specific filter buttons on the radar display. This capability is especially useful if a user receives a verbal communication about a suspected target of interest. Filters provide a way for the user to sort through the myriad of tracks provided by radar and focus on the tracks of greatest interest. In addition, color-coded arrows direct user attention to potential threats and threat levels of tracks currently out of the user's view. That is, if the user has zoomed into an area, threats may be occurring farther out from the current location being viewed. These arrows alert the user to threats that otherwise would have remained unseen. They also alert the user to distant, fast-moving tracks that are headed toward the ADIZ and predicted to violate the ADIZ, but are not shown at the current zoom level.

Enable Rapid Access to Information. Rapid access to controls and the resulting information is essential while minimizing the number of input actions, such as keystrokes, cursor movement, mouse clicks, and embedded menu selections required. Hot keys enable one-click access to the most frequent (e.g., initiating video tracking) and time-critical commands (e.g., initiating use of the VWS). Important commands can be executed from the display that is in use to minimize the amount of cursor movement between displays. For example, a camera can be released from either of the dual-head displays (radar display or the main video display). User preferences for particular input devices are also accommodated. For example, some users prefer using the mouse, while others prefer keystrokes.

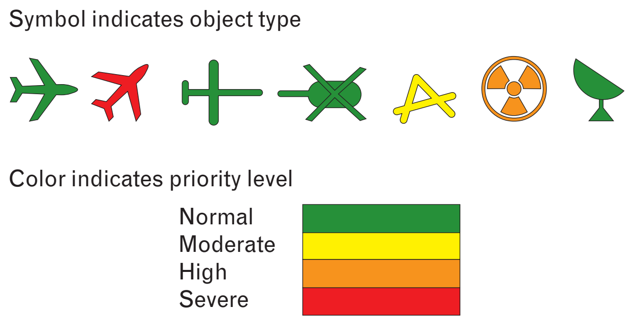

FIGURE 7. Use of symbols and color-coded alerts. These aid users by indicating types of aircraft and facilities and levels of threat.

One example of providing quick access to information is seen in the ability to select a track, and then data tabs are automatically populated with information related to that particular track. As seen in Figure 8, a track has been selected on the radar display (the yellow circle indicates the track that has been selected by the user). In the upper right portion of the display, the user sees that tabs are populated with information related to the selected aircraft. The first four tabs pertain to the selected aircraft of interest and are cross-linked. The tabs provide the available information about the aircraft, an image of a three-dimensional model that matches the aircraft type of the selected track, and a list of cameras that can be selected to view the object. The user does not need to search for this information. Instead it is readily available once a track has been selected.

Provide Shared Situation Awareness. Shared situation awareness is provided through a common operating picture (COP). ERSA is a distributed system; i.e., various users have access to the information through the

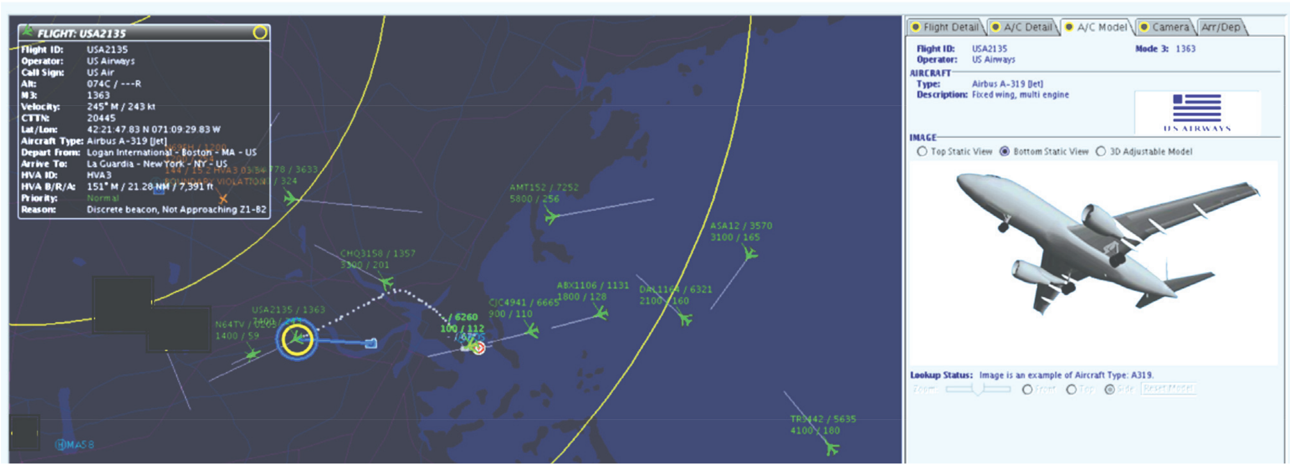

FIGURE 8. Rapid access to information capability. When a specific radar track is selected, it is highlighted in yellow and further information about the aircraft is displayed on the right side of the screen.

ERSA workstation and can see the same picture of the situation. This capability aids in coordination and communication among the various DoD, civil, and law-enforcement agencies. By having ERSA, all users are referring to an aircraft of interest by the same identifying number and are seeing a common air picture and video output. Previously, each user group either had no display of the situation or had a unique/dedicated display not shared with the other user groups. A common operating picture alleviates the negative impact caused by such disparity, namely, confusion, time delay, and error when a user is coordinating a response to a potentially dangerous event.

Provide Appropriate Level of Detail. ERSA supports users who are tasked with the air defense of the NCR, while also supporting users who are tasked with the protection and evacuation of personnel. The tasks of these two user groups are different and, therefore, their specific information needs are different. For example, the NORAD operators want all the information needed in accruing evidence and assessing intent. The details are provided in a drill-down capability that provides a quick indication to alert the user of a threat situation and also provides more details on the threat. Figure 9 shows information represented at three levels of detail. As seen in the figure, the first indicator that the threat level has been raised to severe is indicated by the red aircraft symbol. For more information, the reason for the alert is indicated. In this case the reason is "BOUNDARY VIOLATION." Additional detail on the reason for the alert is seen in the lower right, the Selections Details window. In this case, the alert is "Discrete beacon, Approaching Z1-B1, Altitude below 18,000 ft."

Provide Intuitive User Interface. An easily understandable interface is a key element in facilitating operator learning and ongoing effective and efficient operational use of the information provided by the system. No matter how innovative the technology solutions, if this technology cannot be easily understood and used, then the technology may be underused, used improperly, or not used at all—leading to errors, accidents, and mission failure. ERSA display design focused on making the user interface both useful and usable.

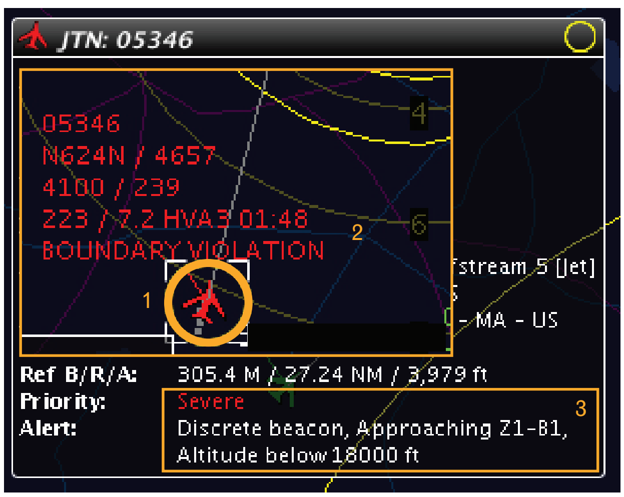

FIGURE 9. Level of information detail. The red aircraft symbol (1) indicates the threat level is severe and (2) shows the cause for the alert to be a boundary violation. The lower right box (3) displays the details of the violation.

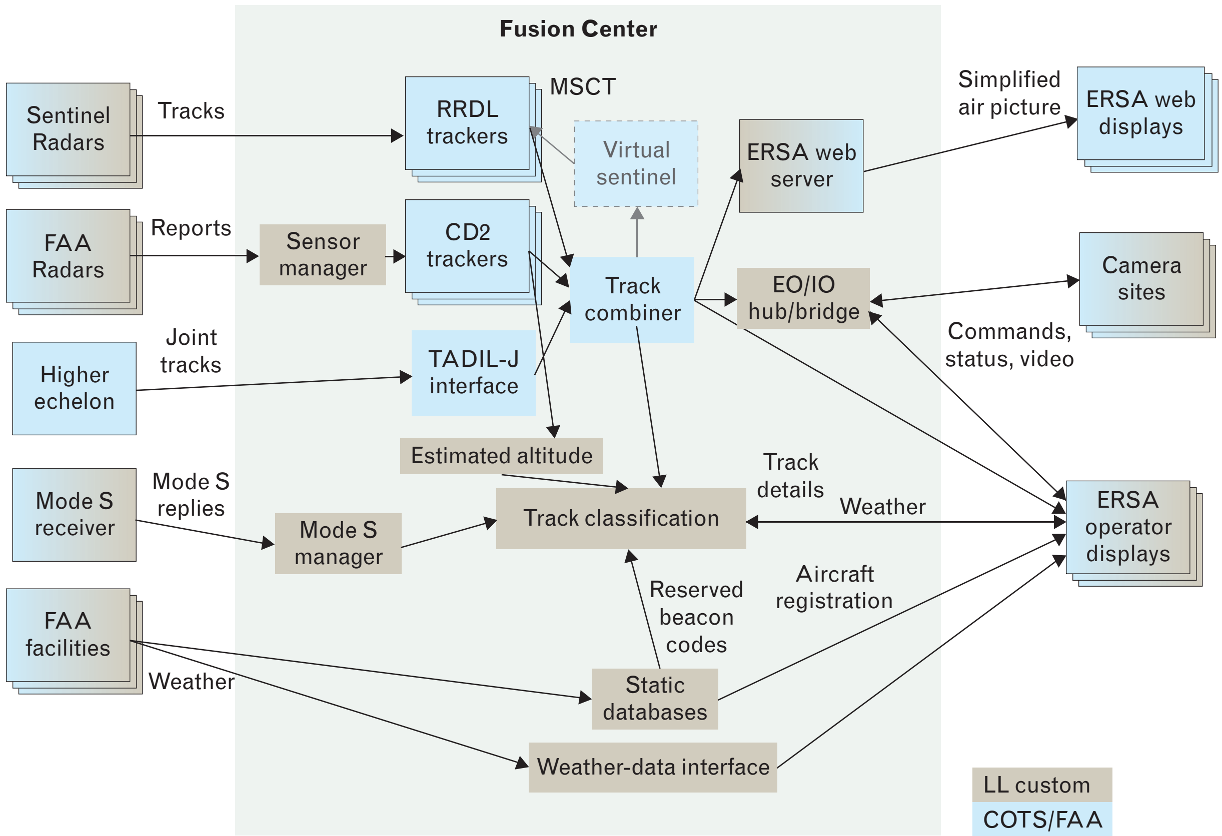

FIGURE 10. Fusion-center data flow. Data are initially handled by the sensor manager. The track classification component evaluates the sensor, Mode S, and static databases, determines threat levels, and presents the data to the operators.

A human factors and software-development team worked together closely to move from design concepts to detailed design documents, to prototypes, and finally to finished, working displays. An ERSA style guide was developed to ensure cross-display consistency. The style guide focuses on major areas of usability, including format and layout, user interaction techniques, widgets, information coding (e.g., symbols and color), and error handling. The guide consists of a set of human factors design guidelines based on Java design guidelines [4] and applicable industry standards and guidelines [5, 6]. The standards take into account the unique aspects of the COTS TDF upon which the ERSA radar display was developed. The development of ERSA design standards ensured that capabilities designed and developed over iterative software spirals provided a graphical user interface consistent in appearance and behavior.

ERSA Fusion Center

An ERSA fusion center comprises COTS hardware and custom software. The ERSA-developed custom software consists of approximately 500,000 lines of code developed to provide radar data processing, recording, and transport; air-picture customizations and threat conditioning; and camera-related functions, including camera allocation, camera cueing, video tracking, video display, and video encryption. The fusion center acts as the major hub in the ERSA network. Figure 10 shows the interrelationships between the various components of the fusion center.

FAA radar data are fused with Sentinel radar data by the Raytheon Solipsys Corporation Multi Source Correlator Tracker (MSCT) fusion tracker in order to create the threat-conditioned, single integrated air picture. These fusion tracks are then associated with higher-echelon (TADIL-J) data by the MSCT. The fusion tracks are then associated with the database, at which point threat conditioning is performed.

Other, off-board decision support algorithms associate ancillary data with the fusion tracks. Estimates of altitude for non-Sentinel radar-only tracks are produced and associated with the fusion tracks. The sidebar "Altitude Inference Using Multilateration and Multiangulation" shows how ERSA can cost-effectively leverage existing surveillance information. Mode S data is received with an omnidirectional antenna and correlated with tracks of known transponder-equipped aircraft.

In Figure 10, the sensor manager is responsible for radar data preprocessing, i.e., preparing the radar data for delivery to the MSCT. The static databases are parsed for later use in the track classification component, which then associates the data with active fusion tracks.

Track classification performs both association between fusion tracks and reserved beacon codes, and threat-level estimation. From a graphical point of view, threats are color coded (yellow, orange, red), indicating an increasing level of threat. The reason for the assigned threat level is also passed to the ERSA air picture for every track. Threat processing occurs on every track for every track update—with the MSCT, each track is updated at most once per second. These track data are processed at the central fusion center; the outcome of threat classification is sent to remote user displays via a custom protocol.

Camera allocation and cueing are performed at the fusion center. Camera allocation is conducted on the basis of static user priorities. The highest-priority user may use any camera at any time, in effect stealing any camera from lower-priority users. Camera allocation also uses symbolic names for camera sites—the Domain Name Service is used to map between symbolic names and Internet Protocol (IP) addresses. Network Time Protocol is used to synchronize all computers to a common time.

Cueing sends a subset of the fusion track information to each camera site (sensor node), so that when a user requests that a camera be pointed at a track (slaved,

ALTITUDE INFERENCE USING MULTILATERATION AND MULTIANGULATION

To control costs, it was important that the system architecture include existing radars whenever possible. Radars used for ATC employ cooperative (beacon radar) surveillance techniques that elicit aircraft transponder replies to ground-based interrogators. These replies include altitude and identification codes used for ATC. Consequently, the existing ATC radars, with the exception of the ARSR-4 long-range radars, do not measure the altitude of aircraft by using the primary radar. Existing ATC primary radars are track-while-scan radars that resolve aircraft positions in range and azimuth. An innovative technique was developed to determine aircraft altitude by using 2-D radars.

Altitude estimates of potential threats are a critical component of intent assessment. The FAA relies on the automatic self-reporting of barometric altitude from Mode C transponder-equipped aircraft. Mode C is not required in all classes of airspace, and this equipment sometimes fails. Furthermore, a person at the controls of an aircraft can also simply turn off the transponder. Sentinel radars, three-dimensional surveillance and track-

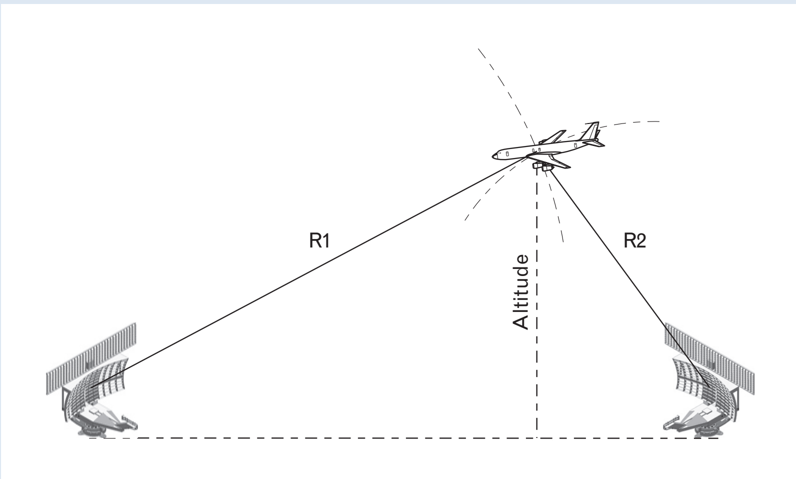

FIGURE A. Example of altitude estimation. Two separate radars can determine an aircraft's altitude by using range measurements.

ing radars, were deployed in the NCR to overcome the 2-D limitation of the FAA primary radar infrastructure. The deployment of these radars largely solves the altitude problem for noncooperative aircraft in the NCR.

ERSA also developed an algorithm that exploits the fact that range measurements almost always have nonzero projections onto target altitude. The basic concept is illustrated in Figure A. For simplicity, both radars and the target are coplanar. Since the distance between the two radars is known, if R1 and R2 are measured, the altitude can be estimated by using basic geometry. The closer the target is to a particular radar, the more its range estimate projects into altitude, making altitude estimates less sensitive to range measurement error.

Conceptually, extending this concept to three dimensions is not difficult. Measurements of range and azimuth define an arc in three space. In the absence of measurement errors, multiple radars give multiple arcs that intersect at the target position. Because the radars are typically coplanar, there is an easily eliminated, but equally valid, intersection below the surface of the earth.

When measurement error is present, the estimated intersection moves, or in the case of more than two radars, there are multiple valid intersections. When the target is far from all radars, the geometry for altitude estimation is poor because the radar ranges do not project significantly in the local up-position direction. Therefore, small range errors will result in large altitude-estimation errors. However, if the target is close to one radar, its range will project significantly into local up- and altitude-estimation error will be reduced.

Performance Prediction

We first investigate predicted algorithm performance for a two-radar geometry. The target of interest in this case is assumed to be flying at 10,000 feet and the two radars are separated by 30 nautical miles. The standard deviation of the up-position estimate is calculated for a grid of points surrounding the radars. The radars are assumed to be ASR-9 radars with range- and azimuth-measurement-error standard deviations of 150 feet and 0.1 degrees, respectively. Figure B shows the results of the analysis, including that, as expected, the algorithm performance is maximized directly over the radars. The figure shows that if a target at 10,000 feet is within about 10 nautical miles of either radar, altitude-estimation error standard deviation is smaller than 2000 feet.

It should be noted that the above analysis assumes that the measurements from the two radars are synchronous. This is rarely the case, but the effect of this deficiency is small because aircraft movement is highly predictable over short time intervals.

Performance Example

We implemented a Kalman filter around the measurement model described in the Appendix and used the filter to demonstrate altitude estimation performance

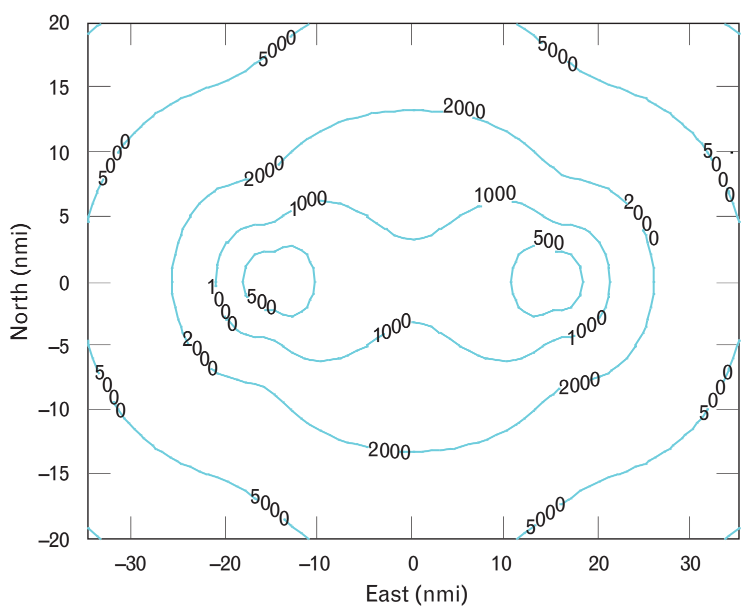

FIGURE B. Estimated algorithm performance, with standard deviation of altitude-estimation error in feet. Target altitude is 10,000 feet. Two ASR-9 primary radars are located at (–15, 0) and (15, 0).

with real radar data. The details of the dynamics model are omitted. We chose a target of opportunity that took off from Washington National (DCA) and flew west while climbing to, and leveling off at, 11,000 feet. This aircraft was Mode C equipped, so self-reported barometric altitude is available for comparison with the AIMM estimates. Only the range and azimuth measurements are used as inputs to the Kalman filter. The Mode C altitude reports are used only for comparison. This aircraft was tracked by six ASR-9/Mode S radars: DCA, IAD, Baltimore (BAL), ADW, Martinsburg (MRB), and Charlottesville (CHO). The aircraft was also tracked by an ARSR-3 radar, The Plains (QPL), but these measurements are not used.

The estimated algorithm performance is also calculated. The analysis of the previous section is repeated by using the trajectory of the aircraft. The range- and azimuth-measurement-error standard deviations are kept at 150 feet and 0.1 degrees, respectively. These parameters are appropriate for ASR-9 radar-only contacts. The radar contacts for this example are actually Mode S beacon reports, which have substantially better error statistics. The analysis is done by using the ASR-9 statistics for reasons that will be made clear by Figure C.

In Figure C, the black trace represents the AIMM estimates of altitude when unmodified radar measurements are used. The green traces are the Mode C altitude plus and minus one-half the standard deviation of altitude error predicted for the algorithm. The AIMM estimates fall outside the one-half-sigma bound most of the time. The one-half-sigma bound is also arguably larger than it should be, since ASR-9 measurement error statistics are used to create it. The reason for this discrepancy is that the error prediction analysis accounts for only random measurement error, not unknown biases. Possible biases include radar position bias, azimuth bias, range bias and, in the case of beacon reports, transponder turnaround bias.

The bias in the AIMM estimate is very likely due to a combination of the measurement biases listed above. The maroon trace in Figure C represents the AIMM estimates when 100 feet is subtracted from every range measurement, simulating the removal of a 100 ft transponder turnaround bias. For most of the track, this measurement bias dominates the other biases to the altitude measurement by more than an order of magnitude. The extreme sensitivity of the algorithm to measurement biases is due to the generally poor geometry. Despite the sensitivity to biases and measurement noise, AIMM does a remarkable job of predicting aircraft altitude with 2-D radar measurements, and provides a usable altitude estimate to the ERSA system.

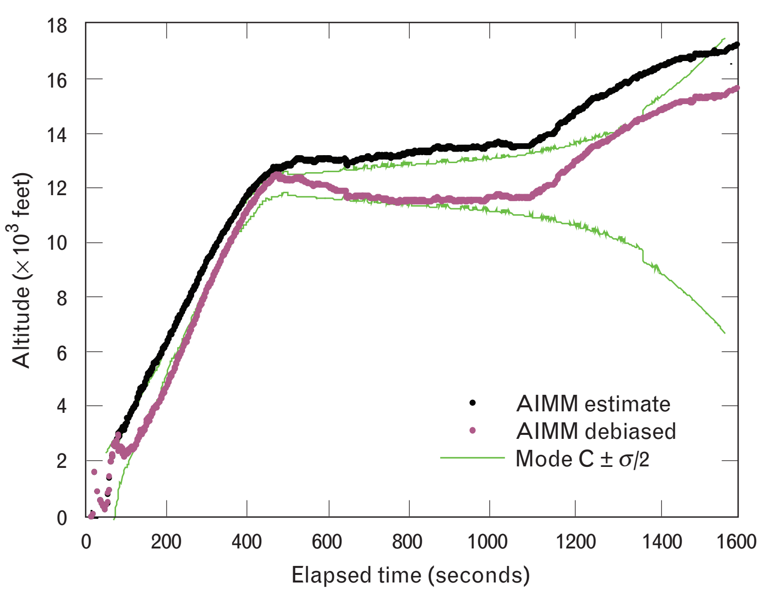

FIGURE C. Estimated and actual algorithm performance for aircraft taking off from DCA airport. The black trace is an AIMM estimate obtained by using unmodified radar measurements. The maroon trace is an AIMM estimate when 100 feet is subtracted from all range measurements. The standard deviation of estimation error is calculated for this particular aircraft trajectory by using the technique described above.

in ERSA terminology), the pointing information will be available. The camera allocation process contains many components of a service-oriented architecture. Each camera-related process registers itself with a centralized directory service. The user processes discover the capabilities of the camera sites and automatically configure the control interface to match the device capabilities for each camera. Some users have the capability to use the VWS; this capability is also configurable on a user-site-by-user-site basis. Users may also be configured for Watch-only. In this mode, steering or slaving a camera is not allowed; however, the video stream is sent to the user station for passive viewing of a camera's output.

Automatic monitoring and control of the real-time system is performed from the fusion center, as well as local recording of FAA radar data, Sentinel radar data, MSCT fusion tracks, the output of altitude estimation, the Mode S identification system processing, and the output of threat conditioning. Video recording is provided in one fusion center on 24 terabyte COTS storage devices.

Sentinel Subsystems

The Sentinel radar is normally deployed as part of an Army Division with two local operators and four maintainers. A major challenge to the ERSA project was making modifications to the radar, and adding support systems to allow the radar to run unmanned, 24/7. Automated fault recovery systems were installed to keep the radar running without human intervention. Radar-track data and status are available remotely to monitor the health of the radars through a network interface.

ERSA Camera Sites

An ERSA camera site is composed of two rack-mounted Dell 2600 dual-central-processing-unit computers. These computers control the camera devices and can dynamically configure their local site to send video from any camera on that site to a user display station. The sensor-node computer contains a frame-grabber card that grabs 30 frames per second of NTSC video. This video can be used to perform software video tracking on video blobs within the field of view of the camera. The sensor-node computer is responsible for controlling all of the cameras at its site. A video node computer configures the equipment at the camera site to send video through the fusion center to the user display.

Follow-On Development

In addition to the usability improvements discussed earlier and inserted through periodic upgrading of the various tracking and assessment algorithms, there are a number of development areas that could significantly extend the capabilities of ERSA.

All-Weather ID

Though the surveillance and track functions of ERSA are all-weather, the visual ID and visual warning components require favorable visibility conditions for their most effective use. All-weather ID could be enhanced by the development and deployment of RF ID sensors. These sensors employ a fine enough range resolution for developing a *range profile* to discriminate target class (large, medium, small) and reveal other features that can be derived from echo strength versus position. Similarly, they employ precision Doppler frequency measurements to discriminate between propeller, jet, and helicopter aircraft, and bird flock signatures. The output of the RF ID system would be integrated into the threat-assessment logic.

Work on radar configurations, waveforms, and processing algorithms is needed to move towards the goal of a low-cost portable RF ID capability. In a positive step, special adjunct ID modes are soon to be made available in upgraded Sentinel radars developed by Raytheon and the Army's Sentinel Project Office.

Deployable Air Defense Components

ERSA was developed specifically to provide situation awareness and intent assessment to support air defense in the NCR. The development of rapidly deployable radar, camera, warning, communications, and fusion-center components would simplify the creation of ERSA-like capabilities in other locales. It could allow for quickly setting up special event coverage for such things as the Olympics, and other high-profile events. Availability of deployable components could also be used to extend the boundaries of the NCR system.

Extension to Other Sites with Less Restricted Airspace

The development of the ERSA decision support logic took advantage of the extended airspace restrictions in effect in the NCR. The success of ERSA has created interest in developing similar decision support systems for use in other regions. In general, however, it will not be feasible to impose extended airspace restrictions in other areas because of their impact on civil aviation. Hence the ERSA intent-assessment tools would need to be modified to operate effectively in less restrictive airspace.

Most of the individual components associated with ERSA—e.g., radars, cameras, and displays—are not dependent on airspace restrictions and could be deployed just as easily in other (on-shore) regions. However, portions of the intent-assessment logic developed for ERSA do take advantage of the NCR airspace restrictions. Many of the alerts generated by ERSA are due to aircraft that violate or are about to violate the rules governing air traffic in that area. Since the airspace restrictions extend 30 or more nautical miles from the center of Washington, alerts can be issued in a timely manner while still maintaining a low false-alarm rate.

Despite the airspace restrictions, there is still a high volume of air traffic in the NCR. In regions with a lower traffic density it is reasonable to expect that ERSA-like intent-assessment logic could be developed without the need for extended airspace restrictions. This expectation is because the false-alert rate depends not only on how well the air traffic is regulated but also on the density of radar tracks. A lower track rate results in fewer opportunities for generating false alerts.

For regions with minimal airspace restrictions, it may be best to replace the ERSA decision-tree-based alerting approach for beacon tracks with the threat-scoring approach currently used for radar-only tracks. The latter approach for radar-only tracks was necessitated because most of the radar-only tracks (e.g., birds) do not abide by airspace restrictions. It is expected that a similar approach could be used for beacon tracks.

Self-Adapting to Decision Support/Assessment Rules

Although the intent-assessment logic developed for ERSA was tailored for the NCR, the development was largely data driven. Alerting was designed around the airspace restrictions in effect and sources of false alerts were examined in detail to determine how they could be prevented while still providing adequate alerting for a true threat. If an ERSA-like system were desired in a large number of regions, it would be best if the intent-assessment logic could be designed to be largely self-adapting.

Developing self-adapting intent-assessment logic should be possible, especially if a threat-scoring approach is applied to all tracks. With such an approach, adaptive alerting thresholds could be employed to limit the number of false alerts (similar to the constant-false-alarm-rate technique used by radar surveillance systems). In addition, self-learning techniques could be implemented to identify and reduce common sources of false alerts. For example, if there is an airport near the protected zone, ground tracks associated with arrivals and departures to that airport may appear threatening. Although other methods would be required to prevent a terrorist from obtaining permission to land at such an airport, logic to self-learn standard approaches and departures could be designed and used to inhibit alerting for aircraft following standard approach and departure routes.

Summary

A rapid development effort undertaken by Lincoln Laboratory for the Air Force Rapid Capabilities Office led to upgrades to NORAD's Integrated Air Defense of the National Capital Region. The Enhanced Regional Situation Awareness system was built upon NORAD's existing components and added additional precision tracking radars, improved radar tracking, cameras for aircraft ID, and a warning system for airspace violators, all integrated with extensive decision support underpinnings. The decision support components were designed to help alert and allow the operators to quickly absorb the available information from the complex airspace surrounding the NCR, to support making the time-critical decisions necessary for counter-terrorism defense.

REFERENCES

1. M.E. Weber, J.E. Evans, W.R. Moser, and O.J. Newell, "Air Traffic Management Decision Support during Convective Weather," in this issue.
2. At https://pilotweb.nas.faa.gov/distribution/home.html, key in FDC for accountability and the NOTAM number 5/4122, and click View NOTAMs.
3. M.R. Endsley, "Situation Awareness Global Assessment Technique (SAGAT)," *Proc. IEEE National Aerospace and Electronics Conf. (NAECON), Dayton, Ohio, 23–27 May 1988*, pp. 789–795.
4. Java Look and Feel Design Guidelines, Version 2.0, Sun Microsystems, Palo Alto, Calif., Feb. 2001, 94303 http://java.sun.com/products/jlf/ed2/book.
5. V. Ahlstrom and K. Longo, "Human Factors Design Standard (HF-STD-001)," Atlantic City International Airport; Federal Aviation Administration William J. Hughes Technical Center, 2003, updated 20 June 2006, http://acb220.tc.faa.gov/hfds.
6. J. Nielsen, *Usability Engineering* (Academic Press, San Diego, 1993; Morgan Kaufmann, San Francisco, 1994).

APPENDIX: A MEASUREMENT MODEL FOR ALTITUDE INFERENCE

To formalize the discussion in the AIMM sidebar, we define a linearized radar measurement model. The goal of the measurement model is to relate a small change in the target's three-dimensional coordinates to a small change in range and azimuth measurements.

Local curvilinear radar coordinates are easily related to east-north-up (ENU) coordinates. The ENU coordinates $\mathbf{x}_L^t$ of the target can be computed by using

$$\mathbf{x}_L^t = \mathbf{R}_L(\mathbf{x}^t - \mathbf{x}_r),$$

where $\mathbf{x}^t$ and $\mathbf{x}_r$ are the positions of the target and the radar, respectively, in earth-centered-earth-fixed (ECEF) coordinates, and $\mathbf{R}_L$ is the matrix that rotates ECEF coordinates into radar-centric ENU coordinates. The radar range measurement is

$$\rho_r^t = \left\| \mathbf{x}^t - \mathbf{x}_r \right\|, \tag{1}$$

and the radar azimuth measurement is

$$\theta_r^t = \tan^{-1}\left(\frac{E_r^t}{N_r^t}\right), \tag{2}$$

where E_r^t and N_r^t are the east and north components of $\mathbf{x}_L^t$. Using a truncated Taylor series expansion, we now linearize Equations 1 and 2 around an initial guess of target position $\mathbf{x}^0$:

$$\mathbf{x}^t = \mathbf{x}^0 + \delta\mathbf{x}$$

$$\rho_r^t\left(\mathbf{x}^0 + \delta\mathbf{x}\right) \approx \rho_r^t\left(\mathbf{x}^0\right) + \left(\nabla\rho_r^t\left(\mathbf{x}^0\right)\right)^{\mathrm{T}}\delta\mathbf{x}$$

$$\theta_r^t\left(\mathbf{x}^0 + \delta\mathbf{x}\right) \approx \theta_r^t\left(\mathbf{x}^0\right) + \left(\nabla\theta_r^t\left(\mathbf{x}^0\right)\right)^{\mathrm{T}}\delta\mathbf{x},$$

where $\delta\mathbf{x}$ is the incremental change in $\mathbf{x}$.

The above linearized approximation is used to produce the measurement model,

$$\delta\rho_r^t = \rho_r^t - \rho_r^t\left(\mathbf{x}^0\right) = \left(\nabla\rho_r^t\left(\mathbf{x}^0\right)\right)^{\mathrm{T}}\delta\mathbf{x} + \varepsilon_\rho$$

$$\delta\theta_r^t = \theta_r^t - \theta_r^t\left(\mathbf{x}^0\right) = \left(\nabla\theta_r^t\left(\mathbf{x}^0\right)\right)^{\mathrm{T}}\delta\mathbf{x} + \varepsilon_\theta,$$

where ε_ρ and ε_θ are the range and azimuth measurement errors, respectively. Carrying out the gradient operator and some algebra yields the following:

$$\delta\rho_r^t = \frac{\left(\mathbf{x}^0 - \mathbf{x}_r\right)^{\mathrm{T}}}{\left\|\mathbf{x}^0 - \mathbf{x}_r\right\|}\delta\mathbf{x} + \varepsilon_\rho$$

$$\delta\theta_r^t = \frac{\hat{\mathbf{v}}_N^{\mathrm{T}}\left(\mathbf{x}^0 - \mathbf{x}_r\right)\hat{\mathbf{v}}_E^{\mathrm{T}} - \hat{\mathbf{v}}_E^{\mathrm{T}}\left(\mathbf{x}^0 - \mathbf{x}_r\right)\hat{\mathbf{v}}_N^{\mathrm{T}}}{\left(\hat{\mathbf{v}}_N^{\mathrm{T}}\left(\mathbf{x}^0 - \mathbf{x}_r\right)\right)^2 + \left(\hat{\mathbf{v}}_E^{\mathrm{T}}\left(\mathbf{x}^0 - \mathbf{x}_r\right)\right)^2}\delta\mathbf{x} + \varepsilon_\theta,$$

where $\hat{\mathbf{v}}_N$ and $\hat{\mathbf{v}}_E$ are unit vectors in the north and east directions with respect to the radar, expressed in ECEF coordinates. Combining the measurements from multiple radars, we get Equation 3, highlighted on the following page, or more compactly,

$$\delta\mathbf{m}^t = \mathbf{G}\delta\mathbf{x} + \boldsymbol{\varepsilon}_{\mathbf{m}}.$$

The first n rows of Equation 3, $\delta\rho_{(i)}^t$, are the multilateration rows, and the second n rows, $\delta\theta_{(i)}^t$, are the multiangulation rows. The last thing that will be needed is a statistical description of the measurement error. We will assume that the measurement error is zero mean and that all measurements are uncorrelated. Therefore, a diagonal covariance matrix, $\Lambda_{\varepsilon\mathbf{m}}$, completely describes the second-order statistics of the measurement error.

$$\Lambda_{\varepsilon\mathbf{m}} = \begin{bmatrix} \sigma_{\varepsilon\rho(1)}^2 & & & & & \\ & \ddots & & & & \\ & & \sigma_{\varepsilon\rho(n)}^2 & & & \\ & & & \sigma_{\varepsilon\theta(1)}^2 & & \\ & & & & \ddots & \\ & & & & & \sigma_{\varepsilon\theta(n)}^2 \end{bmatrix}$$

Least-Squares Algorithm for Performance Analysis

The above measurement model lends itself well to Kalman filtering or interacting multiple model filtering. However, for basic performance analysis, a least-squares

$$
\begin{bmatrix} \delta\rho^t_{(1)} \\ \vdots \\ \delta\rho^t_{(n)} \\ \delta\theta^t_{(1)} \\ \vdots \\ \delta\theta^t_{(n)} \end{bmatrix} = \begin{bmatrix} \dfrac{\left(\mathbf{x}^0 - \mathbf{x}_{(1)}\right)^{\mathrm{T}}}{\left\|\mathbf{x}^0 - \mathbf{x}_{(1)}\right\|} \\ \vdots \\ \dfrac{\left(\mathbf{x}^0 - \mathbf{x}_{(n)}\right)^{\mathrm{T}}}{\left\|\mathbf{x}^0 - \mathbf{x}_{(n)}\right\|} \\ \dfrac{\hat{\mathbf{v}}^{\mathrm{T}}_{N(1)}\left(\mathbf{x}^0 - \mathbf{x}_{(1)}\right)\hat{\mathbf{v}}^{\mathrm{T}}_{E(1)} - \hat{\mathbf{v}}^{\mathrm{T}}_{E(1)}\left(\mathbf{x}^0 - \mathbf{x}_{(1)}\right)\hat{\mathbf{v}}^{\mathrm{T}}_{N(1)}}{\left(\hat{\mathbf{v}}^{\mathrm{T}}_{N(1)}\left(\mathbf{x}^0 - \mathbf{x}_{(1)}\right)\right)^2 + \left(\hat{\mathbf{v}}^{\mathrm{T}}_{E(1)}\left(\mathbf{x}^0 - \mathbf{x}_{(1)}\right)\right)^2} \\ \vdots \\ \dfrac{\hat{\mathbf{v}}^{\mathrm{T}}_{N(n)}\left(\mathbf{x}^0 - \mathbf{x}_{(n)}\right)\hat{\mathbf{v}}^{\mathrm{T}}_{E(n)} - \hat{\mathbf{v}}^{\mathrm{T}}_{E(n)}\left(\mathbf{x}^0 - \mathbf{x}_{(n)}\right)\hat{\mathbf{v}}^{\mathrm{T}}_{N(n)}}{\left(\hat{\mathbf{v}}^{\mathrm{T}}_{N(n)}\left(\mathbf{x}^0 - \mathbf{x}_{(n)}\right)\right)^2 + \left(\hat{\mathbf{v}}^{\mathrm{T}}_{E(n)}\left(\mathbf{x}^0 - \mathbf{x}_{(n)}\right)\right)^2} \end{bmatrix} \delta\mathbf{x} + \begin{bmatrix} \varepsilon_{\rho(1)} \\ \vdots \\ \varepsilon_{\rho(n)} \\ \varepsilon_{\theta(1)} \\ \vdots \\ \varepsilon_{\theta(n)} \end{bmatrix} \qquad (3)
$$

algorithm will do. When the number of radars observing a target is two or more, the three-dimensional position of the target can be estimated iteratively. The general procedure is to guess the target's position, solve Equation 3 to estimate $\delta\mathbf{x}$ by $\delta\hat{\mathbf{x}}$, update the initial guess, and repeat until $\|\delta\hat{\mathbf{x}}\|$ is suitably small.

The measurements from one radar can be used to produce the initial target position. We will have to assume that the local up-coordinate is some reasonable value, say zero. Or if the target is being tracked, the previous estimate of the up-coordinate can be used.

If the number of radars observing the target is two or more, Equation 3 is overdetermined. We will use the least-squares solution

$$\delta\hat{\mathbf{x}} = (\mathbf{G}^T \Lambda^{-1}_{\varepsilon\mathbf{m}} \mathbf{G})^{-1} \mathbf{G}^T \Lambda^{-1}_{\varepsilon\mathbf{m}} \delta\mathbf{m}^t \ ,$$

where $\Lambda^{-1}_{\varepsilon\mathbf{m}}$ is the inverse of the error covariance matrix.

Then we update our initial target position

$$\mathbf{x}^0 + \delta\hat{\mathbf{x}} \to \mathbf{x}^0$$

and start over. The process is repeated until $\|\delta\hat{\mathbf{x}}\|$ converges to a suitably small value. The covariance matrix for the resulting positioning error is

$$\Lambda_{\varepsilon\mathbf{x}} = (\mathbf{G}^T \Lambda^{-1}_{\varepsilon\mathbf{m}} \mathbf{G})^{-1} \ .$$

The covariance matrix can be transformed to local coordinates by using

$$\Lambda_{\varepsilon\mathbf{x}_L} = \mathbf{R}_L \Lambda_{\varepsilon\mathbf{x}} \mathbf{R}_L^T \ .$$

The lower right element in $\Lambda_{\varepsilon\mathbf{x}_L}$ is the variance of the up-position estimate. The standard deviation of the up estimate is of primary interest in helping understand how well this algorithm will work in various geometries.

CURTIS W. DAVIS III
is the assistant head of the Sensor Systems division. He was manager of the Enhanced Regional Situational Awareness rapid development program from the initial definition and development phases through the 2005 Presidential Inauguration deliveries. He received his B.S., M.S., and Ph.D. degrees from The Ohio State University, where he worked on impulse radar for underground probing, building ultrawideband measurement systems, and mathematically modeling radar performance. Since joining Lincoln Laboratory in 1979, he has been involved in analysis and measurement of air defense systems performance, emphasizing instrumentation and experimentation in RF propagation, surveillance and fire control radars, missile seekers, and flight-testing campaigns. He led the design and development of a highly instrumented aircraft for captive-carry seeker testing to investigate seeker and electronic countermeasures performance. His current activities include the design and development of distributed, integrated sensor architectures in net-centric systems.

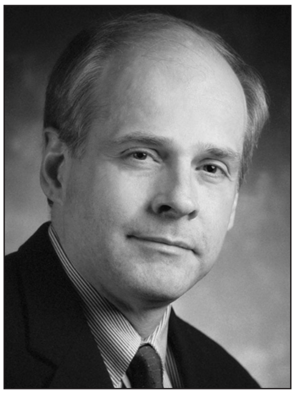

JAMES M. FLAVIN
is the leader of the Surveillance Systems group, which performs research and development to improve the security, safety, and efficiency of the U.S. Airspace System. He came to Lincoln Laboratory in 1986 from AT&T Bell Laboratories, where he worked on the design of adaptive Echo Canceller systems. From 1986 to 1997 he worked as a staff member on the Airport Surface Traffic Automation, Terminal ATC Automation, Weather System Processor, and Kwajalein Modernization and Remoting programs. From 1998 to 2001 he served a two-year tour at the U.S. Army base in the Kwajalein Atoll, first as the radar modernization integration leader and then as the Radar section leader. He received B.S. and M.S. degrees in electrical engineering from the University of Michigan at Ann Arbor.

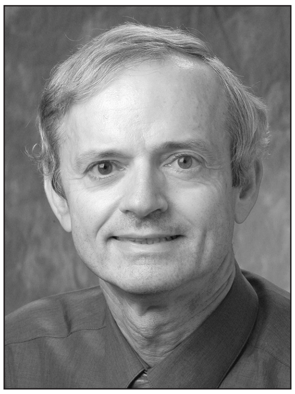

ROBERT E. BOISVERT
is a staff member in the Advanced System Concepts group. Since joining Lincoln Laboratory in 1982, he has provided systems analysis support on a variety of topics, including air defense, ballistic missile defense, civil ATC, and battlefield surveillance. He was also assigned to the U.S. Army base on Kwajalein Atoll, where he developed target discrimination and tracking algorithms, led an effort to upgrade a satellite imaging system, and served as mission director. His current research focuses on the development of intent-assessment and decision support tools for homeland air defense. Prior to joining the Laboratory, he taught for the Georgia Institute of Technology, and he has also taught for the University of Maryland—Asian Division. He received a B.S. degree in mathematics from Emory University, and M.S. and Ph.D. degrees in applied mathematics from the Georgia Institute of Technology. He is a member of Phi Beta Kappa.

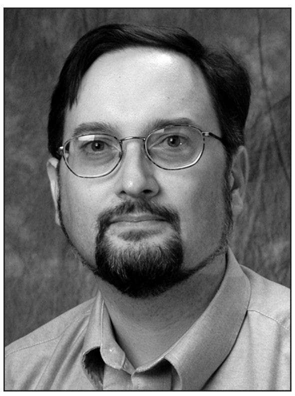

KYLE D. COCHRAN
is a staff member in the Surveillance Systems group. His current areas of research are in real-time systems, service-oriented architectures, air surveillance, and video tracking. Prior to joining Lincoln Laboratory in 1988, he worked at Itek Optical Systems. He received a B.S. degree in computer science/mathematics, an M.S. degree in applied mathematics, and an M.S. degree in computer science/numerical analysis from the University of Illinois.

KEVIN P. COHEN
is the leader of the Tactical Defense Systems group. He joined Lincoln Laboratory in 1995 and worked to develop, flight-test, and analyze electronic counter-countermeasure algorithms. From 1999 to 2003 he was at the Kwajalein Missile Range, where he was the lead of the Mission Support and Data Products section, project lead for the Automatic Target Identification System (ATIDS), and ALCOR test director during the integration of the Kwajalein Modernization and Remoting (KMAR) program. From 2003 to 2006 he worked on the Enhanced Regional Situational Awareness program, initially as lead for sensor data fusion and analysis, and later as the Lincoln Laboratory program manager. He holds a B.S. degree from Boston University in biomedical engineering, and M.S. and Ph.D. degrees from the University of Wisconsin at Madison in electrical and computer engineering.

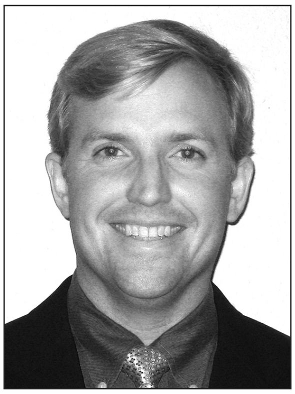

TIMOTHY D. HALL
is a staff member at the Ronald Reagan Ballistic Missile Defense Test Site on Kwajalein Atoll. He is currently working on algorithms to analyze the performance of range safety systems. Before joining the team at the Kwajalein field site, he worked in the Surveillance Systems group developing radar fusion tracking algorithms. Prior to joining the Lincoln Laboratory full time in 2002, he worked on hardware and algorithm development for radiolocation systems. He received B.S. and M.S. degrees in electrical engineering from the University of Missouri and a Ph.D. in electrical engineering from MIT.

LOUIS M. HEBERT
is an assistant leader of the Surveillance Systems group. He is the manager for the Enhanced Regional Situational Awareness program and currently leads several other homeland air defense programs. Since joining Lincoln Laboratory in 1985, he has worked on the design, implementation, operation, and analysis of ground-based and airborne radar systems, including a role as the lead engineer on board over 300 flights of Lincoln Laboratory test aircraft. He received both a B.S. degree and an M.S. degree in electrical engineering in 1983 from Case Western Reserve University, with specialties in electromagnetic theory and communications. He is a member of the IEEE.

ANN-MARIE T. LIND
is a staff member in the Surveillance Systems group and leads the Human Factors team. During seventeen years at the Laboratory, she has designed and evaluated command and control systems for civil and military use. In applying Human Factors principles to these systems, she helps to develop intuitive, user-friendly systems that facilitate situation awareness and decision making. She has supported programs within multiple Laboratory divisions—Air and Missile Defense, Tactical System Technology, Communications and Information Technology, Aerospace, and Sensor Systems. Before joining the Laboratory, she conducted human performance studies at Alphatech Inc. and UNISYS. Her degrees are in psychology and education, including an M.A. from the College of Saint Rose (Albany, New York) and an Ed.D. from Boston University. She is a member of the Human Factors Society and the American Psychological Association.