
Demonstration of Runway Status Lights at Logan Airport

James R. Eggert

■ Lincoln Laboratory has developed a prototype runway-status light system (RSLs), designed to prevent runway incursions and accidents. These status lights will tell aircraft pilots and surface-vehicle operators when runways are unsafe to enter or unsafe for departure. This status information will improve the situational awareness of pilots and vehicle operators, thereby reducing the number of runway incursions and accidents. The goal of the RSLs Logan Demonstration is to use automatic processing of surface primary and approach secondary radar data to drive simulated runway-status lights in a real-time but off-line surface-traffic automation system. This article presents a description of the design motivation, methodology, and implementation for the RSLs Logan Demonstration; it also provides an overview of the entire system on a functional block scale and gives introductory descriptions of the various subsystems.

IN THE LAST TWO DECADES, at least nine major airport surface-conflict accidents have occurred in the United States. In a thirteen-month period in 1990 and 1991, three airport runway-conflict accidents together resulted in the loss of forty-three lives. Runway conflicts start with runway incursions, and approximately two hundred runway incursions are reported every year at U.S. airports. With domestic air traffic predicted to increase by 3% annually over the next decade, the airport surface is expected to become increasingly crowded. The runway-incursion problem must be addressed if future runway-conflict accidents are to be prevented.

A first step toward addressing surface-traffic control is better surface surveillance. Indeed, the FAA is currently engaged in the procurement of about forty ASDE-3 (Airport Surface Detection Equipment) radars, which will provide improved surface surveillance with higher resolution, reduced clutter, and better performance in rain than the older ASDE-2 radars now in operational use. Significant additional performance enhancements are also possible by using appropriate processing technology.

Better surveillance by itself, however, will not address the whole problem. Controllers typically do not

rely on surface radar in times of good visibility. Nor does surface radar by itself offer conflict identification or prevention information to the controller. Also, the surveillance information is not directly available to the aircraft pilots and vehicle operators on the airport surface. This lack of direct access to information is especially important in time-critical conflicts in which the information must be immediately available to both pilots and controllers. What is needed is an automation system that provides conflict detection and aircraft identification aids to the controllers and runway-status information to the pilots.

Progress toward the design of such automation and the prevention of surface accidents requires an understanding of the basic nature of the surface-traffic problem. This problem can be placed into three increasingly encompassing classes—accidents, high-hazard incidents, and runway incursions. Accidents, though great in consequence, are relatively few in number, so an analysis of airport surface accidents benefits greatly by the statistical inclusion of additional incidents that were dangerous but did not result in accidents. High-hazard incidents are those in which at least one aircraft was moving at high speed, and where the minimum separation was a hundred

feet or less. Runway incursions represent a larger class of events than accidents and high-hazard incidents. A runway incursion is defined as any occurrence at an airport involving an aircraft, vehicle, person, or object on the ground that creates a collision hazard or results in loss of separation with an aircraft taking off, intending to take off, landing, or intending to land. Clearly, preventing runway incursions is an effective way to prevent a certain set of airport surface accidents, and a good airport surface-traffic safety system must be effective at reducing runway incursions.

Many of the fundamental concepts of such a surface-traffic automation system have been discussed previously [1]. This article concentrates on the realization of a real-time but off-line surface-traffic automation system at Logan International Airport in Bos-

ton. A much more detailed description of this system and its components is given in a separate report [2].

System Design

A complete airport surface-traffic safety system should include three products that together can address all the major airport surface-conflict scenarios. These three products are runway-status lights, controller alerts, and enhanced controller displays. A runway-status light system (RSLs) will provide current runway-status information to pilots and vehicle operators, indicating when the runway is unsafe to enter or unsafe for takeoff (Figure 1). The information provided by these lights will prevent many runway incursions before they happen. Controller alerts will be used to direct controllers' attention to existing




FIGURE 1. Runway-status light system (RSLs) concept. The runway-status lights indicate to aircraft pilots and surface-vehicle operators when the runway is unsafe to enter or unsafe for takeoff. The system is operated automatically, based on surveillance provided by an Airport Surface Detection Equipment (ASDE) surface radar, an Airport Surveillance Radar (ASR-9) approach radar, and future surveillance systems such as GPS-Squitter or Mode S multilateration.

conflicts between aircraft on or near the runways. Because runway-status lights do not address some of the top accident and incident scenarios, and because controller alerts do not always provide sufficient time for controllers and pilots to correct a situation once it has developed, only a combination of runway-status lights and controller alerts will address all the most common scenarios. (The FAA has contracted for the development of an operational controller-alerting system known as the Airport Movement Area Safety System, or AMASS.) Enhanced ASDE controller displays will present symbology to describe aircraft position, size, altitude, flight number, equipment type, and direction and speed of motion. In addition to airport surface traffic, aircraft on approach to runways will also be depicted on the ASDE displays.

The off-line proof-of-concept RSLs Logan Demonstration incorporates simulated runway-status lights and enhanced controller displays, but does not include a complete controller-alerting system. Runway-status lights provide the greatest part of the protection afforded by the safety system for three important reasons. First, in any time-critical conflict scenario, the most effective safety-system product is one that is directly accessible by the pilots. That direct access is allowed by runway-status lights but not by controller alerts. Second, runway-status lights act to prevent runway incursions before they happen, whereas controller alerts occur only after a conflict has been identified. Third, runway-status lights are effective in a greater fraction of the accident and incident scenarios than are controller alerts. Therefore, for a combination of reasons, including maximizing system effectiveness in the face of developmental schedule constraints and reducing the duplication of research efforts, the RSLs Logan Demonstration does not include controller alerts except for limited demonstration purposes.

Runway-Status Lights

There are two types of runway-status lights: *runway-entrance lights*, which indicate when the runway is unsafe to enter, and *takeoff-hold lights*, which indicate when the runway is unsafe for takeoff. The two types of lights are driven in concert by a single safety logic, and they operate together to prevent runway incur-

sions and accidents. The runway-status lights function fully automatically in response to real-time surveillance. The off-line RSLs Logan Demonstration does not in fact incorporate an actual field-lighting system, but simulates the runway-status lights by the use of an illuminated model board and computer-driven displays.

The runway-status lights have two states: on (red) and off. These lights indicate runway status only; they do not indicate clearance. A green state was specifically avoided to prevent any false impression of clearance. Clearance is to remain the sole responsibility of the air traffic controller, and is not to be provided or implied by the RSLs. An amber state was also avoided because in the case of runway-entrance lights it could tend to be confused with the amber color of the International Civil Aviation Organization (ICAO) standard taxi hold-position (wig-wag) lights. The runway-status lights are designed to be as conspicuous as possible while minimizing the possibility of confusion with other light systems.

Runway-status lights are designed to be generally invisible to pilots of aircraft at high speed. This design decision was made so that red lights, especially lights that suddenly turn red, will not be shown to pilots whose aircraft speed precludes them from making sudden maneuvers. Runway-entrance lights are hooded so as not to be visible to pilots of aircraft on the runway, and they are generally not active at runway-runway intersections. Takeoff-hold lights are also hooded, and they require that an aircraft be in position for takeoff for the lights to be illuminated. The design of the fixtures and light logic thus generally prevents pilots of aircraft at high speed from seeing red runway-status lights.

A proposed fixture for the runway-status lights would be the standard fixture used for ICAO wig-wag lights, with the amber lenses replaced by red lenses and the lamps upgraded to brighter bulbs (Figure 2). These fixtures use redundant light bulbs and other electrical components to minimize the impact of single-component failures on the operation of the system. They are also in current production, allowing off-the-shelf delivery.

Runway-entrance lights will be located at the taxiway entrances to runways and will be positioned on



FIGURE 2. A proposed runway-status light fixture. This fixture is based on the standard International Civil Aviation Organization taxi hold-position (wig-wag) light, with the amber lenses replaced by red lenses and the lamps upgraded to brighter bulbs. An addressable light controller would be mounted in or near the base of this fixture to allow individual control over each lamp.

either side of the taxiway, near the runway edge and well beyond the hold line (Figure 3). Runway-entrance lights will also be located at runway-runway intersections, although they will not always be implemented or actuated there. Runway-entrance lights will be illuminated to indicate to aircraft pilots and surface-vehicle operators that the runway is hot (i.e., it is being used for a high-speed operation like takeoff or landing), and that the runway is currently unsafe to enter at that intersection. Runway-entrance lights will be extinguished when the runway is no longer unsafe to enter at that intersection.

Takeoff-hold lights will be located at takeoff-hold positions and placed on either side of the runway near the runway edge (Figure 4). These lights will indicate to aircraft pilots that the runway is unsafe for takeoff (i.e., the runway is currently occupied or is about to be occupied); they will be extinguished when the runway is safe for takeoff, or if the aircraft in position for takeoff vacates the runway.




FIGURE 3. Runway-entrance lights in operation for an aircraft landing on a runway. The runway-entrance lights in front of the landing aircraft are illuminated red, indicating to the taxiing aircraft that the runway is unsafe to enter. The runway-entrance lights behind the landing aircraft are extinguished, indicating that the runway is safe to enter there.




FIGURE 4. Takeoff-hold lights in operation for an aircraft in position for departure with crossing traffic. The takeoff-hold lights indicate that the runway is currently unsafe for departure because the runway is about to be occupied by a crossing aircraft.

RSLS Enhancements to ASDE Display

The RSLS provides several display enhancements to an ASDE display. An ASDE display without enhancements typically contains radar video with blanking to reduce visible clutter, and line graphics to depict runway and taxiway edges and building outlines. The RSLS provides an iconic depiction of tracked traffic, with symbolic tags for each icon, approach bars for aircraft inside the outer marker, depiction of runway-status-light states, and special markings for aircraft identified as being in conflict. For demonstration and development purposes, additional internal surface-monitor information can also be displayed. The RSLS Logan Demonstration supports both monochrome and color ASDE displays. The off-line RSLS Logan Demonstration does not include radar video on its display. This temporary omission was chosen to reduce development time and equipment expenses, and is not envisioned for a complete RSLS.

Tracks, or indicators of stationary and moving aircraft or other surface traffic, are displayed as icons on the enhanced ASDE display (Figure 5). Each icon represents the position and direction of motion of the track and, for tracks with ASDE image information, is drawn with a size proportional to the area of the ASDE image. Each displayed track has a data tag connected to the icon with a leader line. The ASDE display software selects the leader-line direction to eliminate possible overlapping tags and crossing leader lines. The data tag can be displayed in two formats. The primary tag format shows aircraft altitude in hundreds of feet and track velocity in knots. For example, the data tag 001 122 in Figure 5 indicates an aircraft at 100 feet traveling at 122 knots. The primary tag format also shows aircraft flight code and equipment type when this information is available, the latter alternating with the velocity field. The secondary tag format is meant primarily for system development, and shows internal track numbers, track




FIGURE 5. RSLs enhancements to the ASDE controller display. Arrows indicate position, direction of motion, and size of the radar tracks. Stationary tracks are also optionally depicted by circles. Data tags present the radar track's altitude, velocity, flight number, and equipment type. The approach bar depicts an aircraft on approach to a particular runway; its two endpoints represent approximately five miles of airspace from the outer marker to the runway threshold.

surveillance source or sources, altitude in feet, velocity in knots, and aircraft flight code, when available.

Aircraft on approach to runways and inside the outer marker are displayed on approach bars. The outer marker, which is part of the Instrument Landing System (ILS), is a radio navigational aid located on the runway centerline at the point where an ILS standard approach begins its final descent. An approach bar is a short line segment drawn near the approach end of the runway. It is drawn at a different scale and represents the approximately five-nautical-mile distance from the outer marker to the runway threshold. Aircraft identified as being on approach to a runway are shown as diamonds on the approach bar. When the aircraft is near enough to the runway to appear on the scale of the ASDE display, it disappears from the approach bar and appears as a normally displayed target.

Runway-status-light state information is also rendered to the enhanced ASDE display. It can be drawn in two different symbologies. An illuminated runway-entrance light can be represented by a bar across the intersecting taxiway, and takeoff-hold lights can

be represented as a bar across the runway. Alternatively, runway-entrance lights and takeoff-hold lights can be drawn as acute triangles on either side of the taxiway or runway and oriented to depict the directionality of the actual lights.

If the RSLs safety logic identifies targets as being in conflict, this information can be drawn to the ASDE display. The targets are circled in white and remain highlighted until the conflict is resolved. Additional RSLs internal information can also be shown on the ASDE display. This information includes the target state identification (taxi, stopped, arrival, departure, departure abort, or unknown), the range of predicted target positions produced by the surface monitor, and artificial target (sprite) positions and control information.

The ASDE display enhancements also allow for the possibility that future tower displays could be in color. The color displays show the runway, taxiway, building outlines, and approach bars in green; target icons and tags in yellow; illuminated runway-status lights in red; and conflict alert circles in white. Color in an ASDE display is extremely useful in enhancing

the visibility of the display to controllers, which thus improves the rate and efficiency of information comprehension by controllers.

Controller Alerts

Although the RSLs Logan Demonstration does not supply a complete controller-alerting system, it does have an architecture that supports such an alerting system. To demonstrate this capability we included a single type of alert in the system. The conflict that can be detected is between an arriving or landing aircraft and a stopped target on the arrival runway. When a conflict is detected, the RSLs circles the conflicting targets on the ASDE display and generates a synthesized voice alert. The voice alert gives a warning signal, and then it gives the location and type of the conflict. A complete controller-alerting system would include the capability of detecting perhaps a dozen general conflict types.

RSLs Logan Demonstration Methodology

The main objective of the RSLs Logan Demonstration is to develop a surface-traffic safety system that can prevent most runway incursions and identify impending surface conflicts. This objective required the development of several significant capabilities:

1. An ASDE surface radar to provide radar images with sufficient resolution and scan frequency for tracking surface traffic.
2. A radar interface board to digitize, time-stamp, and limit the radar coverage defined by a downloadable sensor map.
3. A surface-radar processing system to process information from the ASDE radar automatically, performing clutter rejection, target morphological processing, and scan-to-scan association.
4. An interface to the Automated Radar Terminal System (ARTS) computer to provide radar surveillance data for aircraft on approach to the runways.
5. A sensor-fusion process to merge tracks from the ASDE processing system automatically with tracks from the ARTS computer, and perform multipath rejection.
6. A surface monitor to classify and predict aircraft behavior, identify surface conflicts, and drive

runway-status lights and controller alerts.

7. A display system to allow basic evaluation and demonstration of the entire system.
8. A performance-analysis suite to allow a detailed evaluation of the operation of the RSLs.

Figure 6 shows an overview of the system architecture. The analog signal from the ASDE surface radar is digitized and processed in the radar surveillance processing system. Its tracks, along with those derived from Airport Surveillance Radar (ASR-9) radar surveillance using the ARTS tap, are passed on to sensor fusion. The output of sensor fusion is a single set of tracks presenting a coherent view of the airport surface and approach traffic to the surface monitor, which identifies aircraft states, predicts future target positions, determines runway-status-light states, and generates alert commands. The system output is shown on several displays. The various stages of processing are described in more detail below.

Several system requirements resulted in basic engineering design choices. These requirements included the following: (1) December 1992 demonstration, (2) off-line noninterfering demonstration, (3) real-time response to live traffic, (4) minimal system response time, (5) minimal hardware design time, and (6) adequate design flexibility. The RSLs Logan Demonstration was required to be functional in the December 1992 time frame, which precluded the use of the ASDE-3 surface radar at Logan Airport because that radar was not expected to be operational in time. Thus another surface-radar system—a Raytheon Pathfinder X-band marine radar—was installed on the roof of the old control-tower building at Logan Airport for use in the development and demonstration of the system. This radar is called the ASDE-X.

The RSLs Logan Demonstration was required to have no operational impact. Thus there are no actual runway-status lights and no RSLs presence in the control-tower cab, and the RSLs does not interfere with normal FAA or aircraft operations. All demonstration displays and system control screens are located in a demonstration room on the sixteenth floor of the Logan Airport tower, or in other noninterfering areas. All demonstration equipment operates on a noninterfering basis; a failure in any demonstration subsystem cannot result in operational interference.




FIGURE 6. Overview of the RSLS Logan Demonstration architecture. Surveillance provided by a surface radar and a terminal radar is processed separately and then fused to provide aircraft tracks on the airport surface and in the approach space. The surface monitor assesses the traffic picture and drives the runway-status lights. The different displays show the traffic and runway-status light information on a map of the airport and its vicinity.

The real-time nature of the RSLS mandated that the system should have sufficient processing throughput to keep up with peak data loads. For the case of ASDE-X surface-radar processing, this processing requirement demanded the use of a fairly powerful computer. Most of the subsystems operate on separate computer platforms to distribute the computational load and reduce the system impact of a single-point failure.

A real-time surface-traffic safety system must take into account the fact that time-critical situations can occur, making large processing delays intolerable. Several design choices, most notably the order of operation in the ASDE clutter-rejection process, and the design of a dual tap to the ARTS computer, were a result of this consideration.

Because the RSLS Logan Demonstration development overlapped design and implementation, design changes along the way were clearly inevitable. Recogni-

tion of this fact led to the decision that the use of custom hardware would be avoided wherever possible, and much of the system functionality would be performed in software by using commercial off-the-shelf equipment. This decision proved to be of great benefit throughout the system design, and was made possible by the explosion in computer system performance in the past few years. In the case of the ASDE radar interface and certain required improvements to the ASDE-X marine radar, however, custom hardware was required.

RSLS Logan Demonstration Description

The off-line RSLS Logan Demonstration is installed at Boston's Logan International Airport. Figure 7 is a pilot's diagram of Logan Airport showing the runways, taxiways, runway designations, runway dimensions, hold positions, and terminal areas [3]. The demonstration room, which is shown in Figure 8, is

on the sixteenth floor of the Logan Airport control tower, in the Massport conference room. This room provides a clear view of most of the airport's runways and taxiways, allowing good visual verification of the operation of the system. The demonstration room has several displays showing various aspects of the system operation. A Raytheon Pathfinder radar display shows an image of the raw ASDE-X surface-radar surveillance. Two monochrome high-brightness displays (manufactured by Orwin Associates) simulate an enhanced ASDE display and a DBRITE (Digital Bright Radar Indicator Tower Equipment) display. A third high-brightness display uses backlit active-matrix liq-

uid-crystal color technology to demonstrate how a color display could be usable in a high-ambient-light environment.

Figure 9 shows a Logan Airport model board that includes architectural models of the terminal buildings, depictions of the runways and taxiways, and a variety of actively controlled field-lighting systems. The field-lighting systems are simulated by using fiber optics, and they include the RSLS runway-status lights, runway-centerline and edge lights, taxiway-edge lights, approach lights, taxi hold-position lights, and stopbars. These systems are driven actively by an integrated lighting-control system, which is interfaced to the rest of the RSLS Logan Demonstration by using an RS-232 interface. Transition from an off-line demonstration of the runway-status lights using the model board to a real field-lighting system can in principle be performed by unplugging the model board and plugging in the field-lighting controller.

A DECTalk digital voice-synthesizer system generates audible voice alerts in response to the alert commands from safety logic. The DECTalk voice quality is insufficient for a real controller-alerting system, but it is adequate for a demonstration system.

The RSLS Logan Demonstration also has two control displays located in the demonstration room. These are the control displays for the surveillance processing computer and for the sensor-fusion and surface-monitor workstation. The former display can also be used to show real-time radar images either before or after clutter rejection. The latter display functions as an additional color ASDE display (although it is not a high-brightness display), and it is used to generate and control artificial targets.

The other components of the RSLS Logan Demonstration are located outside the demonstration room itself. The ASDE-X radar is located on the roof of the old control-tower building (the building labeled "control tower" in Figure 7), and its associated electronics are located nearby and on the fifteenth floor of the new control tower behind the old control tower. The ARTS interface hardware is located in the ARTS equipment rooms on the sixth and seventh floors of the old control-tower building. The computers used to drive the two high-brightness monochrome displays and to receive the information from




FIGURE 7. Boston Logan International Airport runway and taxiway map. The runways and buildings are shown in black and the taxiways are shown in gray. The taxiway configuration is shown as of 1992; some new taxiway construction and change in nomenclature has occurred since this map was produced.



FIGURE 8. RSLS Logan Demonstration room. The windows offer a sixteenth-floor view of Logan Airport and Boston harbor. The RSLS model board is on the left, while the computer monitors are on the right. The two displays above the model board are high-brightness monochrome displays used to show the traffic in and near the airport.

the ARTS interface are located on the fifteenth floor of the new control tower. Normal operation of the demonstration system includes a startup procedure that takes approximately five minutes. Thereafter, the system is completely functional and normally operates without requiring user input.

Subsystem Descriptions

The RSLS software has three major modules: ASDE surface-radar surveillance processing, sensor fusion, and surface monitor. These three modules are described in more detail here, along with the radar interfaces. Additional modules are used to accomplish the various required ARTS interface, display, playback, and analysis functions. These software modules communicate with one another on the same or different computer platforms by using efficient communications protocols. The system can record all relevant ASDE, ARTS, sensor-fusion, and surface-monitor data simultaneously and in real time. These recorded

data can be played back through parts or all of the RSLS software to review interesting scenarios, evaluate performance, and help refine the various processing algorithms.

Surface-Radar Surveillance Processing

The first task in the development of the RSLS was the creation of a high-quality surface-radar tracking system. This system is, in fact, separately fieldable from the runway-status lights themselves, and represents a major advance in surveillance capability. Because it produces good surface surveillance by using an inexpensive radar and advanced image-processing and tracking techniques, this surface-radar system can be used at airports where a more expensive radar is not justified. The use of this enhanced surface-radar system can make the benefits of surface radar available to more and smaller airports on a cost-effective basis.

To develop the RSLS surface-radar tracking system, we had to overcome several major problems with



FIGURE 9. RLSL Logan Demonstration model board. The model board is an architectural model of Logan Airport with computer-controlled fiber optic lights simulating the runway-status lights; runway approach, threshold, centerline, and edge lights; taxiway centerline and edge lights; and taxi hold-position lights. All of the runway lights are illuminated in this photograph, which is for illustration only and does not represent any real runway configuration.

surface-radar systems—namely, clutter, target splits, shadowing, and multipath. Clutter occurs because the radar transmits energy down toward the airport surface and receives returns from many surface objects in addition to the aircraft and surface vehicles that are of primary interest. Target splits occur because the surface radar has fairly high resolution, and there are portions of an aircraft that reflect essentially no energy back to the radar. Shadowing occurs when one aircraft obscures another aircraft from the viewpoint of the radar. Multipath occurs because the radar signal can bounce off several objects in turn and still return to the radar with enough intensity to be detected, thereby producing phantom outrange targets. These effects make tracking primary surface-target radar returns difficult.

Several techniques were developed to solve these surveillance problems, including clutter rejection,

morphological processing, and merge and scan-to-scan tracking. A dynamic clutter map is used to estimate and remove clutter from the radar images. This clutter map contains the mean and mean square for every pixel in the map and is updated every scan for all clutter pixels. This processing allows the clutter map to accommodate changing conditions such as rain and snow. At Logan Airport the clutter map contains approximately 1.2 million pixels. Morphological, or shape, processing is used to reconnect split targets to avoid multiple-tracking and centroiding errors, and to eliminate small objects that are not target-like in appearance. To decrease computational latency, the surveillance area is split into wedges, and both the clutter rejection and the morphological processing are performed in parallel on these wedges in a multiprocessor computer. Targets output by the morphological processing are pasted together at the

wedge boundaries and tracked from scan to scan by merge processing. In merge processing, tracks that are dropped because of shadowing and other problems are reacquired rapidly (usually within five seconds) by using special reacquisition logic. Multipath is rejected by sensor fusion on the basis of track length and ARTS information. An analysis of the tracking performance of this system, which is presented in a separate report [2], indicates that the probability of tracking an aircraft is approximately 98.6%. The conclusion is that surface traffic can be detected and tracked with high reliability.

Clutter Rejection. The main purpose of clutter rejection is to estimate and eliminate constant or slowly varying clutter from the radar images, detect target pixels that stand out from the clutter, and transfer the target-pixel information in an efficient fashion for later processing. The clutter is estimated by using a linear recursive estimator for the mean $\langle x \rangle$ and mean square $\langle x^2 \rangle$ for each pixel log-intensity measurement x in the surveillance map for each scan i , by using the formulas

$$\langle x \rangle_i = \frac{\tau - 1}{\tau} \langle x \rangle_{i-1} + \frac{1}{\tau} x_i$$

$$\langle x^2 \rangle_i = \frac{\tau - 1}{\tau} \langle x^2 \rangle_{i-1} + \frac{1}{\tau} x_i^2,$$

where τ is the time constant for the two estimators. From the mean and mean square, a threshold t_i is calculated by the equations

$$\sigma_i = \sqrt{\langle x^2 \rangle_i - \langle x \rangle_i^2}$$

$$t_i = \langle x \rangle_i + H(k\sigma_i, l_{k\sigma}, u_{k\sigma}),$$

where the function $H(y, l, u)$ given by

$$H(y, l, u) = \begin{cases} l & y < l \\ y & l \leq y \leq u \\ u & y > u \end{cases}$$

limits the excursion of the thresholds from the mean. Pixels whose log-intensity exceeds the threshold calculated in the previous radar scan are identified as target pixels. Target pixels are not used to update the clutter statistics. Instead, they are grouped together in

radial runs. Runs shorter than three pixels are discarded, as are runs that are not adjacent to any other runs. Discarding these runs produces a radial and an azimuthal prefiltering that does not affect the eventual results, but greatly reduces the amount of information passed on to the morphological processing.

Morphological Processing. Morphological processing is used to coalesce the lists of target-pixel runs produced by the clutter-rejection algorithm into targets representing the outlines of airplanes or surface vehicles. This processing is done in three steps, as shown in Figure 10. First, a morphological opening is performed on the clutter-rejected pixel data, as shown in Figure 10(a). An opening is composed of an erosion followed by a dilation. An erosion has the effect of peeling off one layer of pixels from the outside of every clump of target pixels; this process destroys small clumps and shrinks larger ones. A dilation accretes one layer of pixels onto the outside of every clump of target pixels. (The actual implementation of the erosion and dilation operations does not require that clumps of pixels be identified explicitly.) The net result of the opening is the elimination of salt-and-pepper noise in the detected image and the smoothing of the outlines of the larger images, as shown in Figure 10(b).

The second step in morphological processing is to group the remaining target-pixel runs into connected components. This step is accomplished by using a perimeter-tracking algorithm that steps around the boundary of a component until it returns to its starting point, as shown in Figure 10(c). In this algorithm, every target-pixel run must appear at least once on the boundary of a component. This algorithm fails only for bizarre cases with components inside of components, a pixel run configuration that is essentially never seen in real radar images. The process of grouping runs into connected components produces a representation of all the distinct components visible in the radar image.

The third step in morphological processing is to group components that belong to the same aircraft or surface vehicle, as shown in Figure 10(d). Because the ASDE radar is an imaging radar and because aircraft tend to self-shadow, aircraft images are often broken into completely separate components. A distance cri-




FIGURE 10. Morphological processing example. (a) In the input radar image, the black squares denote clutter-rejected target pixels. (b) Erosion deletes the pixels marked in gray. (c) Subsequent dilation adds the pixels marked in light blue. (d) The two green components are close enough to be grouped as a single object. The blue component is a separate object.

terion is used to identify which components should be grouped together to form one object. The algorithm identifies component pairs that are fairly close and performs a test dilation on them to see if they merge into one component. If they do, then the component pairs before the test dilation are grouped together into a single object. The result of morphological processing is a list of objects detected in the ASDE radar surveillance space, where each object is described by the target-pixel runs grouped into one or more connected components.

Merge and Scan-to-Scan Tracking. A necessary complication in ASDE radar image processing is the azimuthal division of the surveillance region into wedg-

es. This division of the image reduces latency problems and allows the computation to be distributed over several computer processors. An associated complication is that the detected objects on the wedge boundaries must be pasted back together by a single process. This reconstruction of object images is done by a technique called *merge processing*, which carefully merges component segments back together and then groups components into objects correctly across the wedge boundaries. (Another potential difficulty posed at the wedge boundaries by the azimuthal filtering performed by the clutter-rejection process is circumvented by simply not doing azimuthal filtering at the wedge boundaries.)

After merge processing has correctly pasted the objects together across wedge boundaries, the centroid and area are computed for each object. Each object is then compared with tracks of objects computed in previous scans to look for matches. Potential matches are identified by computing a simple two-point projection of the track to the present scan, and accepting targets or objects that are within a certain association radius of the projected position. We must be careful with this process, however, because sometimes more than one target matches a given track, and sometimes more than one track matches a given target. Thus we use a best-available-match algorithm, in which the best match among all target-track pairs is taken first. This matching must be done in real time, even though all the targets may not yet be available for the present scan. The algorithm allows for corrected updates—if a better match is found later, it is used instead, and the previously used match is withdrawn from the track and made available to other tracks.

The tracks are divided into four classes, in order of priority for access to new targets; these four classes are high-confidence tracks, bad-drop tracks, established tracks, and new tracks. High-confidence tracks are those which have passed a lead-in filter, which is a travel-distance requirement used to discriminate between real aircraft or surface-vehicle tracks and those tracks which correspond to false detections or multipath. Bad-drop tracks are former high-confidence tracks that were dropped in regions and at velocities where a track drop is not expected. The algorithm uses special reacquisition logic based on matching target area as a function of range and aspect angle to compare these tracks to targets not matched to high-confidence tracks. Established tracks are those which have not yet passed the lead-in filter. New tracks are those which have been seen only once, and thus have no associated velocity estimate. New tracks are allowed a much larger association radius to allow airborne (and hence quickly moving) aircraft as well as surface traffic to be tracked when they are first acquired. A target that does not match to a track in any of the four groups will start its own new track for the next scan. Thus the result of the scan-to-scan processing is a series of track reports for all the detected objects in the surveillance area.

Sensor Fusion

The tracks from the ASDE radar are fused with tracks from the ARTS system in the sensor-fusion process. These two sets of tracks are paired by sensor fusion to form a combined track containing all the available information on aircraft in the surveillance area. The ASDE tracks are fused with the ARTS tracks by comparing their positions and velocities. If the position and velocity difference of two tracks from different sensors falls within an error ellipse in phase space, then the two tracks are considered to correspond to the same aircraft, and the tracks are fused. The flight number and equipment type are transferred to the fused tracks and maintained even when the ARTS coverage is lost. When the pairing cannot be performed unambiguously, then fusion is not performed. This ambiguity avoidance prevents an ARTS flight number from being applied to the wrong ASDE track.

The sensor-fusion process maintains filtered position, velocity, acceleration, and altitude estimates for all tracks. It also estimates and corrects surveillance clock offsets, and it maintains knowledge of the current barometric pressure for use in correcting the pressure altitudes provided by the ARTS tap. Sensor fusion includes a capability to filter tracks on the basis of position, velocity, altitude, area, track length, track reliability, and surveillance source. This capability is used to suppress multipath and residual clutter, define overlapping radar coverage areas, and reject uninteresting tracks such as boats and overflights. Sensor fusion can also coast tracks to allow for following aircraft through surveillance gaps or glitches. The output of sensor fusion is one coherent picture of the airport surface and approach space, with reliable tracks that include the information required by both the tower controllers and the surface monitor.

Surface Monitor

The fused tracks created in the sensor-fusion process are passed to the surface monitor, which forms an operational view of the airport traffic. The surface monitor first locates the tracks with respect to the network of runways, taxiways, and approach areas at the airport. These areas are defined by bounding polygons,

so a linear search through the list of polygons is used with a point-in-polygon algorithm to identify the correct region. The surface monitor then identifies the present operational state of the tracks, which is one of the following: stopped, taxiing, arriving, landing, departing, landing abort, or departure abort. A state machine with hysteresis in the transitions is used to provide accurate and stable state identifications. Figure 11 illustrates these different track states and their associated transitions.

The surface monitor then projects the future behavior of the tracks. Two projections are made: the first is how far the track must move in a certain time horizon even if it tries to stop, and the second is how far the track might move in the same time horizon if it tries to accelerate. State-dependent assumptions are made for the acceleration and deceleration profiles. The likely future position of the track lies between these two projections. Each projection is allowed to

be multibranching to allow possible turns at every intersection. Impossible turns, in which the track could not make the turn even if it decelerated just for the turn, are not allowed. These projection trees form the basis for the action of the safety logic.

The safety logic determines which runway-status lights or controller alerts need to be illuminated or sounded. The projection trees are used to determine which runway-status lights or abstraction thereof need to be notified of the behavior of a particular track. Once notified, the control logic for that particular light determines the behavior of the light. The projection trees are also used in the demonstration alert logic to identify runway conflicts and sound an audible alert. Using a single surface monitor to generate both light and alert events enables the system to maintain logical consistency for lights and alerts and to avoid contradictory information being sent to the pilots and controllers.




FIGURE 11. Surface-monitor target state diagram. Allowed state transitions are indicated by arrows. The None state is a pseudostate, representing the source and sink for target states. UNK is the unknown state used for an initial indeterminate state or for a don't-care state. ARR is arrival, LDG is landing, TAX is taxiing, STP is stopped, DEP is departure, LBT is landing abort, and DBT is departure abort. The LBT state and the DBT state represent abnormal but not necessarily unsafe aircraft states.

ASDE Radar Interface

A custom radar interface was designed at Lincoln Laboratory to digitize the analog ASDE-X radar signal and send digitized samples of interest to the computer for processing. The radar output is digitized at 42 MHz with an 8-bit A/D converter. The resulting data are subjected to a censoring map that determines which regions are of interest and are to be sent on for further processing, and discards data outside those regions. The censoring map is downloaded to the radar interface at system startup and must be designed separately for each airport. The use of the censoring map at Logan Airport results in a reduction of the data rate to approximately 660 kByte/sec. The censored data can also be recorded on tape for later playback and analysis.

ARTS Tap

Commercial off-the-shelf hardware was purchased to tap the ARTS computer with minimal delay and maximal coverage. The ARTS tap has two parts: a Serial Communications Interface Processor (SCIP) tap that looks at the surveillance input to the ARTS computer, and a Multiple Display Buffer Memory (MDBM) tap that looks at the display information written by the ARTS computer to up to four controller displays. Each part can filter the information to a particular geographical region and type of information desired. The SCIP tap provides position, altitude, and transponder Mode A code for each aircraft in the approach space. The MDBM tap provides position, flight identification, and equipment type for the same aircraft.

Future Improvements

Certain modifications are necessary before the RSLs Logan Demonstration can be turned into an operational field demonstration. First, an actual field-lighting system will need to be installed. This system should include redundant electrical cabling and electrical controllers to maintain high reliability, and a maintenance monitoring facility to shorten down time. Second, a tower-controller interface will need to be implemented. The tower controllers or the controller supervisors will need to input runway configura-

tion information to the system. Third, system processing performance will have to be improved.

The performance of the RSLs Logan Demonstration can be improved by modifying both the system architecture and the various components of the system. Certain system capabilities that might improve reliability were considered out of scope for the research system developed here. These capabilities, which include redundant hardware and software, automatic built-in test procedures, and real-time performance logging, should ultimately be included in a fully operational field system.

Other system-level improvements that should be considered for future incorporation into the RSLs Logan Demonstration concern greater information sharing between the various system components. For example, the ASDE processing component can better initiate new tracks for arriving aircraft if it is given information about these aircraft derived from the ARTS interface. Similarly, sensor fusion can better fuse tracks through surveillance gaps if it is given the arrival runway predictions computed by the surface monitor.

Another way of improving the performance of the RSLs is through improved surveillance. On the system level this improved surveillance would be accomplished by incorporating new surveillance technologies, such as GPS-Squitter, Mode S multilateration, the ASDE-3 radar, or multiple ASDE radars. Surveillance can also be enhanced by improving the performance of surveillance processing, chiefly in the cases of target location in shadows and merged images, and improved tracking.

The RSLs performance can also be enhanced by improving sensor fusion's treatment of ambiguous or conflicting surveillance information. Further improvements can be made in the capability of the surface monitor to estimate the time of future events and to use such estimates to drive lights. A major improvement of the system is also possible by carefully tuning all the available parameters. Some of this system tuning has already been done, although sometimes the parameters used are a compromise between correct results and processing time, and sometimes effective tuning was impossible because of the lack of adequate assessment tools.

Summary

The off-line proof-of-concept RSLS Logan Demonstration showed that the system can detect and track aircraft and surface vehicles on an airport surface by using a primary radar, combine surface primary and approach secondary radar information into one view of the airport and its environs, determine what each aircraft or surface vehicle is doing, predict the possible future positions of each track, and use those predictions to drive runway-status lights. The logical continuation of the development of the RSLS should be to incorporate the discussed design improvements, test the system performance over a wide variety of traffic and weather conditions, and install a set of runway-status lights on the field for an operational suitability test.

Acknowledgments

The RSLS Demonstration at Logan International Airport in Boston was developed by Lincoln Laboratory under an interagency agreement between the Federal Aviation Administration and the U.S. Air Force. The RSLS program was previously called Airport Surface Traffic Automation with Runway Status Lights (ASTA-1). The development of the RSLS Logan Demonstration required the efforts of many dedicated individuals at Lincoln Laboratory, including Walter Brown, Richard Bush, Steve Bussolari, Chuck Crone, Peter Daly, James Flavin, Teresa Hall, William Harman, Walter Heath, Marcia Kastner, Ervin Lyon, Doug Marquis, Ted Morin, Tom Petrillo, Al Reich, Wallace Reid, Richard Sasiela, Kenneth Saunders, Hayden Schultz, Jim Sturdy, Steve Thompson, Jerry Welch, Harald Wilhelmsen, and Dan Wyschogrod.

REFERENCES

1. E.F. Lyon, "Airport Surface Traffic Automation," *Linc. Lab. J.* 4, 151 (1991).
2. J.R. Eggert, R.J. Sasiela, M.P. Kastner, W.H. Harman, H. Wilhelmsen, T.J. Morin, H.B. Schultz, J.L. Sturdy, D. Wyschogrod, and P.M. Daly, "Runway Status Light System Demonstration at Logan Airport," *Project Report ATC-206*, Lincoln Laboratory, 7 Jan. 1994.
3. U.S. Department of Commerce, *U.S. Terminal Procedures, Northeast (NE) Vol. 1 of 3*, effective 30 Apr. 1992 to 25 June 1992.



JAMES R. EGGERT

is a staff member in the Air Traffic Automation group. His research specialty is in airport surface-traffic surveillance and automation. He received a B.A. degree in physics and linguistics from Rice University in 1979 and a Ph.D. degree in physics from Harvard University in 1986. He has also pursued graduate studies in linguistics at the University of Vienna. As a graduate student at Harvard University, he was an I.B.M. Predoctoral Fellow and he received the John Tyndall Prize. Jim is married, with two children, and his hobbies include zymurgy and genealogy.