

---

# Performance of the Runway-Status Light System at Logan Airport

Harald Wilhelmsen, Marcia P. Kastner, Theodore J. Morin, and James L. Sturdy

■ Runway incursions are a persistent problem in airport ground-movement operations. Numerous critical conflicts and several fatal accidents have occurred as a result of unauthorized or otherwise inappropriate entry of aircraft or surface vehicles onto an active runway. Many of these conflicts developed quickly, leaving little time for effective intervention by either the controller or the pilots involved. A reliable system of automatic runway-status lights would be an effective way to prevent such time-critical incursions. The runway-status light system (RSLS) at Boston's Logan International Airport is an off-line proof-of-concept technology-demonstration system designed to show that automatically operated runway-status lights can promptly and reliably transmit runway-status information to pilots and surface-vehicle operators, thereby preventing unsafe runway entry or unsafe takeoff.

The demonstration system does not include actual lights on the airport surface but has relied instead on an illuminated airport model board, which has allowed system development to proceed in a realistic operating environment of live airport traffic without interfering with airport operations. The results of an initial proof-of-concept assessment indicate that the system performs well, even though it is an early prototype. Missed-detection and false-alarm rates are low, and interference with normal airport operations promises to be negligible. The demonstration has shown the technical feasibility of a system of automatic runway-status lights.

**A**S AIR TRAFFIC CONTINUES to increase, traffic loads at major airports are expected to grow proportionally, placing an increasingly heavy burden on the tower controllers who must expedite the movement of this traffic safely and efficiently. Dangerous and occasionally fatal runway conflicts have occurred in the past. Many of these conflicts developed quickly, leaving little time for effective intervention by either the controller or the pilots involved. Some of these conflicts are described in this issue in the article by Harald Wilhelmsen entitled "Preventing Runway Conflicts: The Role of Airport Surveillance, Tower-Cab Alerts, and Runway-Status Lights."

Although surface safety at most of the world's major airports is exemplary, the important task of maintaining and improving safety in the face of increased surface congestion requires continuing effort.

A system of automatic runway-status lights has been proposed as a means of preventing time-critical runway conflicts by indicating runway status to pilots and surface-vehicle operators [1]. These lights would tell pilots when conditions are unsafe for takeoff, and tell both pilots and surface-vehicle operators when a runway is unsafe to enter. Such a warning system, operating continuously, automatically, and promptly without human intervention, represents a last line of

defense against many of those human errors on the part of pilots, tower controllers, and surface-vehicle operators which have in the past resulted in runway conflicts. Such a defense is currently not available on the airport surface.

A proof-of-concept demonstration of a runway-status light system (RSLs) was implemented at Boston's Logan International Airport in 1992 and 1993 by Lincoln Laboratory under the sponsorship of the Federal Aviation Administration (FAA). The objective of this effort was to show that runway-status lights could be operated automatically under the stressing surveillance and operational conditions of a major airport, and to show that these lights could provide reliable, effective protection against inappropriate runway occupancy and runway conflicts without interfering with the normal flow of traffic in the airport movement area.

The Logan RSLs demonstration was conducted with live airport traffic and without interfering with airport control operations. The demonstration system required no input from air traffic controllers and no runway-status lights on the airfield. The lights were simulated on a plan-view computer display of the airport and also on an airport scale model that used optical fibers to simulate the lights.

Surveillance of the airport surface was provided by a modified low-cost marine X-band primary Airport Surface Detection Equipment (ASDE-X) radar, and surveillance of the approach space was provided by listen-only taps to the Airport Surveillance Radar (ASR-9) for beacon returns and the Automated Radar Terminal System (ARTS IIIA) computer for flight identification. This type of surface and approach surveillance requires no new aircraft equipment. The surveillance data were processed by the demonstration system and status-light commands were issued to the model board in real time; processing throughput proved adequate for the heaviest traffic conditions encountered at Logan Airport during the assessment, which approached 110 operations per hour.

A limited but thorough performance assessment of the RSLs was conducted during the spring and summer of 1993 to identify instances of anomalous light operation, quantify the system's end-to-end performance, and verify that the program objectives had

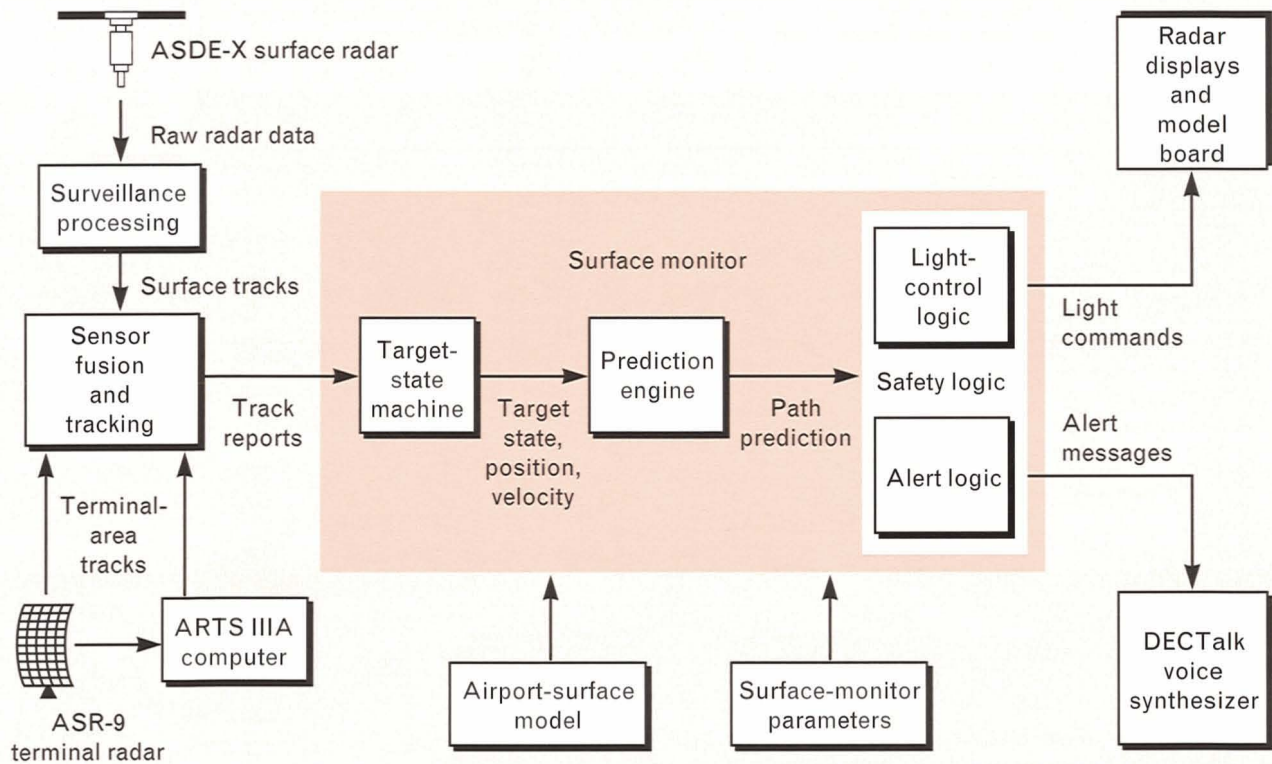
been met. The results are described in this article. The assessment was based on approximately ten hours of recorded traffic, selected to present a well-balanced picture of operations at Logan Airport. All four of the major runway configurations were represented in the data, as were a variety of weather conditions and traffic densities, and most of the busy hours of the operational day.

The primary purpose of the assessment was to quantify the performance of the runway-status light system in its current proof-of-concept state of development. A secondary purpose was to determine the probable causes of the observed anomalies and to identify and prioritize promising corrective actions. Thus the assessment results provide not only a point-in-time description of the system's performance, but also a guide to further development efforts and a basis for judgments about the attainable performance level of a fully developed RSLs.

The article in this issue by James R. Eggert entitled "Demonstration of Runway-Status Lights at Logan Airport" gives a general introduction to the Logan RSLs demonstration system and contains much of the background information needed to understand the descriptions of the various performance anomalies and the performance results in general. Additional detail may be needed, however, with respect to the surface monitor, which is the body of software that monitors and evaluates the state of traffic on the airport surface and in the immediate airspace, and issues the commands that illuminate the runway-status lights. A description of the surface monitor is added in this article to provide an understanding of how light commands are issued.

### **Description of the Surface Monitor**

Figure 1 illustrates the architecture of the surface monitor, which comprises three functional blocks: the target-state machine, the prediction engine, and the safety logic. The surface monitor functions within an airport-specific context; this context is provided by the airport-surface model and the surface-monitor-parameter database. Aircraft-traffic and vehicle-traffic inputs, in the form of track reports, are provided to the surface monitor by the sensor-fusion-and-tracking function. Sensor fusion combines tracks from the



**FIGURE 1.** The architecture of the surface monitor. The surveillance sources for the surface monitor are the ASDE-X surface radar and the ASR-9 terminal radar. Surveillance processing for the ASDE-X radar converts raw radar data into a digital track file for targets on the surface; the ASR-9 radar provides terminal surveillance data for aircraft on approach to the runways. The tracks are obtained from an interface to the ARTS computer, which uses data from the ASR-9 radar. The sensor-fusion and tracking algorithms merge tracks from the ASDE-X and ASR-9 radars and the ARTS system into a single set of tracks. Associated with each track is a track report, which is sent to the surface monitor. The surface-monitor algorithms can be divided into three modules: the target-state machine; the prediction engine; and the safety logic, which consists of light-control logic and alert logic. The target-state machine determines the target's state (for example, whether it is departing, landing, or taxiing). The prediction engine estimates where the target could be in the future by predicting its path on the basis of current position and velocity. The safety logic controls the operation of the safety features, namely, the runway-status lights and tower-cab alerts. Light-control logic determines which runway-status lights should be on and which should be off, and indicates this to the radar displays and airport model board by issuing light commands. Alert logic determines when an existing or potential conflict situation exists between two targets, and sends the appropriate audible alert message to the DECTalk voice synthesizer. As part of its initialization procedure, the surface monitor uses information from two airport-specific databases: the airport-surface-model database and the surface-monitor-parameter database. These two databases contain the site-dependent information from which the prediction engine is built and the target-state machine and runway-status lights are initialized.

ASDE-X surface radar (processed primary radar data), the ASR-9 terminal radar (including beacon code), and the ARTS computer (association of beacon code with flight identification) into a single, seamless, target track-report database containing track number, position, velocity, acceleration, extent (a measure of the size of the target's radar image), and

other pertinent data, which can include altitude, aircraft type, and flight identification.

#### *Surface-Monitor Inputs*

The airport surface presents a challenging surveillance environment, and the processing of the ASDE-X radar returns involves a series of steps designed to

ensure that the digital target tracks passed to sensor fusion accurately reflect the state of traffic on the airport surface. The first step is clutter rejection to eliminate radar returns from the ground as well as from stationary objects on the airport surface that are of no consequence to the RSLs. The second step involves a set of morphological (shape processing) operations on the remaining returns to identify those returns which might correspond to targets of interest, and to determine their centroids. The third step involves scan-to-scan association of these target centroids to form ASDE-X target tracks. Flaws in any of these processing steps, if not rectified in sensor fusion, can lead to erroneous track information being passed to the surface monitor and consequently erroneous light commands.

There are two types of track flaws—missing (dropped) tracks and false tracks. A track drop can be caused by a temporary loss of target return (caused, for instance, by shadowing of the target by a nearby larger target during crowded traffic conditions) or it can be a consequence of failed scan-to-scan association. A false track can likewise result from faulty scan-to-scan association caused by misassociation of unrelated returns from real targets, clutter, or spurious returns. A false track can also be produced by multipath. A multipath return is the result of the radar signal undergoing successive reflections from two or more objects and returning to the radar antenna from the direction of the original reflector but with a time delay caused by the added path length. The result of multipath is a phantom object out-range from the original reflector. Multipath tracks tend to be numerous in particular regions of the movement area, and generally present more of a problem when the ramp and movement area are crowded with aircraft. They can be persistent and difficult to eliminate, because, much like street traffic reflected in a plate-glass window, they tend to behave like real targets.

### *Surface-Monitor Structure*

The surface-monitor algorithms can be divided into three modules: the target-state machine; the prediction engine; and the safety logic, which consists of light logic and alert logic. The target-state machine determines the target's state—for example, whether

the target is departing, landing, or taxiing. The prediction engine estimates where the target could be in the future by predicting its path, based on its current position, velocity, and target state. The safety logic controls the operation of the safety features, namely, the runway-status lights and tower-cab alerts. Alert logic determines when a current or potential conflict situation exists between two targets and sends the appropriate audible alert message to the DECTalk voice synthesizer. Only a limited number of conflict alerts were implemented in the RSLs demonstration to show the ability of the surface monitor to drive a complete safety system incorporating both runway-status lights and audible tower-cab alerts. The alert logic performed well but is not discussed in this article. Further description of the alert logic can be found in Reference 2.

The light-control logic determines the state of the runway-status lights—on (red) or off—and indicates this state to the radar displays and airport model board by issuing commands to the individual lights. The light state conveys runway status only and does not imply clearance from the tower controller. There are two types of runway-status lights: runway-entrance lights (REL) and takeoff-hold lights (THL). The RELs are located at the runway edge on either side of entrances to the runway. The lights are directional and oriented to be visible to entering traffic but not to exiting traffic or traffic still on the runway or on approach. Red RELs indicate that the runway would not be safe to enter at that location. The THLs are located on either side of the runway some distance in front of each normal takeoff location. They are also directional and oriented so they face aircraft in position for takeoff. Red THLs indicate that conditions are not safe for takeoff from that location.

As part of its initialization procedure, the surface monitor uses information from two airport-specific databases in the RSLs: the airport-surface-model database and the surface-monitor-parameter database. These two databases contain the site-dependent information from which the prediction engine is built and the target-state machine and runway-status lights are initialized. The airport-surface-model database contains a centerline description of the runways and taxiways that make up the airport movement area.

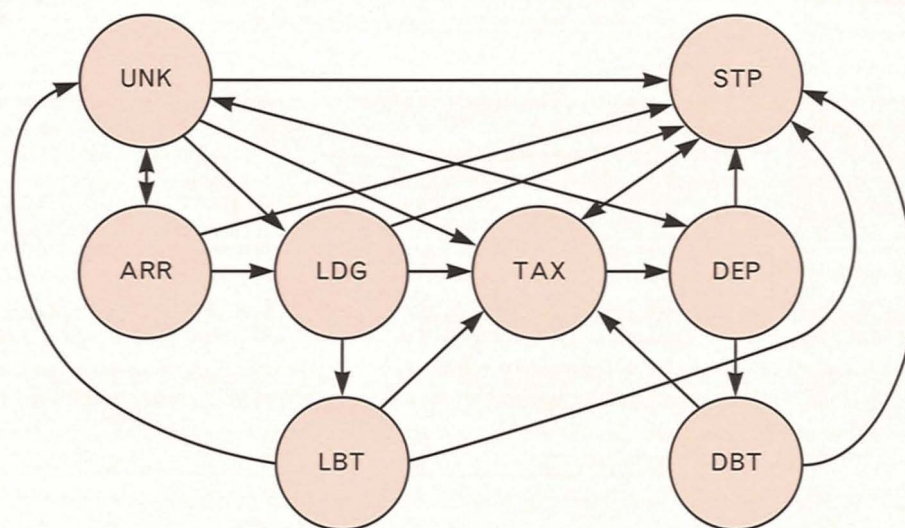
The surface monitor uses this information to define the runway and taxiway structure of the airport and to place the target tracks provided to it by the sensor-fusion function in the proper context. The value of each surface-monitor parameter is generally set for the entire airport but can be individually tuned at any location for optimal performance. In summary, the two surface-monitor databases allow the surface monitor to retain flexibility and site independence.

### *Target-State Machine*

Every target is classified as being in one and only one target state. The target-state machine specifies the transitions *from* one target state *to* another, as opposed to specifying the conditions for being *in* a target state. Figure 2 illustrates the possible transitions between target states, where UNK is the unknown state used for an initial indeterminate 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. These states are described in detail in Reference 2. The following are examples of target-state transitions: (1) the transition from ARR to LDG occurs when an aircraft on approach to an

airport is about to land; (2) the transition from LDG to LBT occurs when an aircraft is unable to complete the landing, such as in a missed approach; and (3) the transition from TAX to DEP occurs when a taxiing aircraft's velocity and acceleration exceed specified parameter thresholds.

There are two advantages to the state-machine approach. First, it avoids the possibility that a target might end up in a situation for which no target state has been defined. With a state machine, a target will always have a defined target state. Second, the target-state machine incorporates hysteresis, a technique for avoiding the problem of jumping back and forth between states because of surveillance errors. For example, the transition from STP to TAX occurs when the velocity of a target that is stopped or barely moving exceeds a parameter threshold. However, the transition from TAX to STP occurs when the target's velocity drops below a lower parameter threshold than for the STP-to-TAX transition. By having two different velocity thresholds separated by an amount greater than the expected surveillance error, small surveillance errors will not cause numerous back-and-forth transitions between states.



**FIGURE 2.** The possible transitions between states in the target-state machine. UNK is the unknown state used for an initial indeterminate 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. A precise statement of the transition rules can be found in Reference 2.


### Prediction Engine

To control the runway-status lights and generate tower-cab alerts, the surface monitor requires knowledge of the predicted positions of all tracked targets. The set of algorithms that determines where a target could be in the future is called the prediction engine. Because the prediction engine takes into account that targets can make turns at intersections of runways and taxiways, a target can have more than one predicted path. Rather than come up with a single set of predicted paths, the prediction engine computes *two* sets of predicted paths, as illustrated in Figure 3. One set, shown as dashed green lines, assumes the target will travel as far as reasonable within a given look-ahead time. The other set, shown as a solid red line, assumes the target is attempting to stop (but not a panic stop) within a given look-ahead time. Thus the two sets of predicted paths, or *trees*, compute reasonable bounds on future target position. In other words, the RSLs path-prediction algorithm does not answer the question, “Where will the target be in a given  $t$  seconds?” but rather the question, “Where *could* the

target be in a given  $t$  seconds?” As we will show below, these two models are essential for the safety logic.

To determine the path-prediction trees, the prediction engine requires prediction models of target motion. There are two models for each target state. The first one, called the *acceleration model*, is used to determine the first set of predicted paths (i.e., the farthest the target could be). The second one, called the *deceleration model*, is used to determine the second set of predicted paths (i.e., the nearest the target could be). Note that these models by themselves do not determine future position because they do not take into account the airport geometry (i.e., which paths can be taken) and how targets would turn from one path onto another path at any intersection.

One of the tasks of the prediction engine is to implement the approach logic, which determines whether an airborne target is on approach to a runway and, if so, which runway. As part of this logic, the prediction engine projects the flight path of an arriving target and determines whether the target could land on a particular runway. To make this determination, the prediction engine uses a turning model that



**FIGURE 3.** Path-prediction trees. The prediction engine computes two sets of predicted paths that show where a target could be in the future, based on two models of target motion, which are functions of target state. One model, called the acceleration model, assumes the target will travel as far as reasonable within a given look-ahead time; its predicted path is represented by the dashed green lines. The other model, called the deceleration model, assumes the target is attempting to stop (but not a panic stop) within a given look-ahead time; its predicted path is represented by the solid red line (the red line is superimposed on the dashed green line). Thus the two sets of predicted paths, or *trees*, compute reasonable bounds on future target position.

allows for S curves as well as straight-in approaches and single-curved approaches.

Sometimes an aircraft can land on more than one runway, for example, in the cases of closely spaced parallel runways or runways that share a portion of an approach path. If the approach model determines that an aircraft on approach can land on more than one runway (i.e., the prediction engine projects it onto more than one runway), then the aircraft is called an *ambiguously projected target*. The logic rule for choosing the runway an ambiguously projected target will probably land on (for purposes of turning on runway-status lights) is to assign the aircraft to the runway in use for landings in the current configuration; this runway is called the *primary* runway.

#### *Logic for Runway-Entrance Lights*

This section describes the operation of RELs at intersections along active and inactive runways. A runway is *active* if it is being used primarily for takeoffs or landings or both in the current airport configuration. A runway is *inactive* if it is not active but can be used temporarily for takeoffs and landings at any time.

Red RELs at runway-taxiway or runway-runway intersections indicate that the runway is unsafe to enter at that intersection because a high-speed target (e.g., departure or landing) is moving along the runway toward that intersection. Controllers call the runway *hot* under these circumstances (we can think of the RELs as being similar to railroad-crossing signals).


Figure 4 illustrates the logic for the runway-entrance lights, which is based on three concepts [2]. The first concept is the *target hot zone*, which is an area ahead of a high-speed target, measured from the front of the target's extent (i.e., from the nose); this area should be free of other targets. Thus entry into the hot zone by other aircraft, surface vehicles, or personnel would be unsafe. The second concept is the *REL activation region* located at runway intersections. Each set of RELs has its own REL activation region associated with it, so that the RELs are operated independently of each other. The idea behind the logic rule for RELs is that RELs are red if a hot zone overlaps their associated REL activation region, and off otherwise. In particular, the RELs behind a target are

always off. An exception to this logic rule is necessary because of the application of a third concept, called *anticipated separation*, which is described below.

In Figure 4 the RELs are represented by rectangles at runway-taxiway intersections but not at runway-runway intersections. The runway-runway intersection poses a problem because high-speed targets taking off or landing on one runway could see RELs along an intersecting runway. An important underlying principle of the RSLs development is not to show red lights to a high-speed target, because there is no assurance that the lights would lessen the danger in such a situation, and they might in fact increase the danger. Thus the light-control logic does not turn on RELs at runway-runway intersections unless one of the runways is being used exclusively as a taxiway.

There are two types of hot zones, each with a different length. The first type is a *t*-second zone, whose length is the distance corresponding to *t* seconds ahead of the target, where *t* is a function of target state and a particular set of RELs. The second type is a whole-runway zone, whose length is the whole runway ahead of the target. The type of hot zone depends on the target's state. For example, an aircraft on approach has a *t*-second zone. However, in two situations, the aircraft "owns" the whole runway and thus has a whole-runway zone. These two situations are (1) a departure and (2) an arrival that has just crossed the runway threshold (i.e., the near end of the runway as seen by an arriving aircraft) and is not yet in a landing rollout. The length of a *t*-second zone is determined from the target's acceleration model.

The condition for turning off RELs whose activation region overlaps a hot zone uses the concept of *anticipated separation*. In this concept, controllers can issue clearances and instructions to aircraft in anticipation that legal separation between aircraft will exist when required, even though legal separation does not currently exist [3]. For example, an aircraft cannot legally cross a runway in front of a departure. However, a controller can issue "taxi across" instructions to an aircraft waiting to cross a runway at a taxiway hold line, even though a departure on that runway has not yet passed the taxiway intersection, in anticipation that by the time the waiting aircraft starts across the runway, the departure *will* be past the



**FIGURE 4.** Hot-zone and runway-entrance light (REL) activation regions. The logic for RELs is based on three concepts. The first concept is the *target hot zone*, which is an area ahead of a high-speed target that should be free of other targets because entry into the hot zone by other aircraft, surface vehicles, or personnel would be unsafe. The length of the hot zone depends on the target's state. It could be a distance represented by  $t$  seconds ahead of the target, or it could be the whole runway ahead of the target. The second concept is the *REL activation region*, located at runway intersections. Each set of RELs has its own REL activation region associated with it, so that the RELs are operated independently of each other. A set of RELs is on (red) if a hot zone overlaps its associated REL activation region, and off otherwise. An exception to this rule is when RELs are off, even though a hot zone overlaps their activation region, because of the application of a third concept, called *anticipated separation*. In this case, the RELs turn off a few seconds before a high-speed target on the runway is predicted to pass the intersection in anticipation that, by the time another target waiting at the taxi-hold line starts to cross the runway, the high-speed target will already be past the intersection.

intersection. The delay in crossing occurs because of the time the waiting aircraft takes to get ready to move before it actually starts across the runway. To avoid interference with the controllers and to allow for surveillance delays, the light-control logic uses anticipated separation to turn off the RELs a few seconds before a target's hot zone is predicted to exit the REL activation region (by using the target's deceleration model) but at approximately the time when a controller would issue the "taxi across" instruction. The use of anticipated separation helps to expedite the flow of traffic on the airport surface without compromising safety.

#### *Logic for Takeoff-Hold Lights*

THLs are operated for both active runways and inactive runways. Because inactive runways may be used

temporarily for takeoffs at any time, the operation of THLs should apply to inactive runways as well as active runways. Also, THLs are installed at locations used for intersection takeoffs, i.e., takeoffs partway down the runway at intersections with taxiways, as well as full-length takeoffs, which are from the end of the runway.

THLs at a given location are red if two conditions are satisfied simultaneously: (1) a target is in position for takeoff or starting its takeoff at this location, and (2) the runway is not safe for takeoff at this location. The first condition implies that a target must be in position to see the THLs for them to be red. This requirement is not true for RELs, which turn red regardless of whether a target is in position to see them. The reason for this difference is that THLs, unlike RELs, are seen by aircraft on active runways. If the




light-control logic did not require a target in position for takeoff in order to turn on the THLs, then a situation could occur in which the THLs are red and an aircraft on approach to the runway could see the red lights. As described in the previous section, the light-control logic was designed to avoid showing red lights to a high-speed target in order to avoid a potentially dangerous situation.

The precise definition of the first condition for a THL to be red can be found in Reference 2, but basically it means that a target is in a specified area of the runway called the *arming region*, and the target's heading is approximately the direction of the runway, as illustrated in both Figure 5 and Figure 6.


The second condition for a THL to be red depends on whether there is another target that could come in conflict with a departure. This condition is also different from the logic for RELs in that THLs turn red based on the actions of two targets, whereas RELs turn red based on the action of a single target. There are two ways to satisfy this second condition, depending whether the two targets are on the same runway or on intersecting runways. Figure 5 illustrates the first way: the second condition is satisfied because a

target is inside the *THL activation region*, which is an area that includes the runway ahead of the lights as well as an extension on either side of the runway. The THL activation region is an area ahead of a target in position for takeoff, or starting its takeoff roll, that should be free of other targets before takeoff commences. In the current logic, this region extends from the THLs to the end of the runway. Because the THL activation region is fixed for each set of THLs, the prediction-engine models of target motion are not used in satisfying this condition.

The second way for the second condition to be satisfied is when there is a potential conflict with a high-speed target on an intersecting runway. The logic for this case is illustrated in Figure 6; it is based on the concept of an *intersection window*, which is an area at the intersection of two runways. The second condition is satisfied (i.e., the runway is unsafe for takeoff) if target A, which is in position for takeoff or starting its takeoff, and target B, which is in any target state except "stopped" and "taxiing," could be in the intersection window simultaneously if A started to take off. The time interval during which target B could reside in the intersection window is calculated as fol-



**FIGURE 5.** Example of takeoff-hold lights (THL) that are turned on when two targets are on the same runway. THLs are turned on (red) if two conditions are satisfied simultaneously: (1) a target is in position for takeoff or starting its takeoff, and (2) the runway is unsafe for takeoff. In this figure, the first condition is satisfied because a target is in the arming region and lined up with the runway. The second condition is also satisfied because another target is in the THL activation region.



**FIGURE 6.** Example of THLs that are turned on when two targets are on intersecting runways. THLs are turned on (red) if two conditions are satisfied simultaneously: (1) a target is in position for takeoff or starting its takeoff, and (2) the runway is unsafe for takeoff. In this figure, the first condition is satisfied because a target (target A) is in the arming region and lined up with the runway. The second condition is also satisfied because another target (target B) is on an intersecting runway, and the light-control logic has determined that A and B could be in the intersection window simultaneously if A started to take off.

lows: the prediction engine uses the acceleration model to determine the earliest time target B could enter the window and the deceleration model to determine the latest time target B could leave the window. Because target A is stopped or moving slowly, its interval in the window is calculated by using only the acceleration model for departures. Note that this logic rule again uses the concept of anticipated separation. Rather than wait for target B to cross the intersection before turning off the THLs, the logic turns the lights off in anticipation that target B will be across the intersection before target A could be in conflict with target B if target A became a departure.

The logic for operating THLs can now be stated simply. If the first condition is satisfied, then the THLs are called *armed*. If the second condition is satisfied in either of the two ways just described, then the THLs are called *activated*. Thus the THLs are red, also referred to below as *illuminated*, if the lights are both armed and activated, and off otherwise.

### RSLS Performance Assessment

Development of the runway-status light demonstration system was carried out entirely off line. That is, although actual live traffic data were used to drive the system, the system itself did not affect either tower operations or airport traffic. This approach was appropriate during the proof-of-concept phase; it did, however, preclude investigation of a number of fundamental questions regarding the operational suitability of the concept, primarily those human-factor issues concerning controller and pilot acceptance and understanding of the system and its effect on controller and pilot work load.

These questions can be conclusively addressed only in an actual field test. A comprehensive, quantitative assessment of the system's performance is a precondition for such a field test. This pre-field-test assessment must, at a minimum, quantify the performance of the runway-status lights, ensure that the system is well

tuned and that it does not appear to interfere with the normal flow of airport traffic, and define the environmental limitations of the system, such as performance in rain.

A limited preliminary assessment of this type was performed in the spring of 1993. The assessment had one primary objective and four secondary objectives. The primary objective was to provide quantitative data on the system's end-to-end performance as it would be experienced by the system's users—namely, pilots, tower controllers, and surface-vehicle operators. These users would want to know if the lights, were they installed, would improve safety and situational awareness, and if they might interfere with the normal flow of traffic.

The four secondary objectives were intended to focus further development efforts. These objectives were: (1) identify, interpret, classify, and prioritize light anomalies; (2) evaluate the appropriateness of the tunable surface-monitor parameters and other adjustable entities; (3) identify prospective solutions to the observed anomalies and estimate the level of effort involved in implementing them; and (4) estimate the level of performance that could be achieved with further development effort.

In keeping with the intent to characterize the system's performance from the point of view of the prospective users, only light anomalies that would have been observable from the cockpit or surface-vehicle cab were counted. The assessment counted three types of light anomalies: (1) lights that were on when they should not have been on (false alarms), (2) lights that were not on when they should have been on (missed detections), and (3) lights that were on longer than necessary (interference). Light anomalies can be caused by flaws in the surface monitor itself or by flawed track inputs to the surface monitor. Flawed inputs include track drops or false tracks resulting from surveillance-processing or sensor-fusion anomalies. Track drops on approach can also be caused by factors external to the surveillance system, such as failure of the aircraft's beacon transponder, although such failures are rare and were not observed in the performance-assessment data. Transponder-related track failures could be remedied by adding a capability to track on the basis of ASR-9 primary returns rather

than secondary returns. The RSLs does not currently have this capability. The assessment counted the three types of anomalies separately for the two types of runway-status lights and the four major runway configurations employed at Logan Airport, and attempted to determine the cause of each anomaly.

### *Performance Measures*

The system's performance is determined by four performance measures that are designed to gauge its ability to enhance safety while remaining transparent to normal operations. These performance measures, which are defined in the glossary below, are *missed detection*, *false alarm*, *interference*, and *light infringement*. The definitions of the performance measures involve the concept of a light threshold, which is an imaginary line near a runway-status light that should not be crossed while the light is illuminated. The concept is similar to but distinct from the taxi-hold lines on taxiways. The first and last of the four performance measures—missed detections and light infringements—address the system's effectiveness, while the other two measures—false alarms and interference—address its transparency. The first three performance measures are collectively referred to as light anomalies. All four measures will be discussed briefly and illustrated with examples. For simplicity, the narrative refers to aircraft but the arguments apply equally to surface vehicles, where appropriate.

*Missed Detection.* A missed detection can be defined in a narrow or broad sense. The narrow definition applies to a failure to detect a developing conflict situation; this definition requires both that the light was off when it should have been on, and that a conflict developed. Because runway conflicts are rare events, we do not expect to see evidence of this type of missed detection in the limited amount of traffic used for the performance assessment. The broad definition is a missed detection of a *potential* conflict situation; this broad definition merely requires that the light was off when it should have been on, and that the traffic picture was such that this failure could have permitted a conflict to develop. The broad definition is used here. Thus a missed detection does not imply failure to detect an actual conflict situation.

A missed detection of a condition that should have

activated a THL could be caused by either a track drop of an aircraft in or about to enter the light's arming region, or a track drop of a target in or about to enter the activation region, or a failure to track or project crossing high-speed traffic properly. A THL missed detection could also in principle, but rarely in practice, have multiple causes. An untracked aircraft holding in position in an arming region is one of the more likely causes of missed detections, because such a situation can persist for an extended period and possibly give rise to several missed detections.

A missed detection of a condition that should have activated an REL could be caused by a track drop of an aircraft engaged in high-speed operations, such as arrival, landing, landing rollout, or takeoff. This type of missed detection is usually brief. Another cause of missed detections is ambiguity with regard to the landing runway. The runway arrangement at Logan Airport is both complex and compact. An arrival on approach to one runway sometimes appears to be lined up with another runway; the landing runway is not apparent until later in the approach sequence. The surface monitor prediction logic can be late in making the determination, with the result that the RELs near the approach end of the landing runway illuminate later than desired.

RELs illuminate whenever activated, regardless of whether anybody is in a position to see them. This mode of operation is unlike that for the THLs, which must be armed as well as activated before they illuminate. The armer is generally in a position to see the illuminated THL. With the RELs a distinction must be made between observed and unobserved illuminations. Only observed anomalies are counted. This choice is consistent with the focus of the assessment, which is on the interaction between the lights and pilots or surface-vehicle operators.

The distinction between the narrow and broad definitions of a missed detection is an important one that must be kept in mind when interpreting the system performance results presented later in this article. Failure to detect an actual (narrow definition) conflict situation could have serious consequences. Failure to detect a potential (broad definition) conflict situation is less serious, albeit clearly undesirable. A missed detection of a potential conflict situation does not mean that a conflict will occur, because no conflict situation usually exists. This type of missed detection implies that, at the specific location on the airport surface where the missed detection occurs, and for the typically short duration of the missed detection, the system offers no protection. Conditions at that location

---

---

## GLOSSARY OF PERFORMANCE MEASURES

---

**Missed Detection** a failure of a runway-status light to illuminate as it should, as judged from the intent of the light-control logic and the state of traffic on the airport surface and in the immediate airspace.

**False Alarm** a runway-status light illumination that does not reflect the real state of traffic on the airport surface and in the immediate airspace.

**Interference** occurs when a runway-status light is on, in accordance with the rules of the safety logic and the state of traffic on the airport surface and in the immediate airspace, while a safe operation is under way that leads to light-threshold crossing within a specified time.

**Light Infringement** an aircraft or surface vehicle advances beyond the light threshold while the runway-status light is illuminated, and this advancement is unsafe—that is, the runway-status light is correctly illuminated in accordance with the intent of the light-control logic and the state of traffic on the airport surface and in the immediate airspace.

then briefly revert to the conditions that exist at all airports today.

*False Alarm.* A false alarm is the illumination of a runway-status light that does not reflect real traffic. A false illumination of a THL might be caused by a false activation track and a real arming track, or by a false arming track and a false or real activation track. As a general rule, observed false alarms are caused by false activation and real arming, but there are exceptions. The definition of an observed false alarm makes no distinction on the basis of the activity or identity of the occupant of the arming region. The occupant could, for example, be an aircraft holding in position or crossing the runway slowly, or it could be a surface vehicle in an arming region on an inactive runway.

If lights were installed on the airport surface, a false alarm might actually cause interference by interrupting or delaying a departure or by prompting a pilot to request verification of a clearance. The performance assessment nevertheless makes the distinction between an observed false alarm and interference, because this distinction is important for a clear understanding of the system's performance.

A false REL illumination could be caused by a false high-speed track on approach or on the runway, generated by multipath or misassociation. A surface vehicle traveling at high speed on the runway might also activate the RELs, but this would not constitute a false alarm because the traffic is real; likewise a helicopter approaching the helipad could be erroneously identified as an arrival to one of the runways. Such instances might or might not be classified as interference, depending on the circumstances. As discussed above, only *observed* false REL illuminations are counted, reflecting the point of view that instances in which an REL illuminates an empty scene are of no consequence to the system's users.

*Interference.* According to the definition, interference is caused by real traffic. It occurs mostly when the light-control logic does not take full account of the operational realities of the airport environment, or when the surface monitor is handicapped by its lack of knowledge of pilot and controller intent. An example is a runway crossing in front of an aircraft on landing rollout. The surface monitor's target-state machine currently does not have a landing-rollout

state, and this limitation causes the RELs to illuminate too far ahead of the aircraft, with the result that the lights on rare occasions interfere with safe crossing operations farther down the runway. Another potential interference situation arises when an aircraft uses a runway to taxi to another runway for departure. If this taxiing aircraft enters an arming region at taxi speed when the runway ahead is occupied or predicted to be occupied, the THLs will turn on even though the taxiing crew has no intention of taking off. Whether this situation would result in interference in an operational setting is unclear at this point.

Interference by an REL can be caused by misprojection or ambiguous projection of airborne traffic, leading the surface monitor to conclude that a landing is imminent on a given runway when in fact it is not. But occasionally the REL is illuminated by an approach that develops into a late sidestep to another runway. This event would not be classified as interference because the projection was correct up to the moment the sidestep was initiated.

*Light Infringement.* A light infringement differs from interference in that the light infringement results in an apparently unsafe situation. We cannot determine with certainty from the recorded surveillance data if safety was in fact compromised in a given situation, because the tower VHF radio channels were generally not monitored during the data collection, nor were video recordings made of the airport traffic. This limitation is of no consequence for the performance assessment, however, because no light infringement was seen in the data. The absence of light infringements is not surprising; such events are exceedingly rare. A search for light infringements was nevertheless part of the performance assessment, because such an occurrence would have provided direct proof of the potential contribution of the RSLs to runway safety.

Because light infringements were not observed, the remaining discussion is limited to the three light anomalies—namely, missed detections, false alarms, and interference—described earlier.

#### *Normalization*

To be meaningful, the anomaly counts must be presented in a way that takes the traffic volume into ac-

count, such as, for instance, false alarms per hundred operations. This normalization scheme is simple and easy to understand; in the context of this assessment the anomaly count per hundred operations can be loosely equated to anomalies per “average” hour, because one hundred operations correspond to approximately an hour of traffic under average conditions at Logan Airport.

Suggesting that there is a simple relationship between anomalies per hundred operations and anomalies per average hour amounts to a statement about linearity. The average number of anomalies per hundred operations is probably not completely insensitive to the density of the traffic in the underlying data. Because the performance assessment looks for *observed* anomalies, the events counted call for one light activator and one observer; that is, they imply interaction between two agents and thus a potential density dependence. (Appropriately, the same functional relationship is true for the problem the RSLS is designed to address, namely, runway incursions.) If such quadratic effects are present, they were not discernible in the assessment data, perhaps because of the limited size of the data sample, but also possibly because the data were intentionally collected under traffic conditions that were representative of typical Logan Airport traffic conditions, which means that more data were collected in near-average conditions than under very low and very high traffic loads.

Because an assumption of linearity appears to be warranted in the context of this performance assessment, the results are presented in terms of anomalies per hundred operations. The assessment results imply a survey of the entire airport. That is, “five anomalies per hundred operations” means that five anomalies are, on the average, observed by the pilot population on the airport in the course of one hundred operations. Thus we may say that five anomalies are experienced airport-wide in the course of one average hour. An individual crew can also expect to encounter five anomalies in the course of one hundred operations but, in general, considerably fewer than five anomalies in one hour of cumulative taxi time. This difference occurs because the total taxi time accumulated by a crew in the course of one hundred operations is much more than one hour.

### *Logan Airport Runway Configurations*

Logan Airport is usually operated in one of four major runway configurations. The choice of configuration is dictated primarily by wind conditions, but also by ceiling and visibility. Other, more restrictive, configurations are employed in special circumstances, such as during snow removal and runway maintenance. The major configurations are described briefly below, with reference to Figure 7. Deviations from the usage indicated can occur upon pilot request and at the discretion of the local controller.

*4/9 Configuration.* Arrivals to Runway 4R, departures from Runway 9, arrivals and departures to and from Runway 4L, occasional departures from Runways 4R and 15R. This configuration is used when the wind is from the northeast or east, as is commonly the case during the winter and spring months, or during deteriorating weather conditions.

*22/27 Configuration.* Arrivals to Runways 22L and 27, departures from Runway 22R, occasional departures from Runway 22L, occasional arrivals to Runway 22R. This is generally a summer configuration, used when the wind is from the southwest.

*33/27 Configuration.* Arrivals to Runways 33L, 27, and 33R, departures from Runway 27. Non-jet departures from Runway 33L at Taxiway G and occasional full-length jet departures from Runway 33L. This is generally a winter configuration, used when the wind is from the northwest.

*15/9 Configuration.* Arrivals to Runways 15R and 15L, departures from Runway 9, occasional departures from Runway 15R. This configuration, used in southeast airflow, is seen less often than the other three.

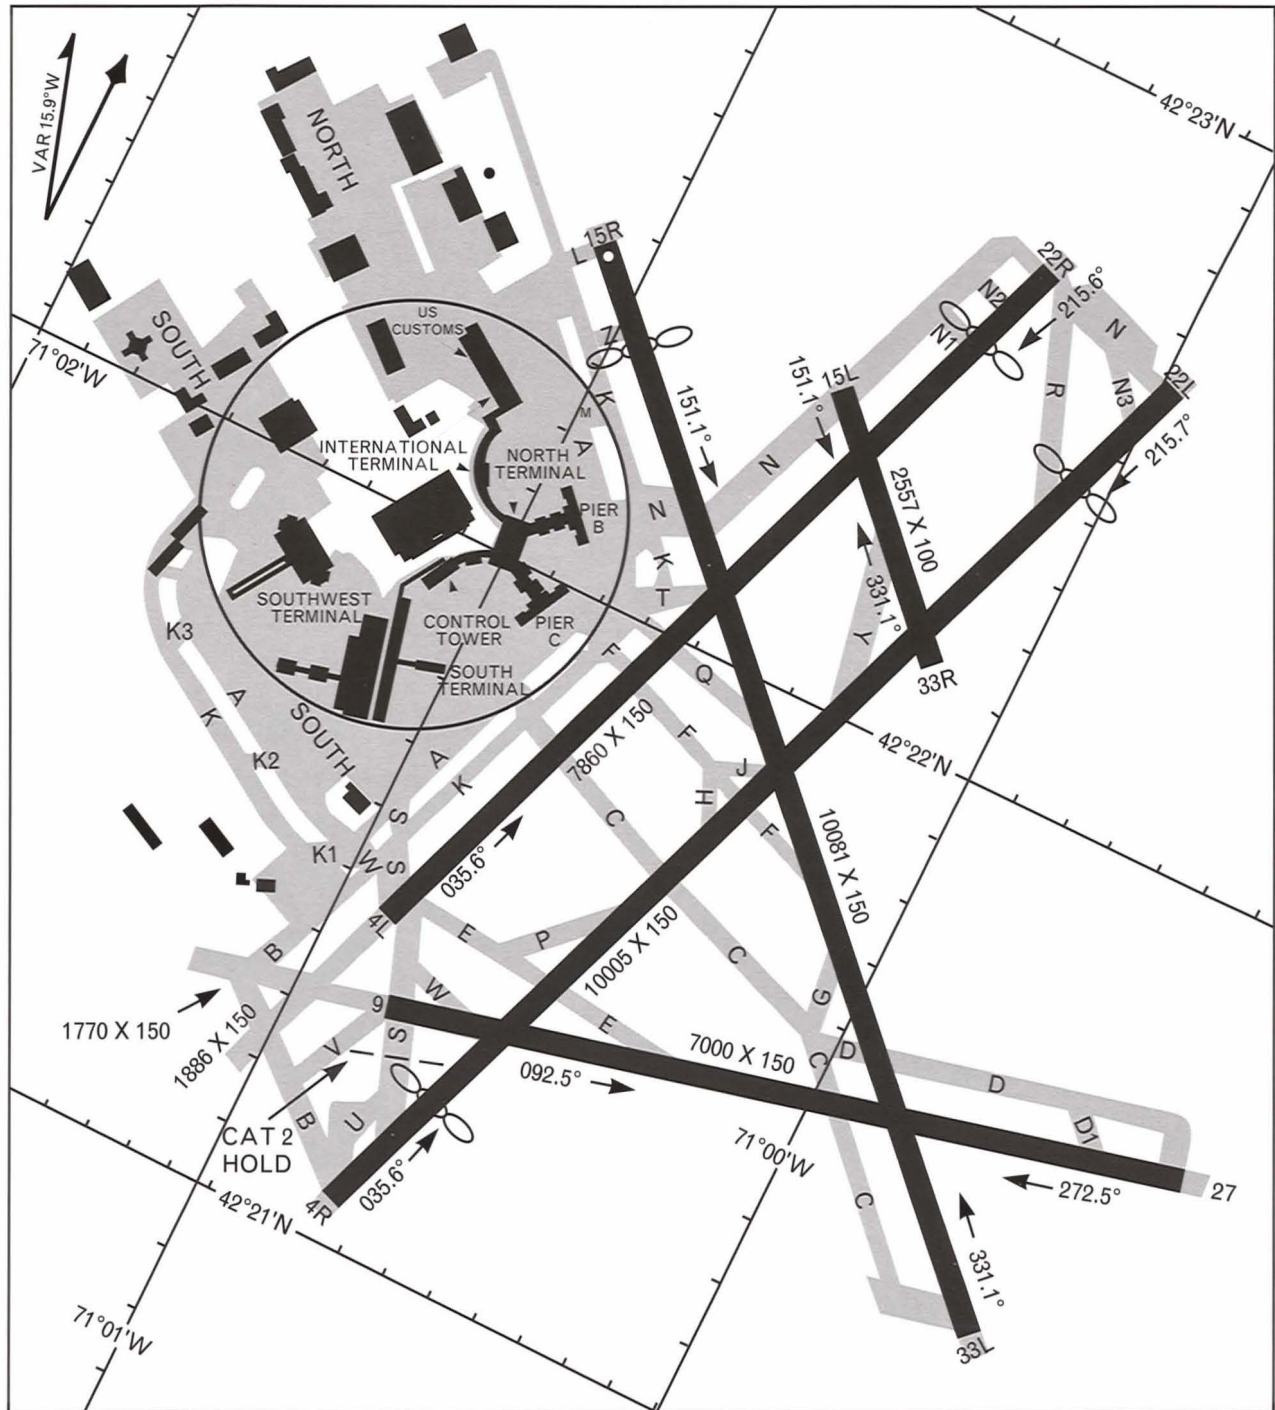
During a period of one year, we can expect to encounter each of the first three configurations approximately 30% of the time, and the fourth configuration approximately 10% of the time.

### *Performance-Assessment Data*

The data for the performance assessment were gathered during a six-week period in March and April of 1993, reflecting predominantly late-winter and spring conditions. In spite of the short period available for data collection, all major runway configura-

tions and most common weather and traffic conditions were captured. Altogether, more than five thousand operations were recorded. A smaller but well-balanced subset of approximately eight hundred operations was selected for detailed assessment.

Table 1 summarizes the performance-assessment data. Listed by time of day, regardless of date and runway configuration, the traffic segments can be viewed as making up the bulk of a composite day of operations at Logan Airport. Thirty percent of the data



**FIGURE 7.** Runway and taxiway configuration at Logan International Airport in Boston, Massachusetts. The primary configurations used are the 4/9 configuration, 22/27 configuration, 33/27 configuration, and 15/9 configuration.

were recorded in instrument meteorological conditions (IMC) and seventy percent in visual meteorological conditions (VMC). There are approximately three hours each of the 4/9 and 22/27 configurations and two hours each of the other two configurations.

The data represent most hours of the operational day and include weather conditions ranging from sunny and calm, to foggy, to rainy and windy, to the early part of the major blizzard of 13 March 1993. The traffic intensity ranges from moderate to heavy, approaching one hundred and ten operations per hour during parts of the eighth traffic segment.

The data-collection sessions included neither tower-channel voice recordings nor video recordings of the airport traffic. Although sufficient information was collected to permit a thorough assessment of the technical performance of the runway-status lights, the lack of a complete record of the operations precludes definitive interpretation of the recorded events from an operational point of view.

#### *Data Analysis*

The ten traffic-data segments of Table 1 were analyzed to identify all light anomalies that satisfied a prescribed set of inclusion criteria. These criteria, or

*counting rules*, are described later in the article. The discussion of counting rules is followed by a detailed description of three specific anomalies, as well as a general summary of the overall assessment results. The three examples serve to illustrate the complexity of the airport surveillance environment and the variety of challenges presented to the RSLs by the physical and operational environment at Logan Airport.

To simplify the presentation, the discussions are presented as if the runway-status lights were actually installed on the field. That is, they refer to lights illuminating "in front of aircraft," "causing interference," and so on, even though there are currently no runway-status lights at Logan Airport.

The current embodiment of the RSLs, with no presence either in the tower cab or on the field, makes possible an intuitive, qualitative assessment of the system's transparency. One indication of a well-tuned and transparent safety system is a traffic display on which the traffic flows as if the pilots and surface-vehicle operators could see, and were responding to, the runway-status lights. Observation of the interplay between the lights and live traffic on the RSLs display clearly shows that the system works well by this criterion, with few anomalies.

**Table 1. Airport Traffic Data Used for the Performance Assessment of the Runway-Status Light System**

<i>Approximate Time Interval (local time)</i>	<i>Configuration</i>	<i>Weather</i>	<i>Operations</i>	<i>Operations per Hour</i>
07:50–08:45	4/9	IMC (fog, drizzle, windy)	70	76
09:05–10:00	22/27	VMC	87	93
09:50–11:00	15/9	IMC (snow, windy)	55	47
10:35–11:35	33/27	VMC	94	94
13:15–14:05	15/9	VMC	63	73
14:50–16:10	33/27	VMC	99	74
16:00–17:15	4/9	IMC (fog, rain, windy)	91	73
17:35–18:35	22/27	VMC	103	103
19:15–20:15	4/9	VMC	91	94
19:50–20:45	22/27	VMC	75	83



Despite this assertion, we must bear in mind that the introduction of runway-status lights into the runway environment raises important human-factors questions and technical questions that cannot be answered conclusively on the basis of an off-line experiment. An operational evaluation with actual lights on the airport surface is required. Pending such an evaluation, however, our observations of the interplay between real live traffic and the simulated lights yield valuable insight into the interactions between the lights and the airport surface traffic. These observations also provide information on the proper tuning of the safety logic. Achieving noninterference is to a large degree a matter of tuning the safety logic so that the lights do in fact operate in harmony with the flow of traffic in all normal operational situations. The RSLs safety logic is specifically designed to be tunable to ensure this smooth interplay between lights and traffic.

#### *Counting Rules*

The counting conventions are consistent with the assessment's objective, which is to characterize the system from the user's perspective. That means asking questions about how the RSLs will affect pilots during ground operations. Answering these questions requires not only an assessment of the operation of the lights themselves, but also an evaluation of the probability that an anomaly will be seen by pilots or surface-vehicle operators. The latter information can be gotten only by observing the interplay between the lights and the traffic. It cannot be determined by assessing the functional elements (radar processing and tracking, sensor fusion, or surface monitor) of the RSLs in isolation.

In keeping with the intent to evaluate the system from the user's perspective, the assessment counts only observed false alarms and missed detections—that is, anomalies that would have been observable from the cockpit. This distinction is a straightforward concept for the THLs; it usually requires simply that an aircraft is positioned in the light's arming region. Although the definition is not much more complicated in principle for the RELs, it is slightly more involved in practice, because there is no well-defined region in which to look for an observer. The assessment

identifies observed REL false alarms and missed detections by the following three criteria: to contribute to the count of observed REL missed detections or false alarms, an aircraft (1) must be within a certain distance of the light during the anomaly, (2) must be first in line, and (3) must eventually enter the runway at that location. The first two criteria address the question of whether the pilot is likely to be looking at the light, the second whether he or she is likely to be concerned with it. These criteria might suggest that the identification of observed REL anomalies is a difficult process, fraught with uncertainty, but as a practical matter the flow of taxi traffic is such that little ambiguity arises.

False alarms and missed detections of less than a four-second duration (one or two scans of the radar antenna) are counted separately from the longer ones. This distinction does not mean that short anomalies are viewed as inconsequential; all anomalies are undesirable, regardless of duration. But the impact of long and short anomalies can be quite different, both in terms of the technical effort required to eliminate them and in terms of their effect on pilots and controllers in an operational setting.

Instances of interference are, by definition, all observed. According to the definition, interference requires actual or imminent crossing of the light threshold while the light is illuminated. All instances of interference are counted, with no distinction attempted on the basis of duration except for a qualitative judgment about severity.

When a surface vehicle causes a light anomaly that affects an aircraft, it is counted as any other. A REL anomaly that affects a surface vehicle is also counted. But the operation of a THL that is armed by, for instance, a service vehicle parked on an inactive runway is not included in the count. Although the affected THL sometimes illuminates as normal airport traffic enters or is projected to cross the runway downfield of the parked vehicle, such an illumination is considered inconsequential because a takeoff by the surface vehicle is out of the question.

The performance assessment described in this article counts only *observed* light anomalies—that is, anomalies that would have been observable by cockpit crews and surface-vehicle operators in the vicinity

of the runways. The decision to count only observed anomalies was made in order to produce results that would be meaningful to pilots and drivers on the airport surface—i.e., to answer the question, “How will these lights affect *me*?”

There were essentially two reasons we decided to count only observed anomalies. First, because the observed anomalies are a subset of all anomalies, the task of characterizing and classifying observed anomalies was manageable, given the resources available. Second, and more importantly, the rhetorical question posed above can be answered only by counting observed anomalies directly. A total anomaly count would not give the pilot or airside surface-vehicle operator any useful information about how he or she might be impacted by the lights, nor would the controller be able to deduce this information. The reason is that the probability that a light anomaly will occur and the probability that that anomaly will be observed are not independent probabilities; they are coupled via the traffic pattern and the surveillance characteristics of the airport surface. Only a direct count of observed anomalies can reveal how the user is likely to be affected by the lights.

#### *Examples of Light Anomalies*

The performance assessment showed that the RSLs perform well, with few false alarms and missed detections and little interference with normal operations, even under conditions of heavy traffic. The level of performance was found to be more than adequate for purposes of technology demonstration. Some instances of anomalous operation were uncovered, however, including the three anomalies described below.

These three examples, selected from the assessment data listed in Table 1, represent the major anomaly categories—that is, THL missed detection and REL and THL interference. They were selected to illustrate how light anomalies can occur, and also to demonstrate the variety of challenges that may be encountered when implementing an automatic safety system for the highly complex surveillance and operational environment of a major airport. In addition, they highlight important elements of the surface monitor and the light-control logic.

#### *Example 1: THL Missed Detection*

Some of the most challenging problems encountered in surveillance processing and tracking are caused by multipath. Multipath returns often exist in great profusion in certain regions of the airport movement area. The placement of the X-band radar in relation to the terminal buildings, ramps, runways, and taxiways at Logan Airport is such that multipath sometimes appears on the active runways, principally at the approach ends of Runways 9 and 15R and the southwest half of Runway 4L/22R.

Multipath-rejection algorithms are generally successful in discriminating between multipath returns and returns from real aircraft and surface vehicles. Thus multipath returns seldom give rise to high-confidence tracks that can cause the runway-status lights to illuminate and produce false alarms.

A more difficult problem exists with regard to misassociation of real tracks with multipath tracks. Misassociation is possible when a multipath track passes close to a real high-confidence track. Track swap can occur, especially when the real track is stationary and therefore does not have a preferred direction of motion. Figure 8 shows an example of misassociation with multipath. This misassociation occurred in the arming region of the THL on Runway 9 while an aircraft was holding in position there. The result was that the track of the aircraft in the arming region was dropped, causing the THL, which was illuminated because of an arrival to Runway 4R, to turn off a few seconds prematurely.

Figure 8(a) shows the situation prior to the track drop. The aircraft with track number 4753 has just entered the arming region and is in position for take-off on Runway 9, and the THL in front of it is illuminated because of the Runway 4R arrival (as yet outside the picture to the right). The holding aircraft has, on the basis of its past track history, been designated a high-confidence track by the surveillance-processing function. It has also been accepted by the sensor-fusion function, as indicated by the superimposed yellow arrow icon. The aircraft on final approach to the intersecting Runway 4R has illuminated the Runway 4R RELs and the Runway 9 THL in accordance with the light-control logic for targets on intersecting run-

ways, and the illuminated THL tells the pilot in position for takeoff that there is now insufficient time to begin takeoff and clear the runway intersection safely ahead of the landing aircraft. A multipath return, shown in blue to signify a low-confidence track, is moving down Runway 27 toward the aircraft holding in position (note that Runway 27 is physically the same runway as Runway 9, but in the opposite direc-

tion). The likely source of this multipath is the aircraft on Taxiway K, with track number 5601, as indicated by the radial line from the radar location. This association between multipath and multipath source is not obvious from the still picture, but the scan-to-scan progression of the radar imagery makes the relationship apparent.

Figure 8(b) shows the situation four seconds later.

---

---

## SURFACE MONITOR DISPLAY SYMBOLOGY FOR FIGURES 8 THROUGH 10

---

**THL Arming Region** blue outline

---

**THL Activation Region** amber outline

---

**Low-Confidence Track from Surveillance Processing** blue circular icon

---

**High-Confidence Track from Surveillance Processing** green circular icon

---

**Sensor-Fusion Track** yellow “arrow” icon. The size of the arrow is proportional to the size of the target’s radar image. The direction of the arrow indicates direction of motion; the end of the arrow shaft indicates the rearmost extent of the target image. The surface monitor recognizes only sensor-fusion tracks.

---

**Runway-Entrance Light** red bar across taxiways at runway edge. The lights are directional and point away from the runway. They will not be seen by traffic on the runway. They could be located on either side of the taxiway, rather than in-pavement.

---

**Takeoff-Hold Light** red bar across the runway. The lights are directional, pointing toward the aircraft in position for takeoff. They could be located on either side of the runway, rather than in-pavement.

---

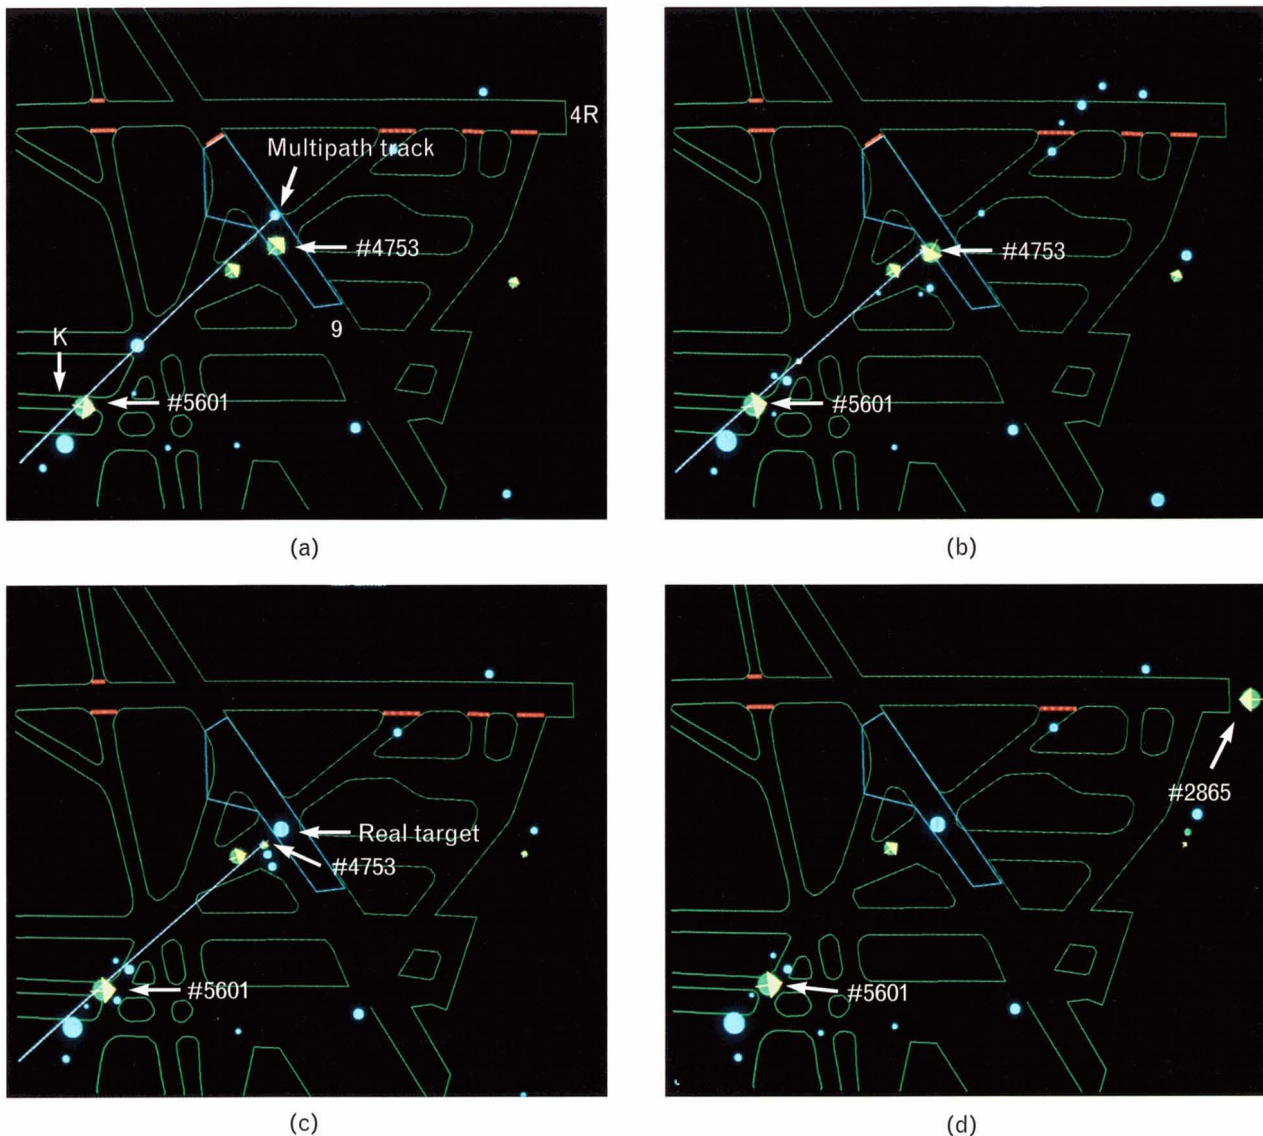
**Data Tags** text information giving track information, such as track number, altitude, speed, and (for arrivals, on alternate surveillance updates) flight identification and aircraft type. Tags are shown in yellow letters and attached to the target with a yellow line.

---

**Path Projections** when enabled, red lines with green extensions projecting ahead of moving targets. Projections are produced by the prediction engine. The red line indicates how far the target *must* travel in a specific time, given a state-dependent deceleration model. The green line indicates where and how far the target *could* travel in the same amount of time, given a state-dependent acceleration model.

---

**Target State** determined by the target-state machine and shown in cyan lettering next to the target whenever the path-prediction lines, or *trees*, are enabled.



**FIGURE 8.** Track drop by misassociation with multipath. This sequence illustrates a track drop in the Runway 9 THL arming region caused by misassociation with multipath. (a) The aircraft with track number 4753 has just entered the arming region and is in position for takeoff. An aircraft outside the picture to the right is on final approach to the intersecting Runway 4R and has illuminated the Runway 4R RELs and activated the Runway 9 THL, in accordance with the light-control logic for targets on intersecting runways. Another aircraft, track number 5601, is taxiing along Taxiway K at lower left in the picture. This aircraft produces multipath returns in the vicinity of the Runway 9 approach end. Surveillance processing has tentatively identified a track from among the multipath returns, which is shown as a low-confidence (blue) return. The radial line from the radar location emphasizes the probable association between the aircraft on Taxiway K and the multipath return in the arming region. (b) Four seconds later, the multipath track has merged with the real target return in the arming region. The reason for the merge is that the return from target number 4753 arrived at the radar antenna at the same time as the (delayed) multipath return from target number 5601. The interference distorts the radar image of track number 4753. This distortion manifests itself as a velocity error, incorrectly indicating that target number 4753 is about to exit the arming region. (c) On the next radar scan misassociation has occurred. The erroneous velocity of target number 4753 has caused the scan-to-scan association function to misassociate track number 4753 with a multipath return that is near the position predicted by the previous incorrect velocity estimate. The real target is still in the arming region but is now incorrectly classified as a low-confidence track. This track is not recognized by the sensor-fusion function and cannot arm the THL because it is invisible to the surface monitor. The THL is therefore not illumi-

A target merge has occurred between the return from the aircraft holding in position (track number 4753) and an overlapping multipath return. This merge has distorted the target image and shifted its centroid so that the target has acquired an erroneous velocity that will take it out of the arming region. It is not yet out of the arming region, however, so the THL is still on.

On the next scan of the radar, as shown in Figure 8(c), misassociation has occurred. The identity of the aircraft holding in position has been transferred to the multipath track, which has left the arming region and is about to disappear. The THL is now off because the real aircraft, which is still in the arming region, is now a new, low-confidence track and thus is invisible to the surface monitor.

Finally, three seconds later, as shown in Figure 8(d), the misassociated multipath track is gone, the aircraft holding in position remains a low-confidence track and hence is invisible to the surface monitor, the THL remains disarmed and is therefore off, and the arrival on Runway 4R (track 2865) is about to land. The THL should have been illuminated at this point to indicate to the pilot holding in position on Runway 9 that the runway is still unsafe for takeoff.

#### *Example 2: Potential REL Interference*

When an aircraft on approach appears able to land on either of two runways, its approach-path projection becomes ambiguous; that is, the surface monitor projects the aircraft to both runways, as described earlier in the section on the prediction engine. For purposes of activating runway-status lights, the surface monitor then assumes that the aircraft will land on the primary runway, which is the main landing runway in the prevailing runway configuration. Occasionally, however, the controller clears the aircraft to land on the other runway. The result can be erroneous illumination of the RELs, which temporarily turn on along the projected landing runway rather than on the intended and actual landing runway.

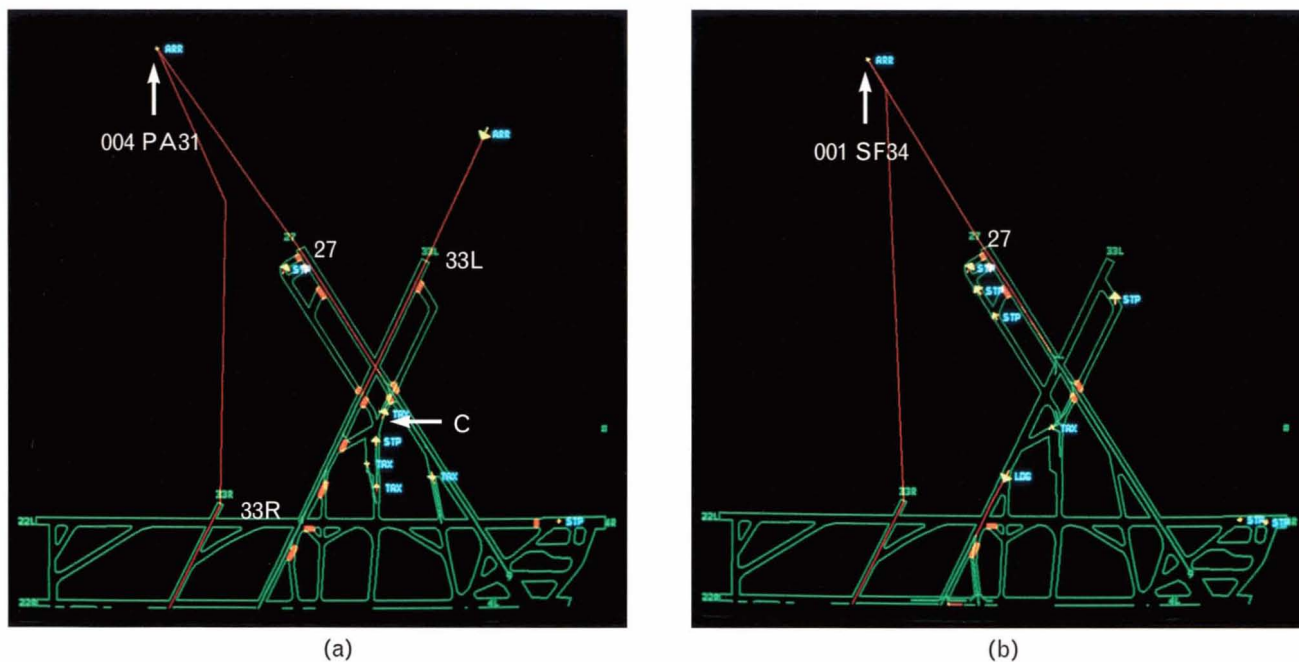
Figure 9(a) shows an example of such a situation, which occurred in the 33/27 runway configuration. A Piper Navajo (PA31) on approach to Runway 33R was lined up with Runway 27. Although projected by the surface monitor to both runways, as shown by the red path-prediction trees from the prediction engine, the aircraft was assumed to be on approach to Runway 27, which is the primary landing runway in this configuration. The RELs illuminated along the first half of that runway until the aircraft turned toward Runway 33R and the correct projection was established. An aircraft was holding short of Runway 27 during the time of REL illumination, and the potential for interference existed, although that aircraft did not move in this particular case. Another aircraft was approaching Runway 27 at Taxiway C during the brief period of REL illumination at that location. The lights turned off just as this second aircraft crossed the taxi-hold line. Had there been runway-status lights on the airfield, the pilot might have delayed his crossing and requested clarification from the local controller regarding the momentary illumination of the RELs.

The difficulty of determining the correct landing runway can be seen by comparing this projection with that of Figure 9(b). This aircraft, a Saab/Fairchild 340 (SF34), was the next one to land, two minutes after the Piper Navajo. In this case, the actual landing runway was indeed Runway 27. The two projections are almost identical; the differences do not become apparent until later in the approaches.

The light-control logic must be operated by surveillance information alone whenever possible. Controller input should be used only when absolutely necessary, in order not to reduce the effectiveness of the runway-status light system as an automatic and independent safety feature. In the case of ambiguously projected targets, such as the example just described, the ambiguity might be resolvable by making use of speed, altitude, and aircraft type, all of which

---

nated in front of the crew holding for takeoff. (d) Three seconds later, the Runway 4R arrival (track number 2865) is about to land. The THL should have been illuminated in front of the aircraft holding in position on Runway 9 to indicate that the runway is not safe for takeoff. This aircraft, however, is still invisible to the light-control logic and is thus unable to arm the THL. The multipath track that caused the misassociation has disappeared. Improved surveillance-processing track-maintenance features will reduce the incidence of track drops caused by target merges and misassociation.



**FIGURE 9.** Ambiguous projection of landing aircraft. (a) A Piper Navajo (PA31) is on approach to Runway 33R but projected to both Runways 33R and 27. The surface monitor chooses the primary runway based on the traffic flow in the current runway configuration and illuminates the runway-status lights accordingly. In this case the primary runway was Runway 27; hence the RELs illuminated along that runway. The correct, unambiguous projection was established on the next scan of the radar and the RELs along Runway 27 turned off. A little later, as the aircraft's hot zone reached Runway 33R, the RELs along that runway illuminated normally. (b) This approach was flown two minutes after the approach shown in part a. Again there is an ambiguous projection, but this time the aircraft, a Saab/Fairchild 340 (SF34), did indeed land on Runway 27, so the primary projection was correct.

are available to the surface monitor from the ARTS. Altitude information is expected to be especially useful, since the threshold of Runway 33R is considerably farther away than that of Runway 27. An arrival to Runway 33R would therefore tend to be at a higher altitude at a given distance from the Runway 27 threshold than an arrival to Runway 27. This fact is indeed reflected in Figures 9(a) and 9(b); the 33R arrival in Figure 9(a) shows an altitude of 400 feet (the entry 004 in the data tag) and the 27 arrival in Figure 9(b) shows 100 feet. Aircraft type (also indicated in the data tag) would also sometimes be helpful in determining the landing runway: whereas the PA31 in Figure 9(a) could have landed on either runway, the SF34 in Figure 9(b) could have landed only on Runway 27 because Runway 33R is too short. Neither speed, altitude, nor aircraft type information is currently used by the surface monitor in this context.

The erroneous illumination of the Runway 27 RELs in Figure 9(a) constitutes a case of (potential)

interference and not a false alarm, because the lights were turned on by real traffic. The error did not cause a corresponding missed detection on the other runway in this case, because the distance to Runway 33R is such that the correct projection was established prior to the point at which the arrival activated the 33R RELs. Thus the RELs along Runway 33R illuminated correctly.

### *Example 3: THL Interference*


To ensure free movement of traffic on the airport surface, the runway-status lights must not interfere with the traffic flow under normal conditions. This requirement is a particular concern under heavy traffic conditions. The surface monitor is implemented to ensure a smooth interplay between runway-status lights and surface traffic. The parameters are location specific and individually adjustable, and the various surface regions are also adjustable. Likewise, the light-control logic is designed to minimize interference. In-

interference does occasionally occur, however, indicating the need for further tuning of the surface monitor as well as algorithmic enhancements.

Figure 10 shows an example of interference by the THLs. This interference occurred under heavy traffic conditions, with simultaneous arrival and departure operations being conducted on both Runway 33L and the intersecting Runway 27. An aircraft (a DC-9, with track number 13505) had just landed on Runway 33L. As it was rolling out and preparing to vacate the runway, another aircraft (track 15948) moved into position behind it for departure at Taxiway G. A third aircraft was on final approach to Runway 33L. Figure 10(a) shows the situation as the departure lined up with the runway centerline on Runway 33L at Taxiway G. The previous arrival was preparing to exit the runway at Taxiway T. It was still on the runway and the THL was therefore still illuminated in

front of the departure at Taxiway G. The controller, who was aware the first aircraft was exiting the runway, cleared the departure for takeoff. The light-control logic had not yet recognized that the runway was safe for takeoff because the first aircraft was still within the activation region for the THL. The result was the situation depicted in Figure 10(b), where the departing aircraft has begun its takeoff with the THL still illuminated. At this time, the RELs are illuminated along Runway 33L, behind the departure because of the next arrival to the runway, and in front of the departure because of the departing aircraft itself.

Had lights actually been installed on the airport surface in this situation, rather than being simulated in an off-line demonstration, the result would most likely have been that the departing crew would have delayed the beginning of their takeoff until the THL turned off. This response might have forced the next



**FIGURE 10.** THL interference. (a) An aircraft (track number 13505) has landed on Runway 33L and is preparing to exit the runway at Taxiway T in the lower portion of the frame. Another aircraft (track number 15948) is lining up with the runway centerline for an intersection departure on Runway 33L at Taxiway G at the upper edge of the frame. A third aircraft, not visible on the screen, is on final approach to Runway 33L. The controller has verified that the first aircraft will vacate the runway and is preparing to issue takeoff clearance to the second aircraft. The THL on Runway 33L at G is illuminated because the prior arrival is still in the light's activation region. (b) The situation nine seconds later. The prior arrival (track number 13505) has exited the runway and is almost out of the THL activation region. The intersection departure (track number 15948) is on its departure roll. It has attained a speed and acceleration sufficient to be classified a departure by the target-state machine and has illuminated the RELs along the runway ahead. The RELs behind it are illuminated by the next arrival, now on close final to Runway 33L. The THL is still illuminated because the activation region is still occupied. This situation would have resulted in interference if runway-status lights had been installed on the airfield. Such interference can be reduced by judicious narrowing of the THL activation region and by the addition of an anticipated-runway-clear feature in the light-control logic.

arrival to Runway 33L to execute a missed approach. Two modifications to the surface monitor will greatly reduce this type of interference. The first involves a narrowing of the THL activation region at the normal runway exit points so that the light-control logic recognizes a runway-clear condition as soon as the traffic is clear of the runway. The second involves adding an anticipated-separation feature for taxiing targets, to enable the light-control logic to anticipate that an aircraft will be exiting the runway.

### *Results*

These three examples give an indication of the variety of challenges encountered during the development of the RSLs and the solutions that were implemented or identified to arrive at a system that is capable of providing effective safety backup without interfering with the normal movement of traffic. A number of other anomalies were uncovered during the performance assessment; some were similar in nature to the examples just described and others were quite different. The overall assessment results of the RSLs demonstration system are described here.

The ten hours and twenty minutes of data used for the performance assessment contained 820 operations, nearly evenly divided between arrivals and departures. A total of thirty missed detections were found in the data, as well as fourteen false alarms and fifteen cases of interference. Ten of the missed detections and eleven of the false alarms were brief anomalies, lasting less than four seconds (one or two scans of the radar antenna). We do not know at present how pilots will respond to such short anomalies; they are therefore discussed separately.

Twenty missed detections and three false alarms lasted more than two scans. They are shown in Figure 11, along with the fifteen cases of interference. The distribution of anomalies is shown separately for the four runway configurations and presented in terms of anomalies per hundred operations. Most of the THL anomalies are missed detections and most of the REL anomalies are interference. This conclusion is valid regardless of runway configuration.

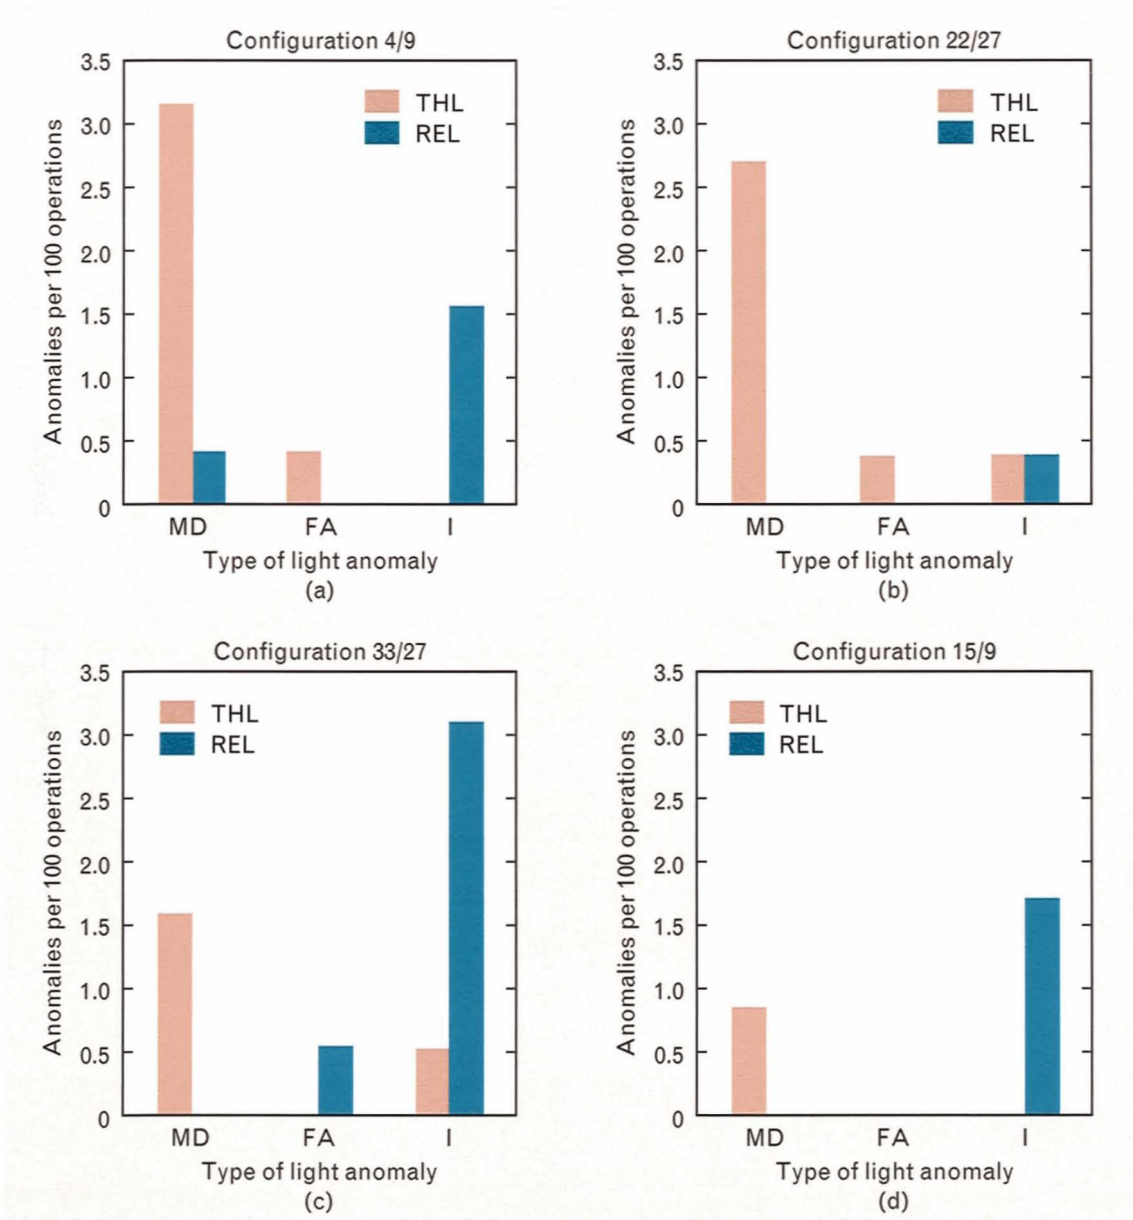
All but one of the one-and-two-scan anomalies occurred in the 4/9 and 22/27 runway configurations. Most of these short anomalies affected the THLs.

Several of these anomalies were caused by multipath or caused by false tracks due to wave action in the harbor channel; the prevalence of short anomalies in these two configurations therefore possibly indicates specific wind and ramp-congestion conditions in the data segments rather than configuration-dependent factors per se.

The four graphs of Figure 11 can be combined with weights of 0.3, 0.3, 0.3, and (for the 15/9 configuration) 0.1 to yield a prediction of the long-term average performance of the RSLs. These weights correspond to the approximate historical frequency of the four runway configurations. Figure 12 shows the results of this weighted combination. The RSLs clearly performs quite well, considering its status as a technology-demonstration system. Averaged over the long term, there is better than a 97% chance that a cockpit crew will not encounter a missed detection of more than two-scan duration during the time they are taxiing in after landing or taxiing out for departure. There is better than a 99% chance that they will not encounter a false alarm of more than two-scan duration. There is about a 98% chance that they will not experience even a mild case of interference.

Most THL missed detections result from one of the following three causes: (1) a track drop while an aircraft is holding in position on the runway, (2) a track drop while an aircraft is taxiing across a runway, and (3) a track drop while an aircraft is on final approach to a runway that intersects the currently active departure runway. A track drop that occurs while holding in position on the runway disarms the THL and prevents it from turning on. The resulting missed detections are sometimes of long duration (ten events averaged twelve scans each). The major trouble spots are the approach ends of Runways 9 and 22R; there are also track drops on Runway 4L at Taxiway S. A track drop that occurs while taxiing across a runway results in failure to activate the THL, so that it does not turn on if armed. These missed detections tend to be of moderate duration (four events averaged seven scans each). The major trouble spots seem to be the intersections of Runway 4L/22R with Taxiways S and E and the intersection of Runway 33L/15R with Taxiway N. These locations are rich in multipath returns. A track drop that occurs while on final ap-

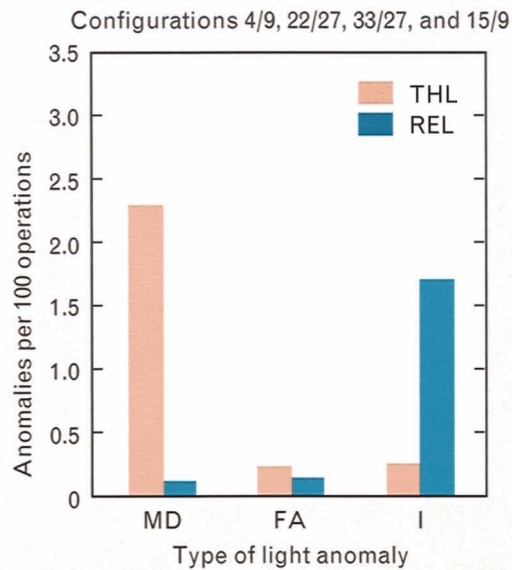




**FIGURE 11.** Anomaly statistics for the RLS in its April 1993 state of development, shown separately for the four major runway configurations at Logan Airport. The graphs show observed anomalous light operations per hundred operations; one hundred operations correspond to approximately one standard hour of operations at Logan Airport. In order to reflect accurately the impact of the light anomalies on personnel in the runway environment, only anomalies that would have affected cockpit crews or vehicle operators are counted. Brief missed detections and false alarms lasting one or two scans of the radar (less than four seconds) are excluded. In each configuration, the major THL anomaly is a missed detection and the major REL anomaly is interference. The results are based on a total of 820 operations, carefully chosen to give a well-balanced depiction of operations at Logan Airport.

proach to a runway that intersects the currently active departure runway sometimes deactivates or prevents activation of an armed THL on the intersecting runway, causing the light to turn off early or, at times,

come on late. The resulting missed detections are usually brief and may be of little direct operational consequence (eleven events averaged two scans each). The major trouble spots are Runways 4R and 33L.



**FIGURE 12.** The expected long-term average performance of the Logan Airport RSLs in its April 1993 state of development. This performance graph is produced by combining the individual graphs of Figure 11 with weights 0.3, 0.3, 0.3, and 0.1 (for the 15/9 configuration). The weights are approximately equal to the fractional usage of the four configurations over a period of one year. As with Figure 11, the underlying database is small but representative of Logan Airport operations.

Both of these runways have approach-light piers, which present a difficult high-clutter surveillance environment.

THL false alarms are generally caused by multipath on the runway; they are usually of short duration (of nine such events, one lasted ten scans, the other eight averaged 1.5 scans each). Both THL missed detections and false alarms mainly reflect shortcomings in surveillance processing and tracking, and, to a lesser extent, sensor fusion. High-leverage areas of algorithmic improvement include clutter rejection, image processing, scan-to-scan association, and multipath rejection. But the fixes should involve the surface monitor as well, because a more robust surface monitor will be less sensitive to surveillance and tracking imperfections.

Most REL interference results from one of the following three causes: (1) an aircraft rolling out after landing illuminates the RELs too far ahead and causes interference with crossing traffic; (2) an accelerating,

high-speed surface vehicle on the runway is mistaken by the safety logic for an aircraft on takeoff roll, thus illuminating the RELs and causing interference with crossing traffic; and (3) arrival traffic is projected to the wrong runway, causing erroneous REL activation. This third phenomenon occurs mostly on Runway 27, in the 33/27 configuration, as illustrated in Example 2 of the preceding section. The observed cases of interference were all caused by the light-control logic. The majority can be corrected either with parameter tuning or with algorithmic modifications.

### Conclusions

Runway incursions are a persistent problem in airport ground-movement operations. Numerous critical conflicts and several fatal accidents have occurred as a result of unauthorized or otherwise inappropriate entry of aircraft or surface vehicles onto an active runway. Many of these conflicts developed quickly, leaving little time for effective intervention by either the controller or the pilots involved. A reliable system of automatic runway-status lights would be an effective way to prevent such time-critical incursions. The Runway Status Light System (RSLs) at Boston's Logan International Airport is an off-line proof-of-concept technology demonstration system designed to show that automatically operated runway-status lights can promptly and reliably transmit runway-status information to pilots and surface-vehicle operators, thereby preventing unsafe runway entry or unsafe takeoff.

As part of the development of the RSLs, we collected real traffic data from Logan Airport and evaluated the system's real-time performance in a variety of traffic situations. This article describes the results of that performance assessment. The main objective of the assessment was to provide preliminary quantitative data on the end-to-end performance of the RSLs, as it would be experienced by the system's users, namely, tower controllers, pilots, and surface-vehicle drivers operating in the vicinity of the runways. A secondary objective was to identify fixes for the observed anomalies. The results are given in terms of light anomalies per hundred operations, and the anomalies are classified as either a missed detection, false alarm, or interference. A description of the surface monitor,

which is the component of the RSLs that receives track reports and sends light commands to illuminate the runway-status lights, was included in this article to enable the reader to understand the assessment results more clearly.

The performance assessment is based on approximately ten hours of live Logan Airport traffic data, recorded during March and April of 1993, and comprising more than eight hundred operations. The data provide a brief but balanced glimpse of operations at Logan Airport. Most hours of the operational day are represented, as are all major runway configurations and a variety of weather conditions.

The RSLs is a technology-demonstration system, and as such it performed well. The system did not employ runway-status lights on the airport surface, but, had such lights been installed, a cockpit crew could have expected to encounter a runway-status light in an incorrect state for more than four seconds only once in thirty-six operations, and an equal number of short anomalies lasting less than four seconds. Interference would have been experienced only once in fifty operations. Much of the interference would have been mild.

The RSLs generally responded correctly to airport traffic, as evidenced by the low missed-detection and false-alarm rates. The light-control logic also proved to be well tuned, resulting in a smooth interplay between the simulated lights and the real airport traffic. Overall, the demonstration system experienced an average of approximately five observed light anomalies per hour. Most of these anomalies can be eliminated with further system tuning and relatively low-risk software enhancements. Interference in particular is generally well understood and probably can be reduced by an order of magnitude.

None of the anomalies seen in this assessment represents a direct safety hazard. One of the fundamental design constraints of the safety logic was that it should not illuminate a runway-status light when doing so might directly influence operation of an airborne aircraft or an aircraft rolling at high speed on takeoff or landing.

Judging from the relatively narrow perspective of an off-line technology demonstration, interference at the level observed in the assessment is unlikely to have

any measurable negative impact on airport capacity. Of course, the validity of this assertion can be determined with certainty only during an actual field test, because the issue of interference raises a number of human-factors questions that could not be addressed in an off-line system.

All anomalies, of whatever type and regardless of duration, are undesirable. They can be reduced or eliminated by implementing the improvements and enhancements identified during the assessment. This work must be completed before the RSLs is used to drive actual status lights on the airfield in an evaluation of operational suitability.

### Acknowledgments

The authors would like to acknowledge the work of James Eggert, who developed the sensor-fusion algorithms and contributed to the development of the surface monitor; William Harman, Hayden Schultz, and Daniel Wyschogrod, who developed the surveillance-processing algorithms; Walter Brown, Charles Crone, and Tom Petrillo, who provided the expertise in air traffic control operations and procedures; Peter Daly, who worked on the radar interface and data recording system; Teresa Hall, who implemented the sensor-fusion algorithms; Steve Thompson, who developed operational concepts and performed conflict analysis; and Richard Bush, Bud Lyon, Al Reich, and Jerry Welch, who provided technical and management support. This work was supported by the Federal Aviation Administration.

---

---

## 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 Transportation, Federal Aviation Administration, "Air Traffic Control," *FAA Handbook 7110.65F* with Changes 1-5 (21 Sept. 1989).



**HARALD WILHELMSSEN** is a staff member in the Air Traffic Automation group. He received an S.B. degree from MIT in 1968 and an M.S. degree from Northeastern University in 1981, both in electrical engineering. His research is in the area of airport runway conflicts and airport surface-traffic surveillance and automation. Before coming to Lincoln Laboratory in 1981, he was employed at GTE as a systems engineer. Harald is married, with two children, and his hobbies are horticulture and jazz history.



**MARCIA P. KASTNER** is a staff member in the Air Traffic Automation group. Her research is in the development of algorithms and data-analysis techniques for air traffic control safety systems and automation systems. She received an S.B. degree in mathematics from the University of Chicago, and M.S. and Ph.D. degrees in applied mathematics from Harvard University. She has been at Lincoln Laboratory since 1989.



**THEODORE J. MORIN** is a staff member in the Air Traffic Automation group. His research is in designing and implementing air traffic control automation systems. He received a B.A. degree in physics from Bowdoin College and an M.S. degree in computer science from the University of Massachusetts at Lowell.



**JAMES L. STURDY** is a former staff member in the Air Traffic Automation group. He received S.B. and S.M. degrees in aeronautics and astronautics from MIT.