Operational Evaluation of Runway Status Lights

James R. Eggert, Bradley R. Howes, Maria Picardi Kuffner, Harald Wilhelmsen, and D. Jonathan Bernays

To maintain safe separation of aircraft on the airport surface, air traffic controllers issue verbal clearances to pilots to sequence aircraft arrivals, departures, and runway crossings. Although controllers and pilots work together successfully most of the time, mistakes do occasionally happen, causing several hundred runway incursions a year—and, less frequently, near misses and collisions—in the United States. With this rate of incursions, it is imperative to have an independent warning system as a backup to the current system. Runway status lights, a system of automated, surveillance-driven stoplights, have been designed to provide this backup function. The lights are installed at runway-taxiway intersections and at departure points along the runways. They provide a clear signal to pilots crossing or departing from a runway, warning them of potential conflicts with traffic already on the runway. Existing FAA-installed radar surveillance is coupled with Lincoln Laboratory–developed algorithms to generate the light commands. To be compatible with operations at the busiest airports, the algorithms must drive the lights such that during normal operations pilots will almost never encounter a red light when it is safe to cross or depart from a runway. A minimal error rate must be maintained even in the face of inevitable imperfections in the surveillance system used to drive the safety logic. A prototype runway status light system has been designed at Lincoln Laboratory and installed at the Dallas/Fort Worth International Airport, where Laboratory personnel have worked with the FAA to complete an operational evaluation of the system, demonstrating the feasibility of runway status lights in the challenging, complex environment of one of the world's busiest airports.

The domestic aviation system in the United States is a wonderfully safe and efficient system. The unparalleled safety record exhibited by U.S. aviation is reflected in the public's attitude toward flying, where the choice of flying versus driving is usually made on the basis of cost, convenience, and speed, but rarely on risk. Such confidence, combined with the increasing affordability of air travel, has led to an explosion of demand for service, particularly at large hub airports. As the number of operations increases, the risk must be continuously reduced to keep the number of accidents to an acceptable level—zero, according to Federal Aviation Administration (FAA) Administrator Marion C. Blakey.

Technological improvements such as the airborne Traffic Alerting and Collision Avoidance System (TCAS) installed on airliners worldwide have dramatically decreased the chance of airborne collisions; similar efforts are under way to counter the hazards encountered by aircrews operating on the airport surface. Although surface operations are generally viewed as less hazardous, the combination of large relative velocities, small separations, and relative lack of maneuverability makes runway incursions or collisions very serious events. In fact, the world's worst aviation accident was the 1977 collision between two Boeing 747 aircraft on Runway 30 of the Los Rodeos airport in Tenerife, Canary Islands. Figure 1 illustrates this accident, which resulted in the loss of 583 lives.

Significant runway conflict accidents occur only about once every two years in the United States, which makes prevention analysis based solely on those events very difficult. However, for every accident there are many more runway incursions*, which resemble accidents closely enough in cause and timeline to provide useful information on common problems and accident risks. Recognition that investigation and prevention of runway incursions is useful in preventing accidents has motivated the National Transportation Safety Board (NTSB) to designate "Stop Runway Incursions/Ground Collisions of Aircraft" as one of their "Most Wanted" aviation safety improvements [1]. In response, the FAA has made the prevention of runway incursions and accidents a priority in their program organization, assigning significant resources to identifying airports with high rates of incursions, developing improved surface markings and procedures, and developing improved situational awareness tools for air traffic controllers and pilots. Indeed, pilot awareness is the key to safe operations at towered airports, according to the FAA's Runway Safety Program Office [2]. In particular, the FAA has tasked Lincoln Laboratory to develop and test a system of automated status lights, located at runway-taxiway intersections and runway departure positions, which will provide direct indication to pilots when the runway ahead of them is unsafe for use because of conflicting traffic. The design and testing of this new runway status lights (RWSL) safety system is the subject of this article.

Runway incursions are often time-critical events; the transformation from a normal, safe operation to an imminent hazard may occur within a few seconds and involve movement of less than a few tens of meters by one or both aircraft. For instance, as illustrated in Figure 2, an aircraft that crosses a runway without clearance is indistinguishable from a normal safe operation until it becomes apparent from its taxi speed and direction of motion that crossing the runway hold line is inevitable. There is minimal time remaining to alert the pilots in either the taxiing aircraft or the high-speed aircraft already on the runway. In this case the only practical strategy for accident prevention is to provide direct indication to the taxiing pilots that the runway ahead of them is unsafe for entry, allowing them to stop short of entering the runway.

The design of a system of automated runway status lights had its origins at Lincoln Laboratory in the

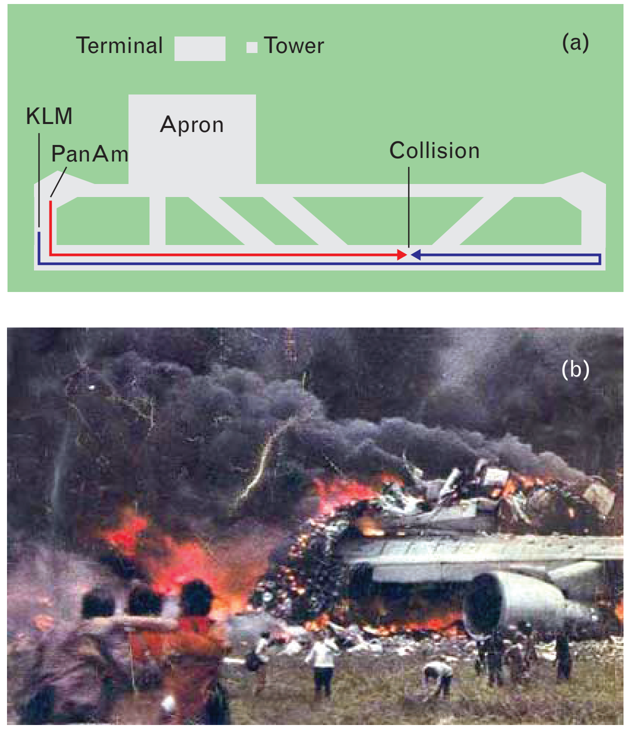

FIGURE 1. (a) Airport plan of Los Rodeos Airport, Tenerife, Canary Islands, detailing the movement of the two planes before the accident on 27 March 1977. (b) Survivors and fire after the Tenerife accident. Miscommunications between crews and the air traffic control (ATC) tower led to a collision between two heavily loaded Boeing 747 airliners. Not realizing that the runway was blocked by a taxiing aircraft (Pan Am Flight 1736) on Runway 30, the crew of KLM Flight 4805 began their takeoff roll after believing that ATC had issued a takeoff clearance. Visibility was so poor that neither crew saw the other until a collision was unavoidable. In the collision and resulting fire, 583 people perished. (Photograph courtesy of www.1001crash.com.)

* The FAA defines a runway incursion as "any occurrence at an airport involving an aircraft, vehicle, person or object on the ground that creates a collision hazard or results in a loss of separation with an aircraft taking off, intending to take off, landing, or intending to land."

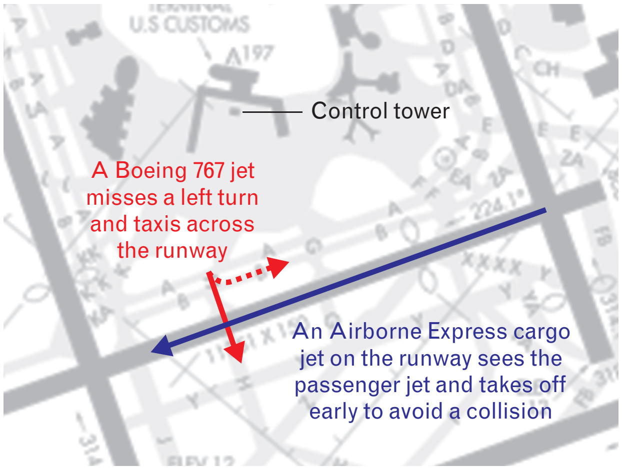

FIGURE 2. A representative runway incursion at New York's JFK International Airport. On 6 July 2005, an Israir Boeing 767 taxied across an active runway instead of turning to follow the taxiway (dashed line), and blundered directly in front of an Airborne Express DC-8 that was taking off. The Airborne Express aircraft lifted off early and cleared the Boeing 767 by less than a hundred feet. Poor visibility in rain and fog prevented either aircraft from being visible from the air traffic control tower. (JFK airport graphic courtesy of the FAA.)

1980s as an outgrowth of a comprehensive study of techniques to improve safety on the airport surface [3, 4]. Figure 3 shows a timeline of the important milestones in the development of the runway status light system. A Lincoln Laboratory team developed and implemented (for off-line demonstration) a system of automated, surveillance-driven status lights with safety logic adapted to Boston's Logan International Airport [5–7]. Although promising, actual operational use of the system was deferred in response to excessive false alerting rates; the false alerts were largely the result of merged images and time-varying multipath in the surface surveillance radar providing the aircraft position information.

In light of these results, the FAA and Lincoln Laboratory redirected their efforts to improving the quality of the surface surveillance, largely through the incorporation of a multilateration system, which uses multiple receivers and differential time-of-arrival processing to estimate the position of aircraft on the basis of signals transmitted by their air traffic control beacon transponders. The sidebar entitled "Airport Surveillance in the Past Decade" summarizes this work in more detail. This system, while subject to some errors resulting from multipath and reflection corruption of the range measurements, is largely immune to false targets caused by radar clutter, and to merged and shadowed targets. In 2000 the FAA issued a contract for twenty-six improved surface surveillance systems to be installed at high-traffic airports throughout the United States.

Meanwhile, NASA completed parallel simulation studies, establishing that pilots could understand runway status lights and would find them acceptable [8]. The validation of the utility of an automated surface traffic advisory system, combined with the expected availability of reliable and comprehensive surface surveillance sensors, motivated the FAA and Lincoln Laboratory to restart the development of a status light system in 2001.

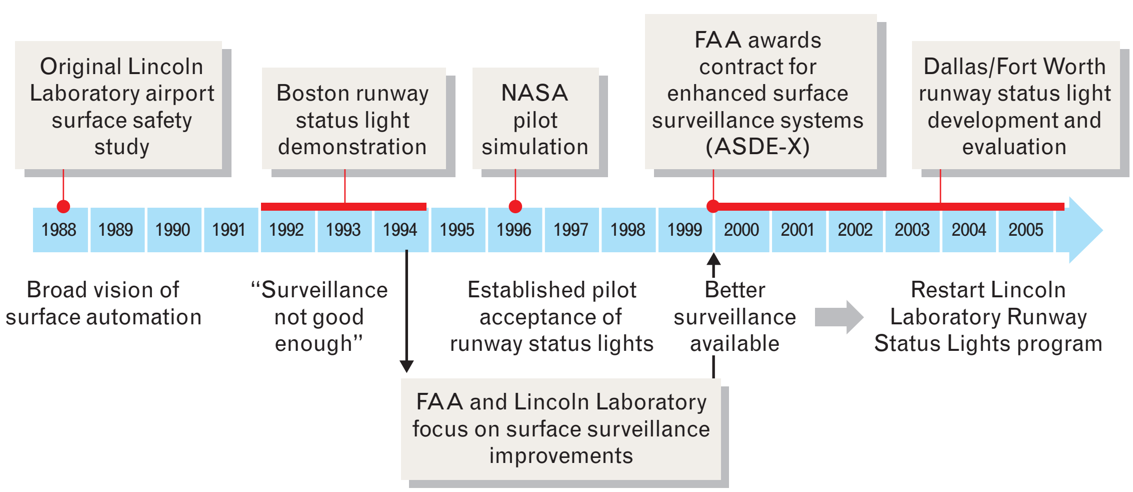

FIGURE 3. Major milestones in the development of a workable runway-status-light pilot-alerting system.

AIRPORT SURVEILLANCE IN THE LAST DECADE

Air traffic controllers direct aircraft and vehicles on the airport surface to assure that safe separation is maintained between all vehicles and aircraft. Controllers depend on a combination of air/ground communications and a diligent visual scan from an air traffic control tower to maintain situational awareness of surface movements. However, even in good weather it can be difficult for a controller to know the exact surface location of an aircraft, particularly when the crew is uncertain of their position. The identification problem becomes increasingly difficult in bad weather, when visibility is poor. Airport Surface Detection Equipment (ASDE) radars (Ku-band and X-band skin track radars) were developed to complement controllers' vision in bad weather, but they do not provide identification information.

For decades, engineers sought to augment the controllers' situational awareness with a digital display that included aircraft identification as well as position on the airport surface. Most concluded that to make the system economically viable it should use existing aircraft transponders and time-difference-of-arrival multilateration. Bendix Corporation demonstrated the technical viability of this approach in a 1974 demonstration at Logan International Airport in Boston, where positional accuracies of fifteen feet were reported [1]. The aircraft transponders in use at the time were the Air Traffic Control Radar Beacon System (ATCRBS) Mode A/C-only systems such that all aircraft within an interrogation beam replied and many replies were garbled at the receiver. Bendix overcame this problem by using large antennas on the periphery of the airport to selectively interrogate by geographic area. The resulting system demonstrated multilateration using signals from beacon transponders, but it was not economically viable.

A better system of beacon-based, multilateration surveillance grew out of the Traffic Alerting and Collision Avoidance (TCAS) system, which was developed and installed on passenger airliners in the mid-1980s. TCAS required that airliners utilize an improved transponder, Mode S, which was configured to respond to addressed interrogations and whose replies were all tagged with a transponder-specific identification code. Thus it became much easier to sort which replies detected at various multilateration receivers were associated with the same transponder. Additionally, a technique (called whisper-shout) was developed by using sequential interrogations of increasing amplitude to selectively interrogate the older, all-call ATCRBS transponders by taking advantage of differences in installed sensitivity, antenna pattern, and path loss, thus minimizing the number of garbled replies from the now less common ATCRBS transponders.

In the early 1990s, researchers at Lincoln Laboratory concluded that these TCAS-derived improvements were sufficiently mature, and a practical, affordable beacon multilateration system should be possible. As part of an FAA program, the Laboratory contracted with Cardion Corporation in 1994 to develop a system to be evaluated at Atlanta's Hartsfield airport. The demonstration was successful. Sensis Corporation subsequently acquired the Cardion technology, which served as the basis for an integrated surface surveillance system (ASDE-X) that fused reports from an X-band primary radar and a beacon multilateration subsystem, depicted in Figure A. The FAA has contracted with Sensis to deploy ASDE-X at twenty-six airports. An extension to the program to integrate the multilateration subsystem with existing, skin track-only ASDE-3 radars at high-density airports is also planned by the FAA.

A major challenge in using the ASDE-X system as the basis for

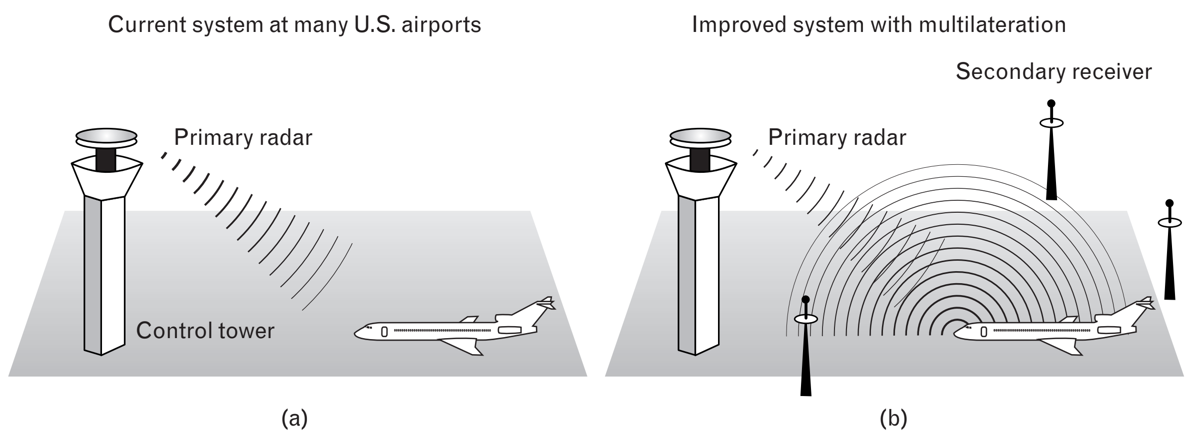

FIGURE A. Airport radar systems help air controllers prevent ground collisions. (a) A primary surface surveillance radar mounted on top of the control tower transmits a short pulse of energy and detects the energy reflected by an airplane. This tracking system is currently in operational use at thirty-four airports in the United States, including Dallas/Fort Worth International Airport, but it is not the most advanced. Position (but not altitude) is estimated from the direction the radar antenna is pointing and the round-trip time of flight of the pulse. The returns can be corrupted by multiple reflections from nearby objects, by clutter from rain, or by shadowing from buildings or other aircraft. For example, snowbanks or rain puddles can be mistaken for aircraft, and the collision warning software can produce false alarms. (b) Alternatively, an improved surface surveillance system can use multilateration, which measures the time of arrival at multiple receivers to locate the point of origin of a coded radio pulse emitted by a transmitter on the aircraft. While the individual measurements of time of arrival can be corrupted by reflections off the ground or off nearby objects, receiver diversity allows cross checking to eliminate these spurious measurements, so in general the multilateration technique yields fewer false alarms. It also collects airplane identification and yields altitude estimates from the coded transponder replies. The multilateration data can also be fused with primary radar data to handle aircraft or vehicles whose transponder is off or unavailable.

driving automated status lights was in developing the fusion algorithm to minimize the duplicate, false, or missing tracks due to biases and uncorrelated errors in the various sub-elements of ASDE-X. For example, the skin track radar senses the centroid of the aggregated radar return, while the beacon multilateration system detects the transponder antenna position. The difference between these two detected positions can be as much as tens of meters for a large aircraft, and varies with aspect angle. Although the effort to find a full resolution to these challenges delayed the deployment of ASDE-X, the system has been highly successful, and systems are currently being deployed at major airports because controllers like the clutter-free color display with tags containing aircraft identification information.

The runway status light system described in this article combines the output of a prototype multilateration system (essentially an ASDE-X without the X-band radar) with reports from an existing Dallas/Fort Worth ASDE-3 surface skin track radar to build an integrated surface surveillance picture. The observed detection and tracking performance is representative, but not identical to that seen in fielded systems. The experience gained in mitigating the effects of fusion anomalies has been valuable in improving the robustness of the runway status light safety logic, and has also provided useful feedback in assisting FAA optimization of deployed ASDE-X systems.

Reference
1. A.D McComas and A.I. Sinsky, "Brassboard Model ATCRBS Based Surface Trilateration Data Acquisition Subsystem," Dept. of Trans. Systems Center No. 471-2513-999, Aug. 1974.

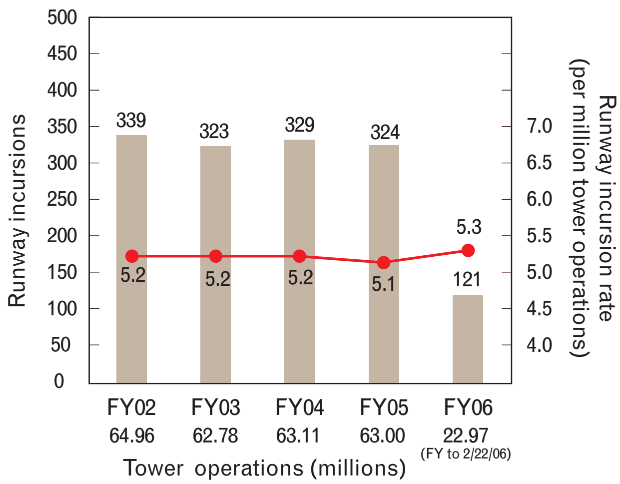

FIGURE 4. The number of runway incursions and the rates of runway incursions per one million tower operations. Runway incursions occur approximately once per day across the United States. While most of these do not represent significant safety hazards, the underlying causes are sufficiently similar to be the basis of valuable analysis to prioritize mitigation strategies.

Runway Conflict Accidents and Incursions

A significant runway conflict accident occurs in the United States approximately every two years. High-hazard runway conflicts, involving at least one major airframe and where the miss distance is a hundred feet or less, occur about twenty times per year. Runway incursions occur about once per day in the United States. Figure 4 summarizes these runway incursions, demonstrating their relatively constant rate. While the rate of incursions is enough to elicit concern, the overall numbers do not differentiate between true near misses and procedural mistakes causing a loss of separation but no risk of an actual collision. Not all runway incursions are equally hazardous; the seriousness varies according to the geometry and timing of the particular scenario.

To provide further insight, a detailed examination of runway incursions reported during a representative four-year period (1997–2000) was undertaken to classify runway conflict scenario types and estimate the potential effectiveness of pilot- and controller-alerting systems [9]. The results, summarized graphically in Figure 5, confirm that the number of truly hazardous events is a small fraction of the reported incursions. Classification of the incursions according to geometry and operational intent shows that high-hazard runway conflicts do not differ materially in geometry or operational intent from other lower-hazard incursions. Furthermore, the division of incursions into operational scenarios does not vary significantly from year to year. This implies that a system to prevent runway incursions in general should equally prevent the most dangerous incursions and should work consistently over time.

Incursion scenarios can also be grouped according to how far in advance a warning could be issued. An incursion can develop slowly enough to allow a controller to react to a warning and issue the appropriate commands to resolve the conflict. Or there can be so little time between detecting the imminent loss of separation and the subsequent collision or incursion that the only viable mitigation strategy is to provide information directly to the pilot in one or both aircraft. Approximately one-fifth of the incidents develop slowly enough to accommodate an alerting system that pro-

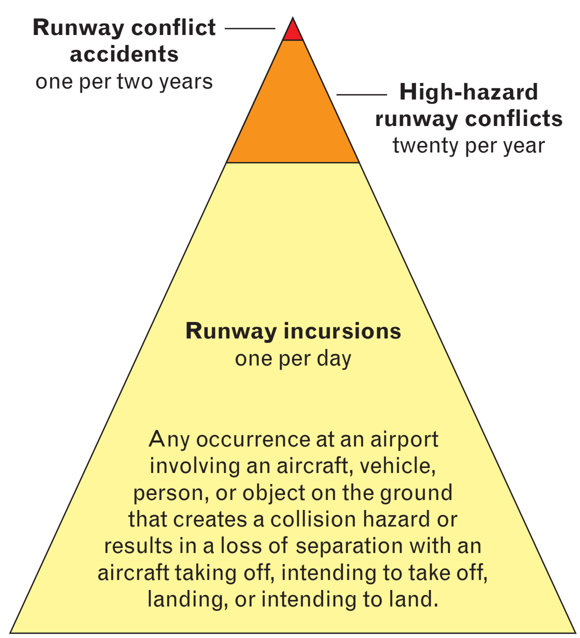

FIGURE 5. Distribution of runway conflict incidents, based on FAA-documented occurrences from 1997 to 2000. A total of 1369 runway incursions were noted at all towered airports within the United States during this period. Of these, subsequent analysis identified 167 high-hazard runway conflicts, meaning that the aircraft involved came within a hundred feet of one another at speeds high enough to assure major damage or loss of life if a collision had occurred. There were two runway conflict accidents during this time.

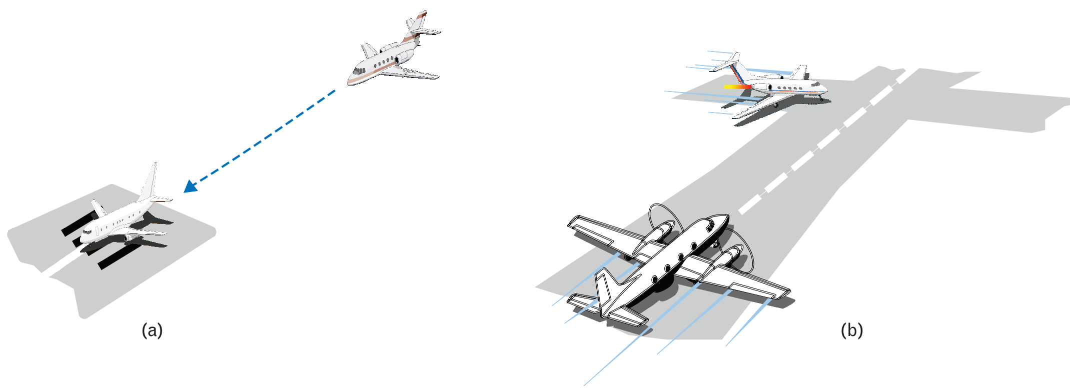

FIGURE 6. Classification of runway incursions. Analysis of 167 high-hazard runway conflicts that occurred between 1997 and 2000 revealed that about one-fifth could be classified as (a) non-time-critical, in which for instance an arriving aircraft is cleared to land on a runway already occupied by another aircraft waiting to take off. Most high-hazard runway conflicts, however, are (b) time-critical, in which for instance an aircraft taxies across a runway ahead of a departing high-speed aircraft. The best defense for non-time-critical runway conflicts is to alert the controller and allow the controller to identify and issue the necessary clearances to resolve the conflict. However, the majority of cases are time-critical, in which the best mitigation strategy is a direct warning to the pilot that the runway ahead is unsafe for crossing or for takeoff.

vides warnings to the air traffic controller, who subsequently corrects the situation through clearances to pilots, as illustrated in Figure 6. In the majority of the cases, however, there would be significantly less than thirty seconds between detection of the imminent loss of separation and a subsequent collision or incursion. In these cases only a direct warning to the pilot has any chance of preventing the incursion or reducing its severity.

Operational Concept

The movement of aircraft on the surface of a major airport is controlled through clearances issued by air traffic controllers to pilots in departing, arriving, and taxiing aircraft. The air traffic controllers have the responsibility to direct and separate the traffic so that aircraft may proceed with minimal delay while remaining clear of any collision hazard. Pilots, in accepting and complying with these clearances, share the responsibility to ensure safety, and must always make the final decision whether it is safe to cross, enter, or take off from a runway. Currently, pilots make their assessment of safety on the basis of their overall awareness of traffic operations, which includes conducting a diligent, visual scan for conflicting aircraft while listening attentively to the clearances being issued on the control frequency.

Runway status lights, illustrated in Figure 7, are a supplement to existing pilot procedures, training, and visual monitoring. The status lights are placed at runway-taxiway intersections (runway entrance lights), where they are visible to pilots about to enter or cross the runways, and on the runways at the takeoff positions (takeoff hold lights), where they are visible to pilots about to depart from a runway (and to pilots about to land, although these lights are not specifically designed to alert such pilots). The lights are controlled via processing of surface surveillance information, and they illuminate whenever the runway is occupied by traffic that would represent a hazard to other aircraft. For instance, the runway entrance lights at a particular intersection illuminate whenever there is high-speed traffic projected to pass through the intersection. Similarly, the takeoff hold lights illuminate whenever there is an aircraft or vehicle on or about to be on the runway ahead of an aircraft in position for takeoff.

The status lights operate autonomously and serve as an independent backup to the clearances issued by air traffic control. They assist pilots in verifying that proceeding to cross a runway or start their takeoff roll will not cause a collision. Status lights at each intersection or takeoff hold position turn on and off as a group, but the groups of lights are controlled independently. For instance, when an aircraft begins its take-off roll,

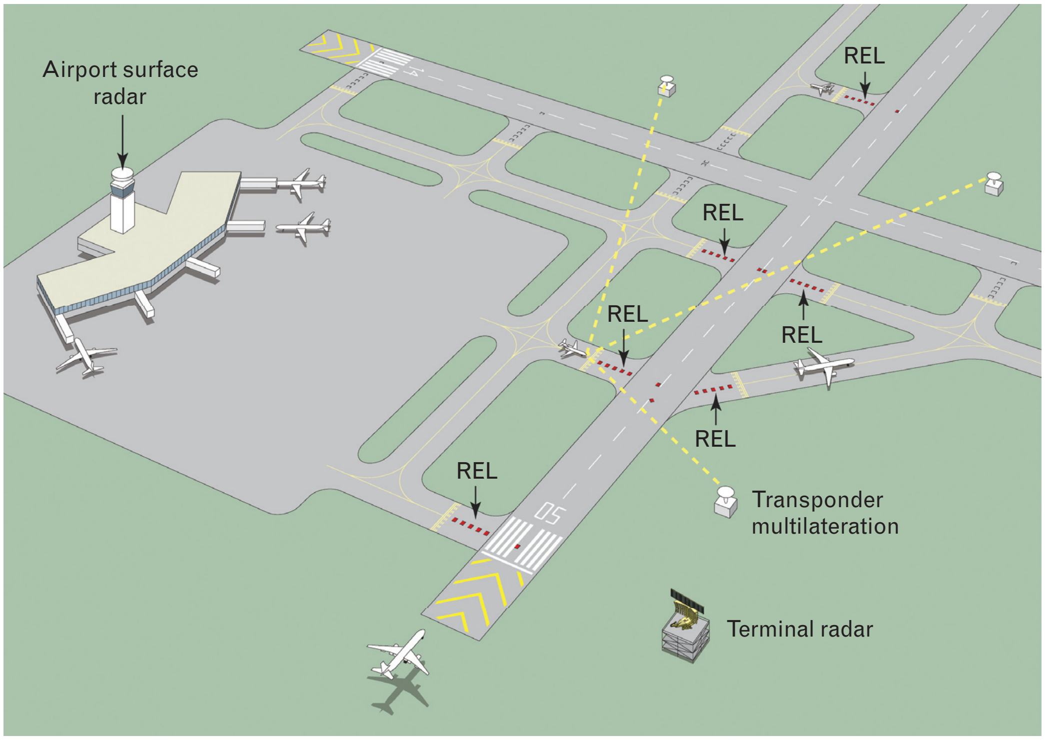

FIGURE 7. Operational concept for runway status lights. The system is comprised of a surveillance source, safety logic that estimates where each aircraft could be within the next twenty to thirty seconds and identifies the conditions under which the lights should illuminate, and a field lighting system. The surveillance source is expected to be a fused combination of terminal beacon radar for airborne targets, local surface surveillance radar to track any aircraft whose transponder is not working on the surface, and the transponder multilateration system for tracking surface aircraft whose transponders are enabled.

the runway entrance lights at all of the taxiway intersections downfield on that runway illuminate, signaling to all crossing traffic that the runway is unsafe for entry or crossing. As the departing aircraft passes by each intersection, the groups of lights at that intersection are extinguished, allowing other aircraft to taxi across the runway once the collision hazard has passed. This automated sequence is repeated for every operation with high-speed traffic on the runway, which may be hundreds of times per day at a busy airport. The challenge in implementing automated status lights is in assuring that their operation does not interfere with normal, safe operations, but still has a high probability of preventing a collision or near miss.

Runway status lights work in concert with other deployed FAA surface safety systems. Currently, the Airport Movement Area Safety System (AMASS) works to detect imminent loss of separation between surface traffic and arriving or departing aircraft. Upon detection of an imminent loss of separation, an aural alert is provided to an air traffic controller in the tower, who then issues revised clearances to resolve the conflict. The design of AMASS makes it particularly well-suited to responding to collision scenarios that develop relatively slowly. Such scenarios typically occur when one aircraft is directed to land on a runway already occupied by another aircraft that is waiting on that runway for a takeoff clearance.

A review of accident and incident reports reveals that such scenarios do occur, and have caused loss of life (e.g., the 1991 accident at Los Angeles International Airport), but are not the most frequent type of runway conflict. Rather, scenarios in which a collision hazard develops with minimal notice occur almost four times as frequently, such as when a taxiing aircraft blunders in front of an aircraft in the midst of its takeoff roll. The scenario is indistinguishable from a normal safe operation almost until the taxiing aircraft enters the runway. By the time it is apparent that the taxiing aircraft is a collision hazard there is no time available for an air traffic controller to warn the pilots in either cockpit to prevent the incursion. Under these more common scenarios, direct notification to the taxiing crew via status lights is a more effective means of prevention than involving an air traffic controller.

Runway status lights depend on both the runway entrance lights and the takeoff hold lights to maximize effectiveness. The concept of operations relies on the ability to warn one of the aircraft in the conflict scenario, but not necessarily to provide warnings to pilots in both aircraft. For instance, when a runway is "hot" (in use by high-speed departing or arriving traffic) the runway entrance lights warn pilots in taxiing aircraft not to cross or enter the runway. No indication or warning is given to the pilot in the high-speed aircraft. Conversely, in a scenario in which a pilot is cleared to take off, and conflicting traffic remains on the runway downfield, the takeoff hold lights provide warning solely to the pilot about to take off. A comprehensive review of incursion geometries revealed that the combined use of runway entrance lights and takeoff hold lights in a runway status light system would provide a warning to one or both of the affected pilots in about 65% of the cases studied. Status lights in conjunction with AMASS would address about 85% of all incursions. The remaining 15% of incursions are due to a variety of circumstances that lie outside most normal operations (e.g., aircraft piloted by inexperienced or impaired crews, willful disregard, or stolen aircraft).

Challenges in Implementing Runway Status Lights

The challenge of implementing runway status lights is in integrating the system into the fast-paced traffic sequence that is characteristic of major traffic hubs. Dallas/Fort Worth International Airport is a typical example of such a hub, with seven runways handling some 2000 arriving and departing flights per day. A single runway at Dallas/Fort Worth may handle about 450 flights per day, interspersed with some 500 runway crossings. The RWSL light activation logic, which has no access to controller-pilot communications, must infer aircraft intent from a combination of the track history, the current aircraft position and velocity, and the geometry and pattern of operations at the airport, and determine the correct timing for as many as 5000 light operations on a single runway during the course of a busy day. The large number of illuminations is a consequence of the fact that in most situations a single operation (e.g., a landing) causes all of the runway entrance lights along a runway to illuminate at least once during some part of the operation.

In planning for the operational evaluation it was estimated that one operation in 2000 encountering a false illumination would be acceptable to pilots and controllers. This rate corresponds to approximately one such error every four days for a busy runway, or an error rate on the order of 0.03% of all light illuminations. This would be difficult enough if the surveillance information available to the safety logic were perfect. Unfortunately, the surveillance includes numerous errors, dropouts, and false tracks that if not taken into account could easily cause an excessive rate of light illumination errors. In the face of these challenges, the RWSL operational evaluation was designed to answer the question of whether a practical RWSL system was possible.

System Architecture

Figure 8 illustrates the major elements of the runway status lights system. The complete RWSL surveillance picture at Dallas/Fort Worth is assembled by fusing reports from several existing FAA surveillance sources. A prototype Sensis ASDE-X (Airport Surface Detection Equipment) system, which includes a beacon multilateration subsystem, an interface to the local airport surveillance radars, and an ADS-B receiver/decoder, provides a combined air picture for surface and airborne aircraft that have operating air traffic control transponders. This picture is combined with outputs derived from a Northrop Grumman ASDE-3 surface-surveillance skin track radar to provide visibility of surface aircraft and vehicles that do not have an operational transponder. The resultant surveillance output is equivalent to a planned FAA system that will combine multilateration, ADS-B receptions, airport surveillance radar, and ASDE plots in a single automation system scheduled for installation at Dallas/Fort Worth in 2009. The surveillance tracks form the input to the Laboratory-developed runway-status-light safety logic that comprises the core of the algorithms. The light

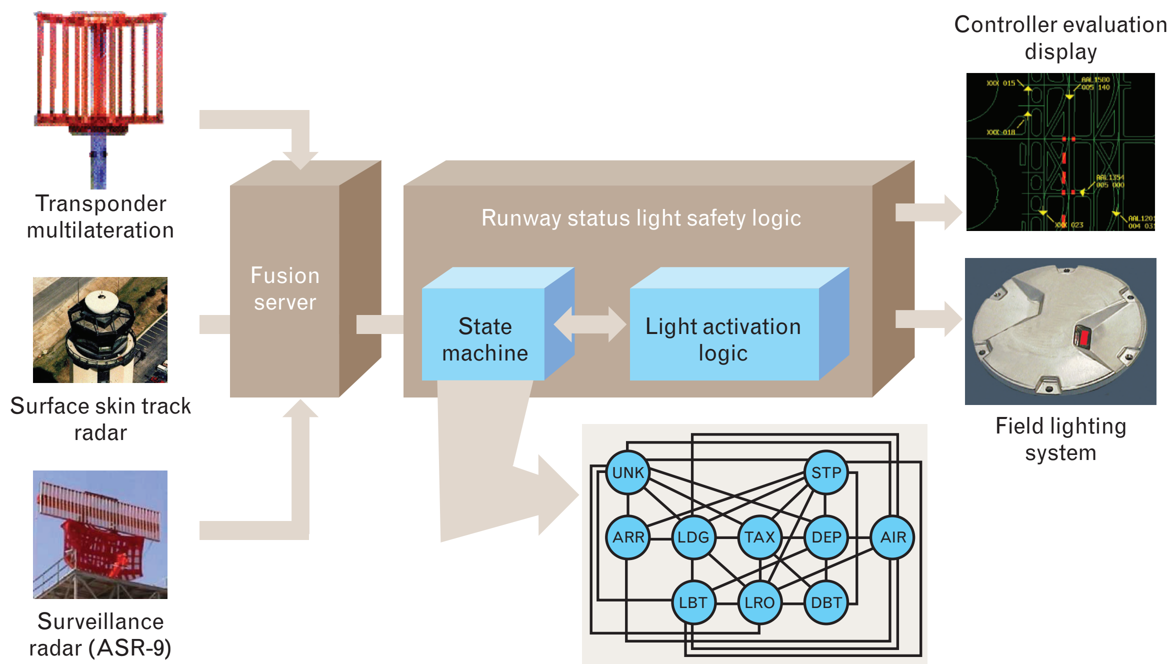

FIGURE 8. Major elements of the runway status light system, as implemented for the operational evaluation at Dallas/Fort Worth International Airport. Surveillance is derived from a combination of transponder multilateration, surface skin track radar, and terminal surveillance radar, fused together to form a single coherent picture of traffic on and near the airport surface. The operational state (taxiing, landing, departing, and so on) of the traffic is estimated by a state machine. The light activation logic determines which runway status lights should be illuminated based on current surveillance parameters and operational state. Light commands are sent to the field lighting system, and traffic and light states are shown on a controller evaluation display.

commands generated by the safety logic are passed to a Siemens-built, FAA-installed field lighting system.

The core software of the runway status light system consists of multiple components designed to segregate functionality and allow interoperability in a flexible environment. Surveillance processing modules handle the surveillance-source specific interfaces and, in the case of the ASDE-X, much of the false track rejection logic. A surveillance fusion tracker fuses the surveillance streams into one coherent picture of traffic on and near the airport surface, and applies rules that further reduce the false track rate, estimate heading for stationary targets, and coast tracks through surveillance gaps. The safety logic process accepts the fused surveillance, determines the operational state of the track (such as stopped, taxiing, landing, or departing), predicts likely future behavior based on the current state, and determines which lights should be illuminated. The field lighting system interface handles the protocol to the field lighting system.

User interface programs include a visual display and a runtime configuration manager. The visual display shows tracked aircraft and vehicles with icons showing vehicle type and tags providing call sign, equipment type, first fix, and altitude data, all presented on a map of the airfield. The runtime configuration manager allows operational air traffic personnel to select the current runway configuration, light intensity (which is also automatically switched to appropriate default settings for daytime and nighttime), and other system configuration information. Additional programs exist to play back and analyze the various recordings generated during RWSL processing and to edit adaptation data.

The RWSL software is a collection of over a hundred programs and twenty-five libraries written in C++. The files contain approximately 380,000 lines of code. The code compiles and runs on the Solaris, Mac OS X, and GNU/Linux systems, and is compatible with any POSIX UNIX system with a standardscompliant C++ compiler. The RWSL software contains over five hundred classes spread throughout the various applications and libraries.

Operational Evaluation Preparation

The operational evaluation of RWSL was divided into two sections: evaluation of runway entrance lights (REL) and evaluation of takeoff hold lights (THL). As previously mentioned, RELs and THLs operate in a complementary fashion. RELs illuminate in response to high-speed traffic on the runway, independently from traffic on the taxiways. RELs may be expected to work properly whenever the surveillance does a good job of detecting and tracking arriving aircraft about to land and high-speed traffic on the runway. In contrast, THLs illuminate only when two conditions are met: an aircraft is in the departure position and traffic is on the runway downfield, or is projected to enter the runway within the next few seconds. Therefore, correct THL operation requires that high-quality surveillance must be maintained for all traffic (including stationary and slowly moving aircraft or vehicles) on the runways and on the taxiways near runway-taxiway intersections.

Because one of the biggest uncertainties in the RWSL evaluation was the effect that surveillance errors would have on RWSL operation, the initial evaluation focused on REL operation, deferring the THL evaluation (with its more stringent surveillance requirement) until it was certain that REL operation would be acceptable. This article presents the results of the REL evaluation.

The overall preparation of the prototype system commenced with a complete engineering design phase, in which the end-to-end system design was completed and tested with recorded surveillance data. Next, the system was installed at Dallas/Fort Worth and evaluated in real time by comparing the light illuminations (as displayed on a computer monitor) with the out-the-window view from the air traffic control tower located in the center of the airfield. Finally, once the light timing had been tuned to the satisfaction of Dallas/Fort Worth air traffic controller supervisors, and verified through flight tests in collaboration with FAA personnel, the REL operational evaluation commenced.

Operational Evaluation Roles and Responsibilities

The REL operational evaluation was led by Lincoln Laboratory personnel, in collaboration with numerous partners. FAA Headquarters personnel from the Air Traffic Organization Planning Directorate (ATO-P) coordinated the overall effort as the program sponsor. Dallas/Fort Worth air traffic controllers worked directly with Laboratory personnel to assure that the REL timing was properly adjusted prior to commencing live operations. Air traffic personnel also exercised final authority for RWSL operation, and were empowered to make all decisions regarding system configuration and status whenever the system was enabled.

Dallas/Fort Worth Technical Operations personnel coordinated all field installations with the field lighting system vendor, and assisted with installation and maintenance tasks. The Dallas/Fort Worth Airport Authority coordinated all training and assessment interactions with the vehicle operators, who work for the Airport Authority.

Airline operations at Dallas/Fort Worth are dominated by American Airlines and its commuter affiliate, American Eagle, together accounting for 85% of passenger traffic and 45% of cargo traffic. The remaining operations are mostly split among a combination of United Airlines and US Airways (for passenger operations) and FedEx and UPS (for cargo operations). Laboratory personnel worked closely with training managers from all these airlines to develop and distribute RWSL training material to pilots, and to encourage pilots to report on their reaction to RWSL via online surveys. Representatives of the two major pilot unions—Allied Pilots Association (APA) and Airline Pilots Association (ALPA)—were also heavily involved in the distribution of training material.

Training

To evaluate RWSL it was essential that the user community—air traffic controllers and supervisors, pilots, and vehicle operators—understand its operation. Forty air traffic controllers and nine supervisors were trained to assess whether RWSL was working properly or not, and how to react to various anomalies that might occur. Their training was implemented through a series of short classes prepared by Laboratory personnel and

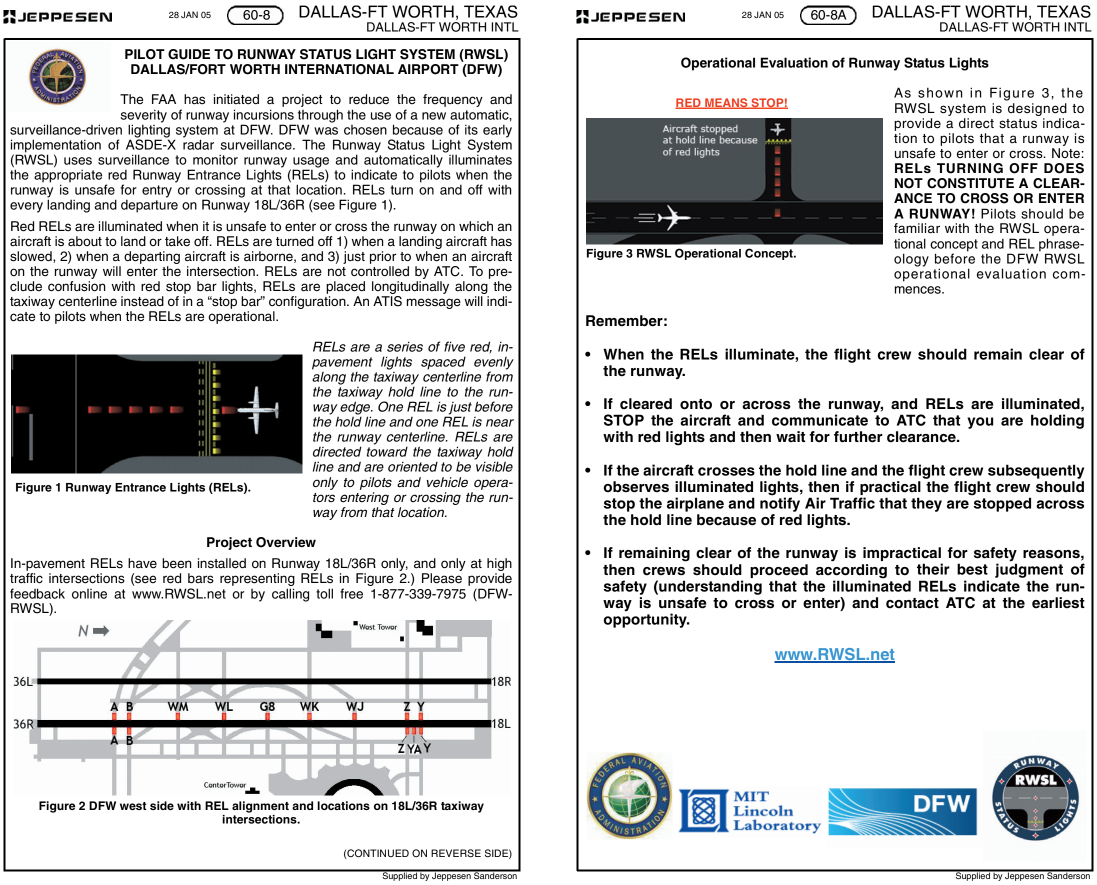
JEPPESEN 28 JAN 05 60-8 DALLAS-FT WORTH, TEXAS
DALLAS-FT WORTH INTL

PILOT GUIDE TO RUNWAY STATUS LIGHT SYSTEM (RWSL)
DALLAS/FORT WORTH INTERNATIONAL AIRPORT (DFW)

The FAA has initiated a project to reduce the frequency and severity of runway incursions through the use of a new automatic, surveillance-driven lighting system at DFW. DFW was chosen because of its early implementation of ASDE-X radar surveillance. The Runway Status Light System (RWSL) uses surveillance to monitor runway usage and automatically illuminates the appropriate red Runway Entrance Lights (RELs) to indicate to pilots when the runway is unsafe for entry or crossing at that location. RELs turn on and off with every landing and departure on Runway 18L/36R (see Figure 1).

Red RELs are illuminated when it is unsafe to enter or cross the runway on which an aircraft is about to land or take off. RELs are turned off 1) when a landing aircraft has slowed, 2) when a departing aircraft is airborne, and 3) just prior to when an aircraft on the runway will enter the intersection. RELs are not controlled by ATC. To preclude confusion with red stop bar lights, RELs are placed longitudinally along the taxiway centerline instead of in a "stop bar" configuration. An ATIS message will indicate to pilots when the RELs are operational.

Figure 1 Runway Entrance Lights (RELs).

RELs are a series of five red, in-pavement lights spaced evenly along the taxiway centerline from the taxiway hold line to the runway edge. One REL is just before the hold line and one REL is near the runway centerline. RELs are directed toward the taxiway hold line and are oriented to be visible only to pilots and vehicle operators entering or crossing the runway from that location.

Project Overview

In-pavement RELs have been installed on Runway 18L/36R only, and only at high traffic intersections (see red bars representing RELs in Figure 2.) Please provide feedback online at www.RWSL.net or by calling toll free 1-877-339-7975 (DFW-RWSL).

Figure 2 DFW west side with REL alignment and locations on 18L/36R taxiway intersections.

(CONTINUED ON REVERSE SIDE)

Supplied by Jeppesen Sanderson

JEPPESEN 28 JAN 05 60-8A DALLAS-FT WORTH, TEXAS
DALLAS-FT WORTH INTL

Operational Evaluation of Runway Status Lights

RED MEANS STOP!

Figure 3 RWSL Operational Concept.

As shown in Figure 3, the RWSL system is designed to provide a direct status indication to pilots that a runway is unsafe to enter or cross. Note: **RELs TURNING OFF DOES NOT CONSTITUTE A CLEARANCE TO CROSS OR ENTER A RUNWAY!** Pilots should be familiar with the RWSL operational concept and REL phraseology before the DFW RWSL operational evaluation commences.

Remember:

- **When the RELs illuminate, the flight crew should remain clear of the runway.**
- **If cleared onto or across the runway, and RELs are illuminated, STOP the aircraft and communicate to ATC that you are holding with red lights and then wait for further clearance.**
- **If the aircraft crosses the hold line and the flight crew subsequently observes illuminated lights, then if practical the flight crew should stop the airplane and notify Air Traffic that they are stopped across the hold line because of red lights.**
- **If remaining clear of the runway is impractical for safety reasons, then crews should proceed according to their best judgment of safety (understanding that the illuminated RELs indicate the runway is unsafe to cross or enter) and contact ATC at the earliest opportunity.**

www.RWSL.net

Supplied by Jeppesen Sanderson

FIGURE 9. Runway status light information page for Dallas/Fort Worth International Airport, supplied as an update page to the Jeppesen chart binder used by almost all commercial pilots as their comprehensive library of approach and departure procedures, and for detailed airport information.

presented by FAA personnel to the Dallas/Fort Worth controllers and supervisors.

Pilot training was more challenging because the pilot community is much larger (American Airlines alone employs more than 12,000 pilots), and direct classroom-based training for all of the pilots that could be passing through Dallas/Fort Worth was not practical. Instead, training material was disseminated to pilots through page additions (illustrated in Figure 9) to the navigational binders carried by all commercial pilots, a web site [10], and e-mail broadcast reminders sent out by each of the major pilot unions. Laboratory personnel also worked with airline training managers to integrate RWSL training elements into the six-month recurrency training every commercial pilot is required to complete. Pilots were provided a summary overview of the RWSL system, and were given specific guidance to stop upon seeing illuminated status lights, and then report to air traffic control that they were "stopped with red lights." However, because RWSL serves as an advisory system and does not change the pilot's statutory responsibility for the safe operation of the flight, the recommended procedure was advisory and not mandatory.

Finally, vehicle operators, who operate under more restrictive rules than pilots, were instructed by the Dallas/Fort Worth Airport Authority that they should stop at all illuminated status lights, and report the status of these light to air traffic control. Under no circumstances were vehicle operators to exercise discretion in proceeding through illuminated (red) status lights. The Dallas/Fort Worth Airport Authority, through the use of computer-based training, was able to document that all drivers authorized to operate in the airport surface movement area had received the basic RWSL training.

Operational Assessment

The goal of the operational evaluation was to determine whether RWSL is compatible with operations at a busy airport and is acceptable to users (pilots, vehicle operators, and air traffic control). A two-part methodology was employed: (1) technical operational assessment, which counts and classifies the light illumination anomalies, and (2) operational feedback, which uses surveys to elicit opinions from the various classes of users on the usability, acceptability, and benefit of RWSL. These are discussed in the following sections.

Technical Operational Assessment

For RWSL to function seamlessly within the normal airport routine, the operation of the runway status lights must account for the full complexity of airport operations. Ensuring that the lights work as intended calls for exposing the system to large amounts of real traffic data and a correspondingly large effort in performance assessment. For that reason, a comprehensive, automation-aided performance analysis and assessment capability was envisioned from the outset of RWSL development. Two features of RWSL make this possible. First, the software maintains a complete record of all system inputs and outputs as well as internal transactions in order to facilitate post-analysis of all events of interest. Second, a versatile playback capability allows replay of synchronized traffic data, simulated light operation, and recorded air traffic control radio transmissions, as well as visualization of normally hidden internal data and other information of interest to the analyst.

These two features do not by themselves add up to a practical assessment capability. The stringent performance requirements imposed on the RWSL system call for analysis of large amounts of data to ensure sufficient statistical confidence. Although it is in principle possible for an analyst to carry out such analysis by careful observation of the traffic and light display, this process is complex, exhausting, and extremely time consuming. Therefore, as a practical matter, playback-based assessment of large amounts of data is out of the question. An automatic assessment capability that prescreens the data for events of interest allows more efficient analysis of the requisite volume of data, either in automation-aided mode or—ideally—by a fully automatic process.

To verify that the runway status lights operate as desired at all times and under all traffic conditions at the airport in question, the performance assessment must be realistic. From the perspective of the prospective system users and beneficiaries (pilots, vehicle operators, and air traffic controllers) this means that the lights must respond appropriately to the *actual* traffic situation on the airport. In general it is not sufficient that the assessment verify that the lights respond as intended to the *representation* of the traffic that is provided as input to the RWSL safety logic by the surface surveillance and target tracking functions.

The airport surface is a difficult surveillance environment, and perfect representation of the actual traffic cannot be assumed. In terms of the end result, it makes little difference to the user whether the runway status lights respond inappropriately to a correct track or whether they respond as intended to a track that turns out to be false or otherwise incorrect. Either way, incorrect light operation may result. The assessment process must identify incorrect light operation irrespective of whether this is caused by faulty light logic or faulty depiction of the traffic. Doing this is hampered by the lack, in any extensive assessment, of direct traffic 'ground truth.' The non-causal nature of the assessment process compensates for this difficulty, by allowing events to be interpreted in the context of future developments.

The Assessment Process

The assessment should identify and characterize all instances of incorrect light operation. However, the wide range of possible anomaly causes and manifestations makes it inadvisable to attempt this assessment in a single step. Direct anomaly detection risks missing unanticipated anomalies that fail to match the preconceptions inherent in any one-step detection protocol. Therefore, the performance assessment was carried out as a series of increasingly selective extractions of generalized 'events of interest,' which were subjected at each stage to progressively refined analysis. This process was initially deliberately non-selective—that is, it was biased toward false positives in order to avoid false negatives. The former were eliminated in the analyses

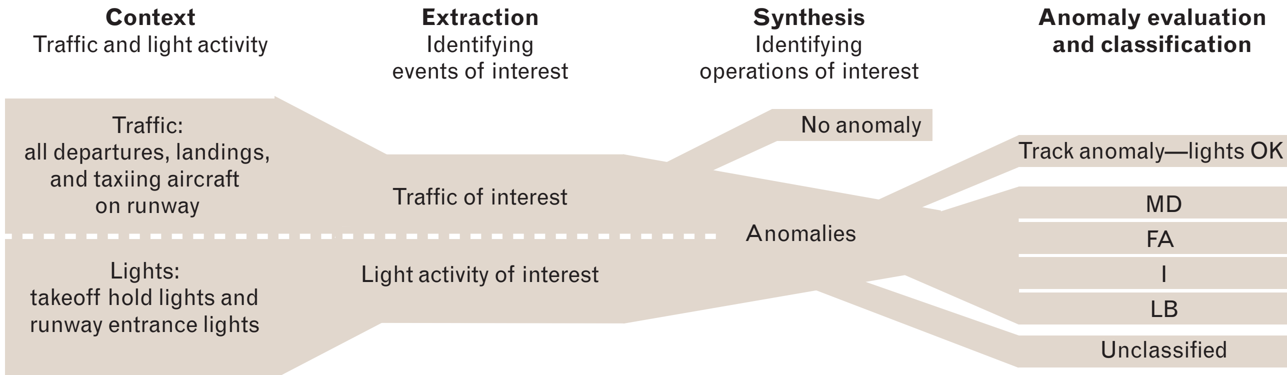

FIGURE 10. The four stages of the performance assessment process. Stage 1 summarizes traffic and runway light activity and screens for events of interest. Stage 2 extracts events of interest as potential anomalies; Stage 3 correlates the potential anomalies and places them in the operational context. Stage 4 classifies the synthesized anomalies. Three light performance metrics are used in the anomaly evaluation and classification stage: missed detections (MD), false activations (FA), and instances of interference (I). Light busts (LB), a subset of interference, are also identified during anomaly classification.

of subsequent stages; the latter would represent undetected anomalies. Figure 10 shows the four-stage assessment process schematically.

Stage 1. Context: Traffic Summaries and Light Activity

Traffic summaries place the assessment of light operation in the context of aggregate traffic and allow all operations to be screened on an individual basis to ascertain whether further analysis is warranted. All runway operations (departures, landings, and crossings) are described in terms of selected salient events and quantities. Departures are described by their initial takeoff and subsequent liftoff parameters.

In a full RWSL installation the light activity includes both REL activity and THL activity. Only the former is relevant to the 2005 Dallas/Fort Worth operational evaluation. Unlike THL activity, a detailed REL activity summary tends not to be useful to the analyst for two reasons: (1) the number of illuminations is very large and (2) what constitutes normal operation, as reflected by the illumination duration, varies widely, depending on context. The REL activity summary was therefore limited to an overall statistical description, supplemented by additional detail for illuminations judged *prima facie* of interest.

Stage 2. Extraction: Identifying Events of Interest

All tracks and illuminations were evaluated in isolation, and potential track or illumination anomalies calling for further investigation were flagged according to defined criteria. Examples of track anomalies included short duration or inappropriate start or end points, state-sequence anomalies, or speed/state incompatibility. Operational curiosities included speeds, headings, times, locations, and durations that deviate from what would be expected, as well as states that may point to operational situations of interest. Some examples of light anomalies are questionable illumination durations for the context, inappropriate illumination gaps, and light state changes coinciding with track gaps.

Stage 3. Synthesis: Identifying Operations of Interest

The events of interest were correlated with respect to track number and in some instances also with respect to time and location. A large number of anomalies strongly suggests a defect but may also signal merely an unusual operational condition or a case of incomplete surveillance information. A single anomaly is less suggestive but may nevertheless indicate a reportable defect. Anomaly classification can often be performed directly at this stage by a human analyst or by a simple classification algorithm. Not all anomalies correspond to incorrect light operation, but track anomalies are generally of interest even if they do not affect light operation negatively.

Light anomalies can be sorted at this stage into observed and unobserved anomalies. The distinction was made because only observed anomalies affect traffic. The fraction of anomalies that is observed cannot be

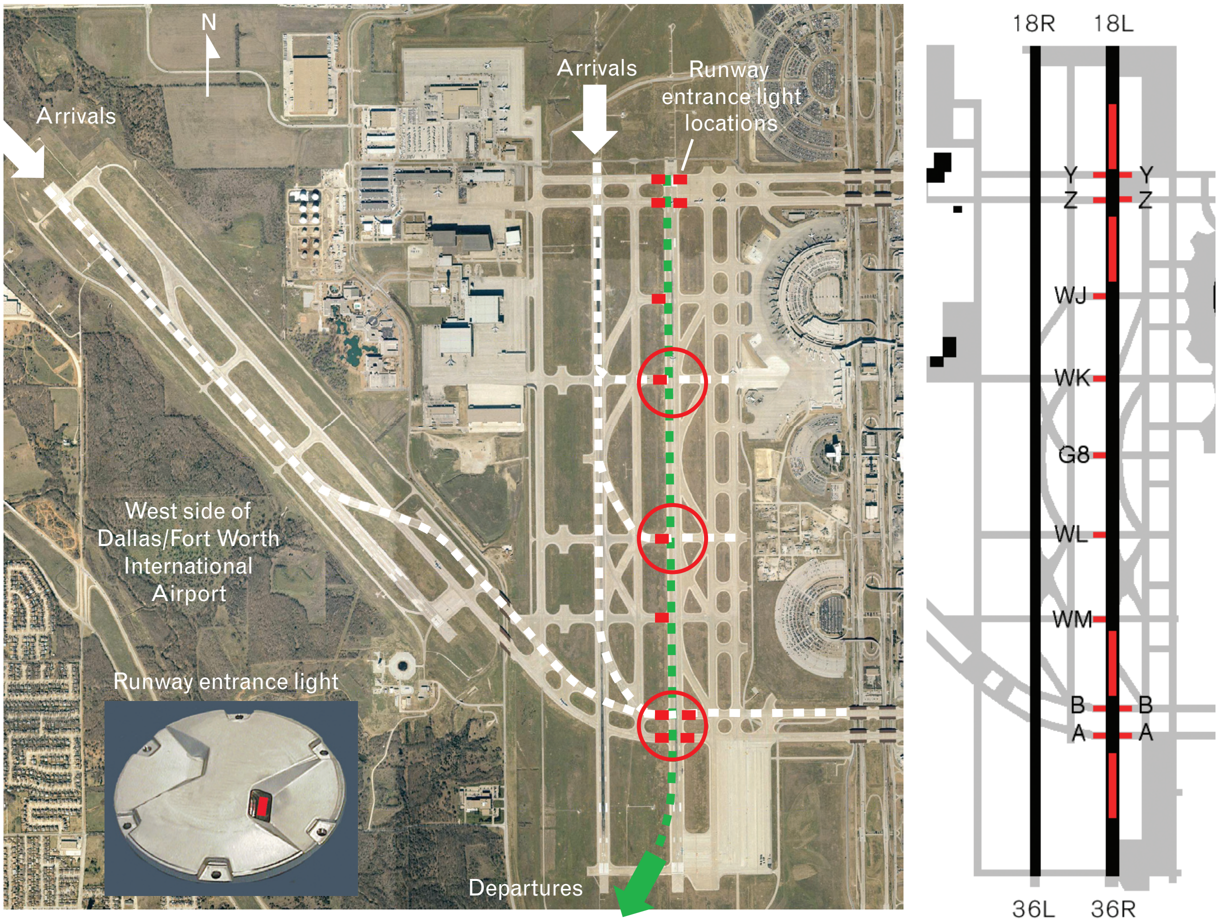

FIGURE 11. Traffic in south flow on the west side of Dallas/Fort Worth International Airport. Aircraft are landing to the south or southeast; departing aircraft take off from the inner runway, Runway 18L/36R. South flow designates the runway as 18L; north flow designates the runway as 36R. Arrivals mostly use the outer (westerly) runways and usually cross Runway 18L/36R at the intersections highlighted by red circles to reach the gate areas. Red boxes illustrate the locations of runway entrance lights and takeoff hold lights along Runway 18L/36R. The airport plan at the right gives the taxiway names. The inset photograph in the lower left shows an in-pavement runway entrance light fixture with red lens.

determined *a priori* by theoretical arguments and is apt to vary with location and traffic conditions, and from airport to airport. Unobserved anomalies cannot be ignored; they are important to a general understanding of overall system performance and may also be viewed as indicators of potential observed anomalies. Still, counting all anomalies without distinction would be misleading because it would ignore an inherent characteristic of RWSL operation and present a false picture of the effect of light anomalies on airport traffic.

Stage 4. Anomaly Evaluation and Classification

The final step in the assessment process was classifying the extracted anomalies in terms of the defined anomaly types, cause, and effect, and estimating the anomaly duration. Given the necessarily contextual illumination rules governing REL operation, precise determination of anomaly duration is inherently difficult, but sorting, by agreed-upon quantitative criteria, into 'short' and 'long' is always possible. Short anomalies are often, but not necessarily, operationally insignificant; their exact duration is, however, almost always unimportant, as well as impossible to quantify precisely. Long anomalies, on the other hand, are of operational importance, and an estimate of duration is desirable.

Automated evaluation and classification of anomalies is a development task that is not yet complete. Pending completion of this work, the process was be-

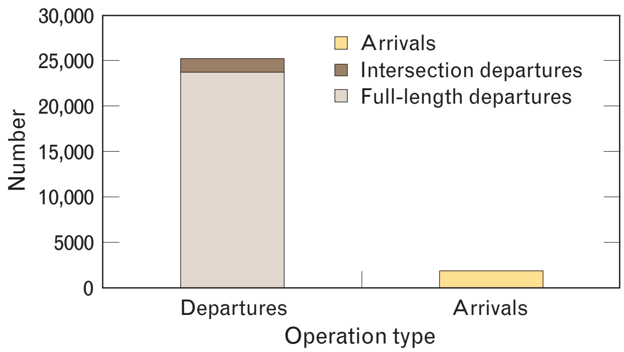

FIGURE 12. Total number of operations on Runway 18L/36R. About 93% of the operations on 18L/36R during the 2005 operational evaluation were departures, 88% were full-length departures, and 5% were intersection departures. Only 7% of the operations were arrivals.

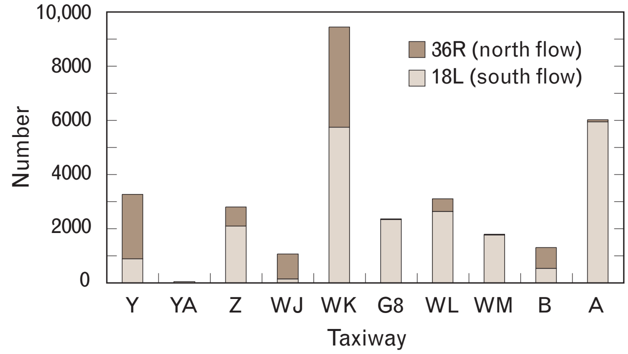

FIGURE 13. Number of crossings and entries at instrumented taxiways for Runway 18L/36R. The location of these taxiways is shown in the airport plan in Figure 11. Every runway crossing and entry at an instrumented taxiway (those equipped with runway entrance lights, or RELs) was counted during the operational evaluation. The most common crossing points were WK and A; the latter was used almost exclusively in south flow (i.e., aircraft departing on Runway 18L).

ing carried out in automation-aided manual mode, using event playback as needed.

The full assessment process depicted in Figure 10 needs to be carried out only for a system that performs at close to an operationally acceptable level of performance. Prior to that point, the processes of system development, adaptation, and tuning are served adequately by the first two stages of the assessment, which flag the dominant anomalies for evaluation and corrective action.

The Dallas/Fort Worth Operational Environment

The operational evaluation of the RWSL RELs was restricted to Dallas/Fort Worth Runway 18L/36R, which is shown in Figure 11. The red boxes in the figure illustrate the locations of the runway status lights along this runway. The dashed white lines show examples of how arriving flights on the outer runways must cross Runway 18L/36R—an inner runway used mainly for departures—to reach the terminals. The airport plan on the right side of the figure gives the taxiway names for the crossings where the runway status lights are located. Figure 12 summarizes the operations on this runway, clearly indicating a large ratio of departures to arrivals. The mix of departures and arrivals is similar in south operations (18L) and north operations (36R), but the pattern of taxi flow across the runway differs between the two airport configurations. Figure 13 summarizes the crossing statistics according to taxiway crossing location along the runway.

In addition to logging the departures, arrivals, and crossing traffic at each crossing point, the assessment also records REL activity and whether a crosser is in a position to observe an illumination. Such a crosser—called an affected crosser—may observe one or several illuminations. Neither the number of affected crossers nor the observed illuminations can be reliably estimated *a priori*, so the assessment provides a detailed count based on the actual interplay between traffic

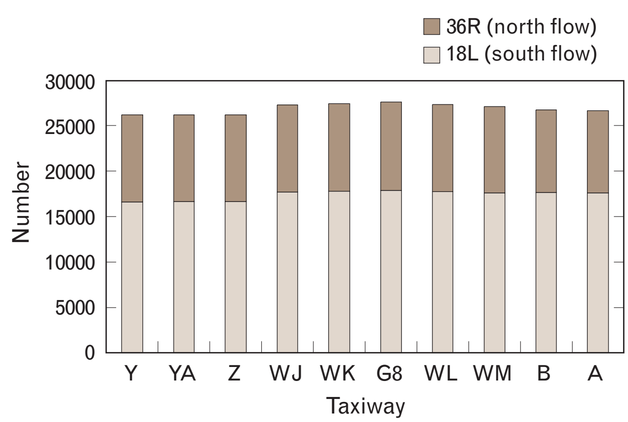

FIGURE 14. Illuminations at each instrumented taxiway for Runway 18L/36R. Because all RELs forward of a departure or an arrival are illuminated, most RELs illuminate for every operation, and light activity is nearly the same at all instrumented points. Midfield locations such as WK, G8, and WL see additional illuminations, reflecting the effect of intersection departures and landers re-illuminating RELs in landing rollout state.

and lights. The significance of these quantities is that they convey how the pilot population experiences the runway status lights.

The total number of illuminations recorded during the operational evaluation was nearly the same for all locations, ranging from 26,000 to 28,000, corresponding to about 25,000 departures and 2000 landings. The number of crossers affected by the lights at the various taxiway locations depends on overall crossing activity, as well as the relative timing of crossing operations and light activity, and the typical illumination duration at the location. Figures 14, 15, and 16 summarize the statistics for the number of illuminations, the number of crossing aircraft affected by the lights, and the number of observed illuminations, respectively, for each of the taxiway crossings along Runway 18L/36R. Overall, about 30% of the 36,000 crossers were affected, but the fraction ranged from 15% at WM to 45% at adjacent taxiway B. Observed illuminations averaged about 5% of the total overall, but exhibited considerable variability, from 13% at WK and 11% at A to 3.5% at B and 1.5% at WM. There are further differences between south and north operations.

Technical System Performance Assessment

Three light performance metrics were used: missed detections (MD), false activations (FA), and instances of interference (I). In a missed detection the lights were not on when they should have been, in a false activation the lights were on incorrectly, and in an instance of interference the lights were on and appeared to interfere with a safe crossing operation. In addition to evaluating the correctness of light operation, the technical system performance assessment also evaluated the impact of the lights on controller-pilot communications. The assessment also made special note of all instances of pilots entering the runway over illuminated RELs. These events are called light busts (LB); most represent a subclass of interference in which the pilot proceeded in spite of the lights' warning.

These definitions, and associated terminology, are similar to those used in earlier work at Lincoln Laboratory [7], except that, at the sponsor's request, the definition of interference was broadened to include FA illuminations that interfered with the flow of traffic. To avoid double counting, these were then excluded from the reported FA statistics.

Interference was also determined somewhat differently than in our earlier work. Instead of judging interference purely by observing the timing of crossing operations relative to REL deillumination (an inappropriate approach when RELs are installed on the airport and visible to the cockpit crew), light-off was related to the controllers' runway crossing instruction. A light turning off after the controller had finished

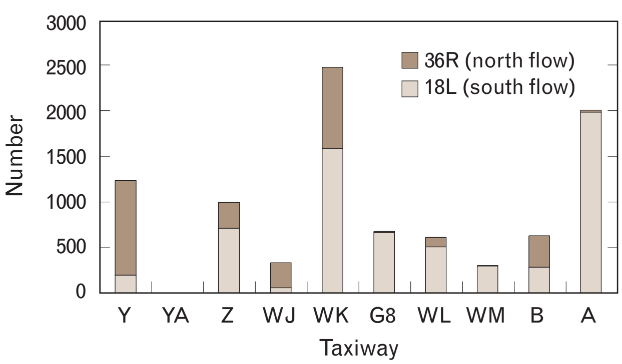

FIGURE 15. Crossing or entering aircraft affected by runway entrance lights. An aircraft or vehicle is affected by runway entrance lights if it is on the taxiway near the lights while the RELs are illuminated, and it subsequently crosses the runway at that location. Affected crossings—about 30% of all runway crossings—represent those which potentially benefit from REL operation. The distribution of affected crossings is qualitatively similar to that of crossings in general.

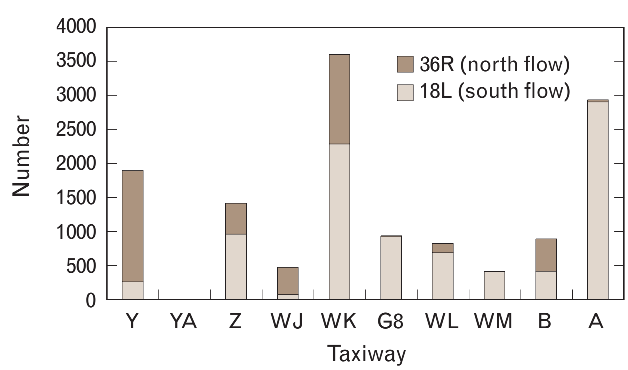

FIGURE 16. Observed illuminations. An observed illumination is an illumination with an aircraft or vehicle in a position to observe it. These differ in number from the affected crossings because a single aircraft may observe several REL illuminations while holding short of the runway. The distribution of observed illuminations across taxiways follows closely that of affected crossings. On average, about 5% of the REL illuminations are observed.

Table 1. Performance Goals for the RWSL Runway Entrance Lights*

Development phase	*MD*	*FA*	*I*	*Total*
Engineering development	320	1500	800	200
Shadow operations	360	1800	900	225
Operational evaluation	400	2000	1000	250

* expressed as inverse rate operations per anomaly

giving a crossing instruction was judged interfering, and if the light turned off four or more seconds late it was counted as interference, regardless of the movement, or lack thereof, of the crossing aircraft. This definition of interference is generally more stringent than the earlier one based on dynamics alone.

The interference check did not require listening to the totality of the voice traffic for Runway 18L/36R. The crossing record of the performance assessment traffic lists includes information on the timing of the crossing operation, relative to REL on-off times. This, the basis for the prior definition of interference, was used to identify those crossings that appeared to be most tightly timed. These were played back with audio, in the expectation that interference with crossing instructions was most likely to occur when the dynamics suggested tight timing.

The performance specifications for each of the three RWSL development phases were stringent, especially the operational evaluation FA specification. Table 1 lists these performance goals. Only anomalies of four seconds' duration or longer were counted against the specification, although all anomalies were logged. FAs and MDs were further broken down into observed and unobserved anomalies. Instances of interference are by definition always observed. In addition to observed anomalies, we also counted anomaly observations to take into account the multiplicity of observers at different crossing points.

The gradual improvement reflected in the performance specification has continued after completion of the operational evaluation; the anomaly rate eight months later is better than one per four hundred operations.

Anomaly Counts

The anomalies were classified by anomaly type and according to identified cause. A total of 114 anomalies of all types were identified in the data encompassing 27,000 departure and landing operations and 36,000 runway crossings. Approximately 40% of the anomalies were classified as MD, 50% classified as FA, and 10% classified as I. The causal breakdown was dominated by early or late off-ground reporting from some aircraft types (36%), resulting in classifications of MD or FA, respectively, and ASDE-X radar track problems (33%), predominantly causing classifications of FA and MD.

The anomaly, observed anomaly, and anomaly observation counts were normalized to the number of operations on runway 18L/36R. When the rates thus obtained (anomalies/operation) are expressed as their inverses (operations/anomaly), they can easily be compared with the operational goals shown in Table 1. Figure 17 summarizes these inverse anomaly rates.

The 114 anomalies resulted in 37 anomaly observations, 25% of which were classified as MD, 40% classified as FA, and 35% classified as I. The dominant causes were the same as for the total anomalies, off-ground declaration anomalies (40%), and ASDE-X radar track failures (24%).

Operational Feedback and Analysis

Feedback was obtained from air traffic controllers, supervisors, pilots, and vehicle operators during the operational evaluation. User observations and opinions were captured through questionnaires filled out by pilots and vehicle operators, as well as comments, notes, and observations made by controller supervisors, voluntary controller participants, and test team members. The operational feedback determined when and to what extent RWSL affects the normal operation of the airport, and served as a measurement of the operational suitability of RWSL in actual operational use.

Operational Feedback Measures

The assessment process recorded sufficient information to make quantitative estimates of the initial impacts of RWSL on system capacity, controller/pilot communi-

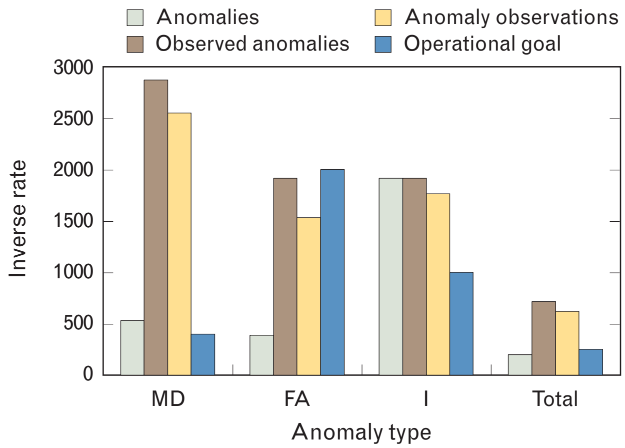

FIGURE 17. Inverse anomaly rates during Dallas/Fort Worth RWSL operational evaluation. The missed detection and interference performance goals were uniformly exceeded. The false activation goal was not achieved if all anomalies are counted. If only observed anomalies or anomaly observations are counted, however, the false activation goal was very nearly achieved.

cations, and pilot and vehicle driver situational awareness. Elicited observer feedback included suggestions for system tuning and safety enhancement. The pilots were asked to answer survey questions confidentially to assess the operational suitability, safety, and effectiveness of the runway status lights.

Voice transmissions between air traffic controllers and pilots and vehicle drivers over the local and ground frequencies were recorded for confidential *post hoc* analyses, including incident-by-incident investigation of any cases in which extended pilot-controller communications concerning some element of RWSL took place.

Pilot Survey

The pilot survey comprised eighteen yes/no response statements presented in a positive and negative counterbalanced order, with additional comments encouraged. We employed three survey methods—web site (www.rwsl.net), telephone, and paper (placed near posters in operations centers). Most pilots used the web site method to respond to the survey; the phone method was rarely used and is not included in this report because of insufficient response.

A total of 220 responses were received, 167 via the web and 53 in paper form. (The survey was identical for both data entry modes.) An additional 13 incomplete surveys were received, but these were not included in the final analysis. Surveys were collected over the whole operational evaluation period of three months, starting 1 March 2005 and ending 31 May 2005. About 75 completed surveys were received each month.

Overall reaction to the RWSL program and the RELs was very favorable among the participants. An overwhelming majority of the respondents (92%) felt that RELs would help reduce runway incursions, and 88% would recommend that RELs be installed at other airports. Only 26% of the respondents felt that the system needs some fine tuning in such areas as the configuration of the lights relative to the taxiway hold line, the timing of the lights, and the intensity of the lights. Only 6% of the respondents were expressly negative about the concept and/or its implementation.

Pilots' Comments about RELs

Many respondents expressed strong opinions about their experience with the RELs. Of the 220 respondents, 129 respondents—59% of the total respondents—elected to add unstructured and open-ended comments to the survey. These comments for the most part reflected the pilots' personal attitude toward and concerns about the RWSL system. This is a relatively high percentage of added comments, indicating the general interest and overall enthusiastic attitude that pilots felt about this system (typically only 15 to 20% of the respondents add free-form comments to such surveys). The majority of comments were quite positive and consistent with the given ratings.

Pilot Survey Statement Aggregated Results

To summarize the survey results, responses to specific statements were aggregated into three indices: understanding, operational effectiveness, and acceptance (as shown in Figure 18). The acceptance score was based on responses to statements asking whether RELs would enhance situational awareness and help reduce incursions, and whether additional installations would be recommended. The understanding score was based on responses to statements asking whether a pilot will cross a runway if RELs are red and whether RELs off indicate clearance to cross. The perceived operational effectiveness score was based on responses to state-

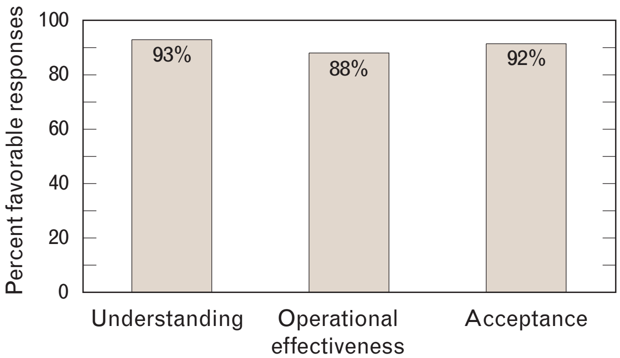

FIGURE 18. Aggregated favorable responses to the eighteen-question pilot surveys, grouped into three index categories.

ments asking whether the lights are (1) conspicuous; (2) consistent with air traffic control clearance; (3) not functioning; and (4) off when they should have been on, or on when they should have been off.

These indices were formed after the survey had been administered and were based on the logic behind the different statements. Significant correlations found between the responses to each statement within each index verified the logical groupings and added legitimacy to the choice to aggregate the statements into these indices. Furthermore, with increased exposure to RELs, the favorable responses of the pilots tended to increase across all three indices, as summarized in Figure 19. (Note that the operational effectiveness score for pilots with no exposure to RELs is missing because those individuals were asked to skip statements pertaining to this topic.) Finally, trends developed over time wherein situational awareness was rated less favorably by pilots who had issues with the presentation of the lights—i.e., conspicuity or configuration—although the former was later mitigated by using brighter default intensity settings for both daytime and nighttime in coordination with the Dallas/Fort Worth tower supervisors.

Overall the respondents rated RWSL very favorably. The large majority of the respondents (93%) understood or comprehended the operating procedures associated with RELs. Nearly all of the respondents (98%) stated that they would not proceed through illuminated RELs, and 89% stated that they would not interpret the off state of RELs as clearance to proceed. Indeed, not a single instance of a pilot and aircraft proceeding without clearance on light deillumination was noted either during or after the operational evaluation. This reaction is critical for the operational acceptability of the RWSL concept, because it must prevent more runway incursions than it causes. The perceived operational effectiveness index was also high, with responses for this index averaging about 88%. The overall acceptance of the system, based on the survey, was above 90%. Thus the pilot feedback indicates that RWSL is well understood and effective, and should be implemented.

Vehicle Operator Survey

Vehicle operators permitted to drive on and across runways were also surveyed for their reaction to REL operation. The pilot and vehicle operator surveys were identical, although the order of the statements differed somewhat between the two surveys. Although the sample size was small in comparison to the pilot respondents, the responses from the vehicle operators in most cases were similar to those of the pilots. A larger proportion of the vehicle operators felt that RELs would increase their verbal response time, most likely because they are trained to contact the tower if they were prevented from crossing a runway if the RELs were red when the clearance was issued (although there is an exception to the rule here, and controllers are allowed to tell vehicle operators to disregard the lights if a safe quick crossing is achievable). Also, compared to the

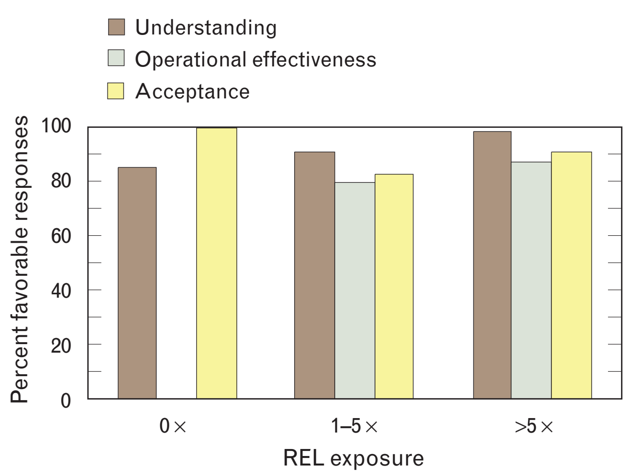

FIGURE 19. Positive responses to pilot survey questions, grouped by three aggregated index categories and by number of times each pilot was exposed to the RELs. Clearly, favorable responses increased as pilots with some exposure to the lights became more familiar with them.

pilot survey, a smaller percentage of vehicle operators felt that RELs would help reduce runway incursions and/or would recommend additional installations. Yet overall, the vehicle operators also rated RWSL highly on the understanding, operational effectiveness, and acceptance indices.

Summary of Operational Evaluation

The essential results of the operational evaluation of RWSL are contained in the two types of assessment: technical assessment, which measures system correctness; and operational feedback, which measures user opinions of the system. The technical assessment indicates that the system is working at the 99.5% level and improving in performance. The operational feedback indicates that users are satisfied with the level of performance and think that RWSL will improve runway safety.

It is important to note that these assessments can provide only a gross estimate of the safety benefit of RWSL in the operational environment, since the assessment period was too short to allow a statistically significant number of potential high-hazard events to occur on one runway at one airport. Obtaining useful operational feedback from participants in such events that were (or were not) prevented by RWSL is difficult because few users are likely to admit they would have done something dangerous if the RELs hadn't indicated red. There is at least anecdotal evidence that a nighttime crossing clearance was rescinded because red RELs were observed from the tower (thereby preventing both a potential operational error and possible runway incursion). In addition, there are cases in which crews have questioned a crossing clearance or even refused to begin a crossing upon observing red RELs in front of them.

Furthermore, since the operational evaluation of RELs began over a year ago, there has been only one reported runway incursion at the lighted intersections (which were selected in part because they had been problem intersections in the past). The one incursion was due to pilot error and occurred during the first week of the operational evaluation, when pilots had minimal experience with the runway entrance lights. In fact, the majority of pilots, controllers, and air traffic control supervisors have developed a growing trust of the system, and the Dallas/Fort Worth tower management feels certain that red RELs have stopped unsafe crossings. Given the visual indication to pilots, the Dallas/Fort Worth operations manager has stated, "I feel a whole lot better with that system running because I know those lights do what they're supposed to do!"

The operational evaluation proved the RWSL concept meets the key high-level requirements: that the runway status lights operate automatically, that no controller action is required for their operation, that the lights accurately depict runway status to pilots and vehicle operators, and that the lights do not interfere with normal safe surface operations. Although it works well, the system currently does not work perfectly. That the system is well received by a large majority of controllers, pilots, and vehicle operators indicates, however, that the RWSL system is currently operationally suitable even at a busy airport.

Future Work

The current operational evaluation of RWSL has concentrated on runway entrance lights. Takeoff hold lights are being tested in ongoing work at Dallas/Fort Worth, with an operational evaluation that has begun and will continue throughout 2006. Pending successful completion of this evaluation, the FAA has indicated that additional runways at Dallas/Fort Worth may be tested, as well as other airports. The Dallas/Fort Worth FAA management has also expressed a keen interest in having RWSL extended to include protection to landing aircraft. Meanwhile, an operational evaluation of RWSL at San Diego International Airport is scheduled to commence later in 2006, and Lincoln Laboratory is currently supporting the FAA in that effort.

Significant research questions have yet to be answered; these questions will require testing at other airports. First, different airports have different surveillance systems and environments that may affect the rate of RWSL anomalies. Second, takeoff hold lights for crossing runways require special logic that has been designed and implemented but not tested extensively. Third, some airports use a procedure called Land and Hold Short Operations (LAHSO) that will require a change in runway entrance light operation at entrances downfield from the hold short point. In addition, there are significant human factors and interoperability issues specific to the presentation and operation of the lights (e.g., their configuration, visibility, and distinctiveness) that necessitate further exploration. For example, the runway entrance lights and takeoff hold lights should be assessed when implemented in conjunction with airports that have active Surface Movement Guidance and Control System (SMGCS) red stopbars, which are currently used for low-visibility operations at selected airports. All these questions will need to be answered before RWSL can be fielded at a wide variety of airports.

Ongoing research is also focused on improving RWSL performance at Dallas/Fort Worth. Two sorts of improvement are being pursued; improvements in surveillance to provide a more accurate picture of airport activity, and improvements in safety logic to provide a more faithful representation of runway status based on that picture. These are both components of the relentless pursuit of improvement in our nation's aviation transportation system.

Conclusion

Results from the operational evaluation of runway entrance lights at Dallas/Fort Worth support the conclusion that runway status lights shown directly to pilots and vehicle operators offer the potential to reduce runway incursions and runway conflict accidents by increasing overall situational awareness of the dynamic runway environment. The operational evaluation phase of RWSL was a live test with actual traffic and a limited deployment of a field lighting system, along with a presence in the operational tower in order to prove the runway status lights concept. The operational evaluation test period provided critical technical performance and operational feedback information required to assess the operational suitability and correctness of operation of RWSL in providing an important safety function. The operational evaluation proved that the RWSL concept meets the key high-level requirements. The results of the operational evaluation will serve as validation to continue further evaluation of runway status lights at other busy airports in the National Airspace System. The eventual deployment of RWSL as an operational system has specifically and repeatedly been requested through feedback from the end users (especially airline pilots), who have commented and continue to comment on their favorable interactions with the system under test.

Acknowledgments

The success of this project is the result of the efforts of many individuals and organizations. We would like to especially acknowledge Jaime Figueroa, Vincent Chu, Peter Hwoschinsky, Ed Feigenbaum, Paul Donaldson, Gary Birdwell, Victor Nartz, Ed Runyon, Dave Olster, William Mino, Stan Bissell, Jack Eppard, Ron Nichol, Elizabeth Mauer, Gabriel Spitz, Richard Williams, and Eric Shank. This work was sponsored by the Federal Aviation Administration.

REFERENCES

1. "NTSB Most Wanted Transportation Safety Improvements 2006," brochure, November 2005, www.ntsb.gov.
2. FAA Runway Safety Program Office, "The Human Factor: The Major Factor in Runway Incursions is Human Error," Runway Safety Southwest Region, May 2003.
3. W.M. Hollister, "Airport Surface Traffic Automation Study," Lincoln Laboratory Project Report ATC-156 (9 May 1988).
4. E.F. Lyon, "Airport Surface Traffic Automation," *Linc. Lab. J.* 4 (2), 1991, pp. 151–188.
5. H. Wilhelmsen, "Preventing Runway Conflicts: The Role of Airport Surveillance, Tower-Cab Alerts, and Runway-Status Lights," *Linc. Lab. J.* 7 (2), 1994, pp. 149–168.
6. J.R. Eggert, "Demonstration of Runway-Status Lights at Logan Airport," *Linc. Lab. J.* 7 (2), 1994, pp. 169–186.
7. H. Wilhelmsen, M.P. Kastner, T.J. Morin, and J.L. Sturdy, "Performance of the Runway-Status Light System at Logan Airport," *Linc. Lab. J.* 7 (2), 1994, pp. 187–214.
8. S.D. Young, R.W. Wills, and R.M. Smith, *Pilot Evaluations of Runway Status Light System,* NASA Technical Memorandum 4727, Langley Research Center, Hampton, Va., Oct. 1996.
9. H. Wilhelmsen and J.R. Eggert, "Runway Conflicts at Major US Airports 1997–2000: Patterns and Prevention," ATC Project Memorandum 42PM-RSLS-0001, Lincoln Laboratory, 7 October 2002.
10. Runway status lights project web site, www.rwsl.net.

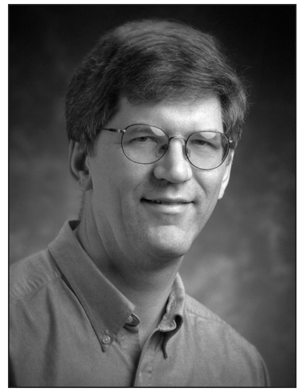

JAMES R. EGGERT
is a staff member in the Surveillance Systems group, and the technical lead of the Runway Status Lights program. His research specialty is in airport surface-traffic surveillance and automation. He received a B.A. degree in physics and linguistics from Rice University in 1979 and a Ph.D. degree in physics from Harvard University in 1986. He has also pursued graduate studies in linguistics at the Universität Wien. As a graduate student at Harvard University, he was an IBM Predoctoral Fellow and he received the John Tyndall Prize. He is married, with two children, and his hobbies include genealogy and notgeld collecting.

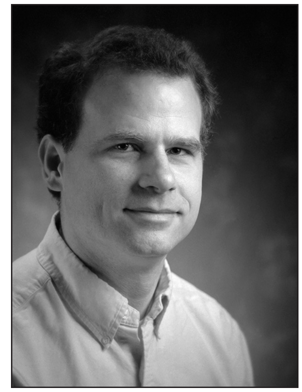

BRADLEY R. HOWES
is a staff member in the Surveillance Systems group, where he is the software lead for the 331 SideCar effort. He received a B.A. degree in English literature and music composition from Sarah Lawrence College. Prior to joining Lincoln Laboratory in 2000, he worked for various startup companies and large Fortune 100 firms as a software developer on a variety of tasks, such as automated analysis tools for the Human Genome Project; a media searching application using speech recognition; a parts layout program for large computer numeric control (CNC) machines; and a corporate-wide intranet collaboration and document storage tool for Motorola. His only regret so far is that he turned down a job offer from Microsoft in 1987. He is an avid sailor, and races in the Soling and Sonar classes.

MARIA PICARDI KUFFNER
is a specialist in the Surveillance Systems group. She has designed and analyzed advanced computer/human interfaces for air traffic control, developed and conducted training of pilots and controllers, and initiated and continues to manage the project web site. Her ongoing concentration is providing technology and training to prevent runway accidents, specifically by supporting field evaluations of runway status lights at Dallas/Fort Worth and San Diego airports. Before joining the Laboratory, she worked for Raytheon Company for five years as a senior engineer. Prior to Raytheon, she was a member of technical staff in the Human Engineering Laboratory at GTE Laboratories. She received an M.S. degree in cognitive/experimental psychology from Arizona State University on a Regents scholarship. Earlier, she served as a research assistant for five years at the Psychophysics Laboratory of Harvard University while completing graduate courses in psychology and law, and she worked as a programmer/analyst for the Air Force Geophysics Laboratory throughout her undergraduate program at Regis College. She is a past president of the Human Factors and Ergonomics Society's New England Chapter, and she holds a private pilot license. She was awarded a Certificate of Appreciation in recognition of contribution to the National Plan for Aviation Human Factors by the FAA.

HARALD WILHELMSEN
is a former staff member in the Surveillance Systems group. He is now with DAG Consulting in Burlington, Massachusetts. 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, airport surface-traffic surveillance and automation, and automation system quality control.

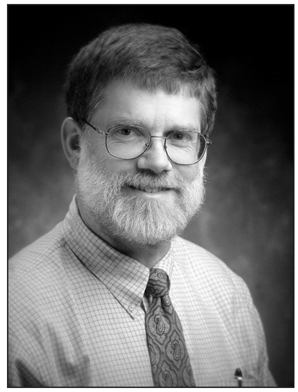

D. JONATHAN BERNAYS
is an assistant group leader in the Surveillance Systems group. He received a B.S. degree in aeronautical engineering in 1979, and B.S./M.S. degrees in electrical engineering in 1981, all from MIT. From 1981 to 1984 he was a member of the technical staff at the IBM San Jose Research Laboratory, and from 1984 to 1986 he was a flight controls engineer for Boeing Military Airplane Company. In 1987 he joined the Optical Communications group at Lincoln Laboratory, where he worked on the Laser Intersatellite Telecommunications Experiment (LITE) program. In 1992 he transferred to the Air Traffic Control Systems group, where he worked on a several data link and beacon surveillance programs. In support of the FAA's Data Link program, he led an effort to develop, flight test, and deploy a radar-based, automated traffic advisory system, now provided by the FAA as Traffic Information Service. He also led teams focused on the development, standardization, and validation of Automatic Dependent Surveillance–Broadcast (ADS-B), a cooperative surveillance system in which aircraft self-report their GPS-based positions. His current responsibilities include oversight of the Runway Status Lights program, as well as a variety of FAA and homeland defense program. He is an instrument-rated private pilot, and enjoys technical mountaineering and coaching soccer in his spare time.