

Experimental Examination of the Benefits of Improved Terminal Air Traffic Control Planning

Airport capacity can be improved significantly—by precisely controlling the sequence and timing of traffic flow—even when airspace usage and procedures remain fixed. In a preliminary experiment, a plan for such sequencing and timing was applied in a simulation to a 70-min traffic sample observed at Boston's Logan Airport, and the result was a 13% increase in terminal throughput. A total of 2.2 aircraft flight hours were saved. Delays imposed upon arriving traffic in the simulation were much more equitably distributed than in the actual traffic sample.

An even greater improvement may be possible if controllers are able to space aircraft more precisely on final approach than was achieved in the simulation. If the plan had been followed precisely, the throughput increase would have been 23%.

Landing rates under instrument meteorological conditions (IMC) are typically half that achievable under visual meteorological conditions (VMC). But airlines tend to schedule flights at about the terminal capacity obtained under good weather conditions. The loss of capacity in bad weather leads to delays, and delays at one airport can cause secondary delays throughout the air transportation system.

The Terminal Air Traffic Control Automation (TATCA) program has been studying the use of improved sequencing and timing to increase the landing rate at major airports, especially under IMC. This has been accomplished by developing a plan for efficient flow of traffic in the terminal area that provides advisory information to controllers to help them achieve the planned timings.

The main objective of the TATCA plan is to ensure that landings are spaced as closely as possible—without violating separation standards or other operational requirements. To maximize its effectiveness, the plan will extend into the en route facility that feeds traffic to the terminal. Preparatory timing of en route arrivals will prevent gaps or bunching, and ensure a steady flow of merged arrivals.

The TATCA approach to flow planning is in some ways similar to the FAA's current and

planned automation of metering, sequencing, and spacing. But there are important differences. Those programs deliver aircraft to the terminal airspace at a manually specified average hourly rate.

TATCA, by contrast, will maximize the efficiency and throughput of terminals (within the limits of operational constraints) by adapting the delivery rate to the continuously changing terminal environment. The resulting rate will be the greatest that current conditions will allow. To achieve this goal, however, TATCA needs more precise observations of winds and more accurate predictions of individual aircraft movements than current metering automation provides.

A number of control advisories will be provided by TATCA to assist controllers in achieving the timings in the TATCA flow plan. In fact, the process of determining which advisories are needed and how best to display them to controllers has been a significant part of the TATCA development program. Among the advisories considered are conformance indicators, selection of the point where the descent from cruise altitude should begin (the top-of-descent point), the speed profile for the descent, speed reduction advisories in the terminal area, and final-approach spacing aids that achieve the planned

separations on the final-approach path.

Although the theoretical capacity gain from precise scheduling of a runway can be determined through calculation, operational considerations will affect results during actual use. Therefore, the impact of improved planning and scheduling on a sample of live traffic data recorded at Boston's Logan Airport was measured.

In the experiment a revised landing schedule was manually produced, and then all the arriving aircraft were relanded by using a simulator. Air traffic controllers with experience in the Boston area assisted in the control of the simulated traffic and reviewed the results for operational acceptability. The actual landing rate was then compared with the results of the simulation. Because realistic controller interfaces were not available, the study did not explore ways of presenting the plan to controllers or its effect on the controller workload.

Environmental Conditions

The traffic data were recorded on 15 December 1987 by the Mode S Experimental Facility (MODSEF) radar beacon interrogator at Lincoln Laboratory. About 70 min of data were recorded, consisting of traffic within 50 nmi of MODSEF, and including Logan Airport and at least a portion of each of the three holding stacks in use at that time.

During the test period, 34 aircraft landed at Logan. At the beginning of data collection, the Logan Automated Terminal Information Service (ATIS) was reporting a 300-ft ceiling, 1/2-nmi visibility, and ice pellets. The wind was 3 knots from 150°. By the end of the recording session conditions had improved to a ceiling of 500 ft, broken, with 4-nmi visibility. The wind then was 15 knots from 130°. Thus the traffic data were recorded during IMC.

During this time there were active holding stacks at each of the three turbojet entry fixes that feed the terminal area. The capacity of one of the stacks (at Providence, R.I.) was exceeded, and overflow traffic had to be sent to an entry fix over the ocean east of Boston. The stacks kept the aircraft arrival "pipeline" full throughout the measurement. Therefore, excessively large

landing intervals could originate only in the stack or terminal areas and could not result from system-wide flow restrictions or en route delays.

During recording of the live traffic data:

- There were two runway changes, from runway 15R to runway 4R and back again.
- The departure flow was handled largely by a separate, single runway (runway 9) that was largely independent of the landing runways, so that there was negligible interaction between departures and arrivals.

Generation of the Revised Schedule

The revised schedule was generated by manually obtaining time-to-fly and speed information from plots of the 15 December data, and then entering these parameters into a spreadsheet program to produce a revised schedule of landing times. Manual entry of information was used in order to evaluate the benefits of the concept—TATCA will perform the same analyses in real time.

Standard Traffic Flows

Under the procedures in effect on 15 December 1987 [1, 2] high-altitude turbojets entered the Boston Terminal Radar Approach Control (TRACON) area (Fig. 1) from one of three fixes. Traffic from the west held at the Gardner (GDM) VOR (VHF Omnidirectional Range—a navigational aid) and entered the TRACON via the LOBBY fix. Traffic from the south held at the Providence (PVD) VOR and entered via the MIXER fix. Oceanic traffic, and some overflow from Providence, held at the SCUPP fix over the ocean east of Boston and also entered via that fix. Aircraft arriving via these fixes generally crossed them at 10,000 or 11,000 ft, unless their entry point was close enough to the approach for the active runway that a lower altitude had to be assigned.

Low-altitude arrivals were also assigned standard entry points and routes. These aircraft generally entered the TRACON at 4,000 ft via the EXALT, FREDO, WILKI, or LOBBY fixes, or over

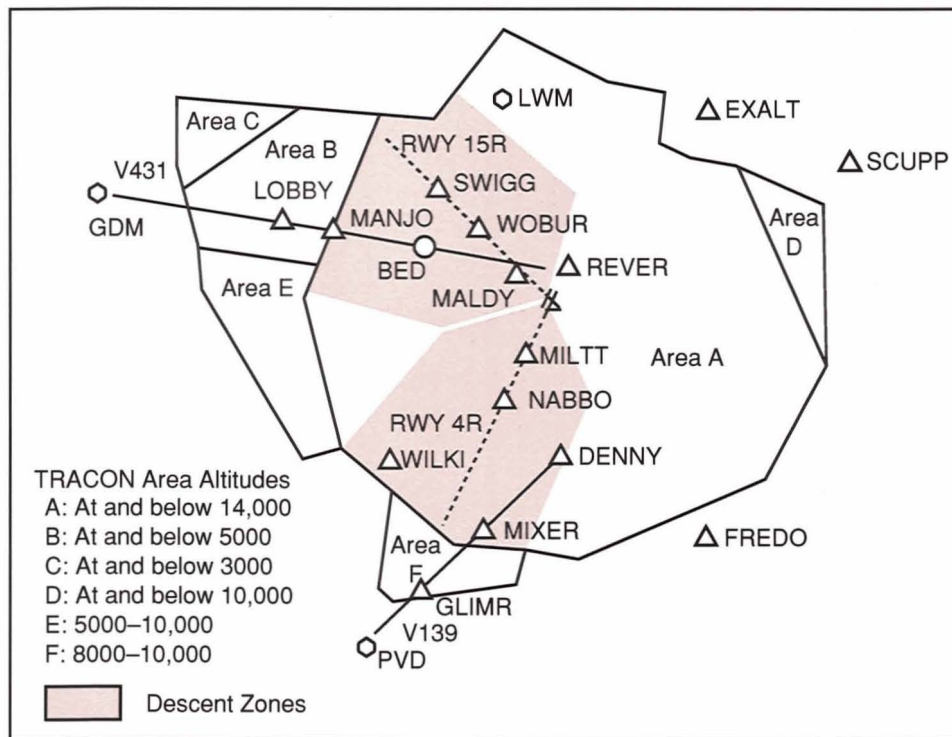


Fig. 1—Boston TRACON. Descent zones and localizer paths for runways 4R and 15R are shown.

the Lawrence (LWM) VOR.

A composite of the paths for the 34 aircraft arriving in Boston during the 15 December data-taking session is shown in Fig. 2. The photograph shows the actual approach paths and holding stacks.



Fig. 2—Composite of arrival paths during data-recording period on 15 December 1987.

Initial Aircraft

The first runway change (from runway 15R to runway 4R) occurred shortly after the start of the measurement. The seven aircraft that landed on 15R were not rescheduled (see Fig. 3).

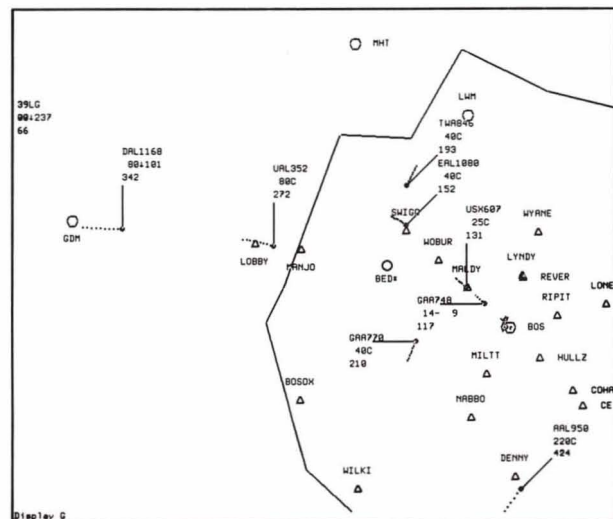


Fig. 3—Initial seven aircraft landing on runway 15R.

Three of these (GAA748, USX607, and EAL1080) were already on the localizer, one (TWA846) was about to intercept the localizer, and one (GAA770) was about to turn to its downwind leg. The remaining aircraft (UAL352 and DAL1168) were about to enter the terminal area, shown by the boundary line. Given this situation, the landing sequence was essentially determined.

Landing Sequence for Runway 4R

For the 27 remaining aircraft, estimates were made of the minimum times needed to fly from their initial observed positions (along the standard arrival routes) to runway 4R. The estimates were derived from the speeds, altitude rates, and flying times observed in the MODSEF data. In actual operations, these post facto estimates would be replaced with estimates calculated in real time by TATCA aircraft-modeling functions. By adding the estimated flying time to the time that the aircraft first appeared, the earliest possible runway-arrival time for each aircraft could be estimated. The aircraft were then sorted according to their earliest possible landing times, which determined the landing sequence.

Appropriate time spacings between successive aircraft had to be maintained. The scheduled landing time for an aircraft was then taken to be either its earliest estimated landing time or the landing time that would give it the minimum allowable spacing behind the previous aircraft, whichever was later. For all aircraft in this data set, the earliest estimated landing time preceded the time required by spacing, and so the scheduled landing times were determined entirely by minimum-spacing considerations.

To determine the scheduled landing time of the first aircraft on runway 4R, a fixed time interval was added to the scheduled landing time of the last aircraft to land on runway 15R. In the real traffic data, the interval was approximately 2 min and 40 s. Initially, the rescheduled plan allowed 60 s between the last aircraft landing on 15R and the time when the first aircraft landing on 4R crossed the runway

threshold. However, after the simulation experiment was completed, air traffic specialists pointed out that during a runway change an adequate separation (3 nmi or 1,000-ft altitude) must be provided for the possibility of a missed approach, and that a 60-s interval was too short. This problem was handled by removing the time interval between the last aircraft on 15R and the first aircraft on 4R from both the observed data set and the simulated data set.

For the aircraft that followed the first one, landing times were calculated so that the minimum aircraft separations were never violated along the final-approach path from the outer marker, which is 5 nmi from the runway threshold, to the runway. The outer marker for runway 4R is the MILTT fix and for runway 15R is the MALDY fix (Fig. 1).

The calculation of landing times took account of the various separations required for ATC wake-vortex separations. Normally, aircraft must be spaced 3 nmi apart within a TRACON. Special wake-vortex separation requirements apply, however, whenever an aircraft is operating directly behind an aircraft in a heavier weight class or following it on an instrument approach. There are three weight classes: heavy, large, and small. A heavy following a heavy must be spaced 4 nmi behind, and a large or small aircraft following a heavy must be separated by at least 5 nmi.

Additional separation is required whenever a small aircraft lands behind a heavier aircraft. At the time the heavier aircraft is over the landing threshold, the small aircraft must be separated by 4 nmi if the preceding aircraft is large, and 6 nmi if the preceding aircraft is heavy. Among the 27 aircraft that were rescheduled, only two were heavy (there were also two in the initial sequence of seven aircraft on runway 15R). The remainder were known to be large aircraft, except for four, which were unknown (no flight strips were available for them). The four unknown aircraft were assumed to be small. Three of the small aircraft landed behind large aircraft, and one followed another small one.

Fix Departure Time

A fix departure time was calculated for each aircraft, to indicate when it should leave its holding fix. This time was obtained by subtracting the fix-to-runway flying time from the scheduled landing time.

Since the low-altitude aircraft often didn't cross a defined entry fix, the time-to-fly could not be measured. Instead, a required delay was calculated by subtracting the earliest estimated landing time from the scheduled landing time. In the simulation, the fix departure time for these aircraft was determined by adding the required delay to the time when each aircraft first reached its desired holding fix.

The Simulation

To demonstrate that the schedule could be implemented without violating any ATC procedures, the landing schedule was tested in a simulation. The air traffic simulation program was one that accepted manual input of ATC clearances. The program ran on a Symbolics 3670 computer. Aircraft clearances were determined while the simulation ran, and also recorded in a script that could be followed automatically on subsequent simulation runs. Thus repeated runs could be made, each time improving clearances that had not worked out properly, but preserving those actions which had been successful.

We could stop the simulation and continue from the same point, so that decision making and clearance entry (via keyboard and mouse) did not have to be performed in real time. Generation of clearances for all aircraft in a reasonable period of time required three people, including a former air traffic controller familiar with Boston TRACON operations. Appropriate ATC procedures, as defined by Refs. 1 through 4 and the appropriate Boston approach plates (Fig. 4), were followed throughout. The actual procedures in use were also verified by checking the MODSEF data.

Initialization of Aircraft

All simulated aircraft were initialized to ap-

pear at the same points as the corresponding real aircraft, with the same initial speeds and headings. Most aircraft appeared at the point where they were entering their holding stacks, and the simulated aircraft entered the stack in the same way. Several aircraft in the MODSEF data flew past Providence to the stack at SCUPP, probably because the Providence stack had filled. In the simulation the corresponding aircraft were made to do the same, and the speeds and altitude rates were matched to ensure that stack entry times and altitudes corresponded to the real data.

Holding Procedures

Once each simulated aircraft arrived in its assigned stack, it did not follow the path of the real aircraft. This was felt to be unnecessary for the purposes of the simulation. The five aircraft initially in the Providence stack, however, were made to match the actual flight paths in the stack quite closely, in part to demonstrate that the simulator was accurately simulating aircraft behavior.

The remaining aircraft were held according to the *Airman's Information Manual* (AIM) standard procedures [4]. The four small aircraft in low-altitude stacks were held at speeds corresponding to those of the real aircraft, if known, or were assigned the holding speeds for propeller aircraft from the AIM (175 knots). Rather than rigorous adherence to the AIM's instructions for the length of inbound/outbound legs, the overall stack length was made to conform approximately to the observed stack size, if known. Aircraft were issued altitude clearances that kept them at least 1,000 ft apart at all times.

The initial goal was to bring aircraft out of the stacks very close to their scheduled fix departure time. However, this procedure didn't provide sufficient slack time to allow for the path variations needed to adjust aircraft spacing. To account for this effect, aircraft were maneuvered so that they crossed their entry fix prior to the scheduled fix crossing time, normally by 30 to 60 s, but sometimes as much as 120 s early. This strategy is known as *delay discounting*. (Ref. 5 uses this term in reference to delays

While aircraft are under the control of the en route center, the required separations are 5 nmi horizontally or 1,000 ft vertically. (Above Flight



Level 290—approximately 29,000 ft—a 2,000-ft vertical separation is required.) These conditions were easily maintained in the stacks, and the few aircraft that flew to SCUPP from Providence had longitudinal separations well in excess of 5 nmi.

Within the TRACON, 3-nmi radar separation or 1,000-ft altitude separation was required. This requirement was usually easy to satisfy up to the point where the streams of traffic from the various fixes merged. The aircraft from any given fix usually had large longitudinal separations.

The one problem with aircraft separation occurred in the Providence stack, because the first four aircraft were scheduled to leave the holding fix very close together. The resulting spacings were exactly 5 nmi. Even though this satisfied the en route constraint, and exceeded the TRACON requirements initially, problems arose as the leading aircraft slowed down during their descent, and also as the simulation controllers issued speed reductions. To avoid separation problems on the localizer, several of the Providence aircraft were flown over delaying paths at entry to the TRACON, which increased their initial longitudinal separations.

In those areas where traffic streams merged, altitude separation techniques were occasion-

ally used to maintain separation until the merge was completed and the planned sequence established. This use of altitude separation techniques did not appear to introduce any problems.

Second Runway Change

The time at which the second runway change (back to runway 15R) was announced was obtained from an audio recording. The announcement occurred approximately 11 min before the first aircraft landed on the new runway. Figure 5 shows the situation at that time. When the runway change was announced, all of the remaining aircraft in the study were in the TRACON. A plan for the runway change was then devised. Because of the small number of aircraft remaining, no scheduling aid was used.

Three aircraft were already in the vicinity of runway 15R and they were chosen to be the first to land. ATL405 was a low-altitude arrival and could easily be sent around to intercept the 15R localizer from the right (southwest) side, landing first. COA96 and TWA810 required more time and space to make their descents and speed reductions. They were therefore kept in their sequence, sent over the 15R approach, and then put on a downwind leg on the northeast side of the localizer.

UAL898 and EAL644 were approaching runway 4R from Providence. They could either land on 4R before ATL405 landed on 15R, or approach 15R behind TWA810. Neither had yet descended below the minimum allowable altitude of 6,000 ft outside the descent zone, and so they could be held at or above 6,000 ft and moved to the 15R descent zone. Moreover, landing the two aircraft on 4R would have required finding a way to delay them while other aircraft landed, and would then have delayed the landing of ATL405, because the simulation controller had fallen approximately 2 min behind the planned schedule at this point. Unlike a real-time TATCA system, the simulated schedule did not adapt to this condition. Because of these potential problems, both aircraft were assigned to 15R.

A scheduling aid such as TATCA would have

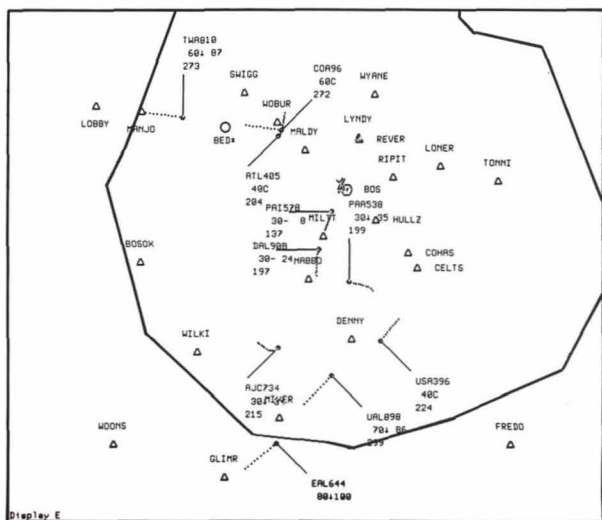


Fig. 5—Situation at the time when air traffic controllers announced the second runway change.




Fig. 6—Composite of paths for the 34 simulated arrivals.

been able to predict whether ATL405 would have been delayed under various circumstances. In the original data, there was a gap of 5 min between the last landing on runway 4R and the first one on 15R; in the simulation the gap was only 2 min and 48 s. Yet even that gap may have been long enough that a scheduling aid might have decided to land UAL898 on 4R and send EAL644 to 15R. However, that decision could have caused a gap on 15R until EAL644 arrived.

An even better procedure might have been to take the aircraft out of order, landing EAL644 on 4R in UAL898's slot and sending UAL898 to 15R. This would require a scheduler that could optimize the choice of runways when two runways are available. None of these alternatives have been investigated in simulation. Also, no attempt was made to account for aircraft that followed the 34 in the study. The later aircraft could have affected the choices at the time of the runway change, because some of them might have been able to reach runway 15R before any of the aircraft in the simulation.

Because the time interval between the last simulated landing on 4R and the first one on 15R was larger during the second runway change than the first, the rule for separation of aircraft during a possible missed approach did not affect the results. The time interval of 2:48 was therefore retained in the analysis data.

Results

Figure 6 shows a composite of the paths of all 34 simulated aircraft. These results should be compared with the paths of the real aircraft, which were presented in Fig. 2. In general, the scheduling procedure appeared to work well—the overall landing rate was 34 aircraft per hour, a throughput increase of 13%.

The difference between the outer-marker crossing times observed in the original data and the outer-marker crossing times of the simulated aircraft gave an average flying time saved per aircraft of 236 s. Much of the time saved was due to one aircraft, PBA3516, which was delayed by an unusual amount in the real data. But even if that aircraft were removed from the sample, the average time saving would still be 170 s per aircraft.

As shown in Fig. 7, only eight of the 34 aircraft landed later in the simulation than they did in the original data; 19 aircraft arrived earlier. By the time the 27th aircraft landed, enough time savings had accumulated that all subsequent simulated aircraft landed earlier than the corresponding real aircraft.

The 13% improvement was achieved despite the approximate nature of the time-to-fly estimates and the landing-time-interval calculations, despite the fact that the schedule was not dynamically updated in response to events, despite not being able to affect the schedule for




Fig. 7—Delay change after resequencing as a function of position in original landing sequence.

the first seven aircraft, and despite the final-approach spacings planned by the scheduler not actually being achieved. Except for the problem in scheduling of the first runway change, the schedule was implemented without violating any aircraft separation standards or other ATC rules and standard operating procedures.

During the simulation, it was noted that minor course adjustments and minor timing adjustments frequently caused significant spacing changes. For example, shortening a downwind leg by 10 s normally shortens the final by an additional 10 s. For a ground speed of 170 knots, the 20-s reduction would result in just under a 1-nmi difference in the position of the aircraft on final approach. This sensitivity made it difficult to achieve the minimum spacings assumed by the planned schedule.

In general, spacings within 1 nmi of the minimum were obtained during the simulation. During the data analysis the distribution of spacings achieved by the real controllers were compared with those achieved by the experimenters. The distributions were quite similar, but many of the very large separations seen in the real data were greatly reduced in the simulation because of the improved metering of aircraft over the TRACON entry fixes.

Because minimum spacings were not achieved, the simulated landing times gradually slipped further behind the ideal scheduled times, falling approximately 2 min behind at the time of the second runway change. If the scheduled landing times had been maintained, UAL898 and EAL644, the last two Providence aircraft and the last two simulated aircraft to land, would both have landed on runway 4R well before the time that ATL405 landed on runway 15R. Thus it would not have been necessary to send them around to runway 15R, which would have improved the results by an additional 271 s and given an overall 23% improvement

in the landing rate relative to the one actually measured on 15 December 1987.

Conclusions

This study has demonstrated that significant capacity benefits can be achieved in the current ATC system by more precisely controlling the timing of traffic flow, even while airspace usage and procedures remain fixed. The TATCA program will now pursue an intensive effort to address the human-factor and detailed algorithm-design issues that must be resolved before the planning system can be placed into operation. This effort will be carried out in a simulation environment, with realistic displays and a real-time version of the planning algorithms.

Acknowledgments

Dr. Steven Boswell of Lincoln Laboratory assisted in analyzing the data from this experiment. James Sturdy of Lincoln Laboratory and Tom Moody of UFA, Inc. spent many hours controlling simulated traffic. This work was supported by the FAA.

References

1. Boston Center/Boston Tower Letter of Agreement, 4 June 1987.
2. *Standard Operating Procedures*, Boston Air Traffic Control Tower (ATCT) (1 Dec. 1987), BOS TWR 7110.35F, Boston ATCT.
3. *Air Traffic Control, Order 7110.65E*, Department of Transportation, Federal Aviation Administration (9 Apr. 1987).
4. *Airman's Information Manual*, Department of Transportation, Federal Aviation Administration. Updated versions are issued several times a year.
5. L. Credeur and W.R. Capron, "Simulation Evaluation of TIMER, A Time-Based, Terminal Air Traffic, Flow-Management Concept," *NASA Technical Paper 2870* (Feb. 1989).
6. B.G. Sakkappa, "A System Concept of Metering and Spacing," *Proc. 22nd Annual Mtg. Air Traffic Control Assoc.*, 10-13 Oct. 1977.



DAVID A. SPENCER is a senior staff member of the System Design and Evaluation Group at Lincoln Laboratory. His focus of research has been in com-

puter software systems related to air traffic control. Dave received a B.S.E. in electrical engineering from Princeton University, and an E.E. and S.M. in electrical engineering and computer science from MIT.



JOHN W. ANDREWS is an assistant group leader in the System Design and Evaluation Group and is currently responsible for technical management of the Termi-

nal Air Traffic Control Automation (TATCA) Program. For 14 years, he was involved in the development of collision avoidance systems. His work on the Traffic Alert and Collision Avoidance System (TCAS) resulted in contributions to aircraft altitude tracking techniques, analysis of computer logic, human-subject flight testing, and the analysis of pilot visual acquisition performance. John has served as a consultant to the National Transportation Traffic Board in the investigation of midair collisions, including the Cerritos accident. He has a B.S. degree in physics from the Georgia Institute of Technology, and an S.M. degree in aeronautical engineering from MIT.



JERRY D. WELCH received an S.B. and S.M. in electrical engineering from MIT and a Ph.D. in electrical engineering from Northeastern University. After joining Lincoln

Laboratory in 1962, Jerry worked on laser radars, integrated microwave receiver circuits, and microwave semiconductor devices. In 1970 he was part of the team that initiated the development of the FAA's Mode S Air Traffic Control Radar Beacon System (ATCRBS). After the Mode S design was completed in 1975, he began work on the Traffic Alert and Collision Avoidance System (TCAS), and became leader of the Laboratory's program to develop surveillance techniques for air-to-air, beacon-based collision avoidance. In 1987 he became manager of the Terminal Air Traffic Control Automation (TATCA) Program at the Laboratory. Jerry is a member of IEEE, AIAA, and ATCA.