W.H. Harman

TCAS: A System for Preventing Midair Collisions

To reduce the possibility of midair collisions, the Federal Aviation Administration has developed the Traffic Alert and Collision Avoidance System, or TCAS. This airborne system senses the presence of nearby aircraft by interrogating the transponders carried by these aircraft. When TCAS senses that a nearby aircraft is a possible collision threat, TCAS issues a traffic advisory to the pilot, indicating the presence and location of the other aircraft. If the encounter becomes hazardous, TCAS issues a maneuver advisory.

When two aircraft collide in midair, the consequences are tragic. Fortunately, such collisions are rare in today's airspace because a number of mechanisms insure safe separation between aircraft—primarily the ground-based system of air traffic control (ATC). To improve on the safety record of the existing systems, the Federal Aviation Administration (FAA) has continued to explore the possibility of adding an airborne collision avoidance system that would serve as a backup to all current provisions.

The Traffic Alert and Collision Avoidance System, or TCAS, is the result of a development program, sponsored by the FAA, that has extended over more than a decade, and is now entering a period of full-scale nationwide implementation. As a result of the development effort, the TCAS design provides reliable air-to-air surveillance, and has been enthusiastically accepted by pilots and other in the aviation community. A federal law passed in 1987 requires that all carrier aircraft install a TCAS by the end of 1991.

The Concept of TCAS

TCAS is an airborne electronics system that employs radio signals for surveillance of nearby aircraft, and in dangerous encounters warns the aircraft pilot by means of cockpit displays and auditory alarms. To detect the presence of nearby aircraft, TCAS transmits interrogations at a steady rate, nominally once per second, and employs a receiver to detect replies to these interrogations from the transponders on nearby aircraft (Fig. 1). The resulting surveillance consists of three components.

Range, or the distance between the two aircraft, is determined by the time between the transmission of the interrogation and the reception of the reply.

Altitude of the other aircraft is determined by reading an altitude code included in the reply. Altitude is measured barometrically on board the other aircraft and is transmitted by digital code to the TCAS aircraft.

Azimuth, or bearing of the other aircraft with respect to the nose of the TCAS aircraft, is obtained by a direction-finding antenna on the TCAS aircraft.

TCAS uses the interrogation/reply technique to detect the presence and measure the location of all aircraft within 15 miles. All of this information is not displayed at all times to the pilot, however. TCAS activates the display only in a dangerous situation, such as when another aircraft is close or when a distant aircraft is closing rapidly.

Figure 2 shows an example of a TCAS display. The traffic-advisory-display circle denotes a range of 2 nmi. The TCAS aircraft is at the center of the display, and the nose of the aircraft corresponds to the 12-o'clock position. In the center column of this figure TCAS informs the pilot that a nearby aircraft is at 11 o'clock (that is, 30° to the left of straight ahead), at a range of slightly more than 2 nmi, and at an altitude of 200 ft below. The display shows altitude digi-

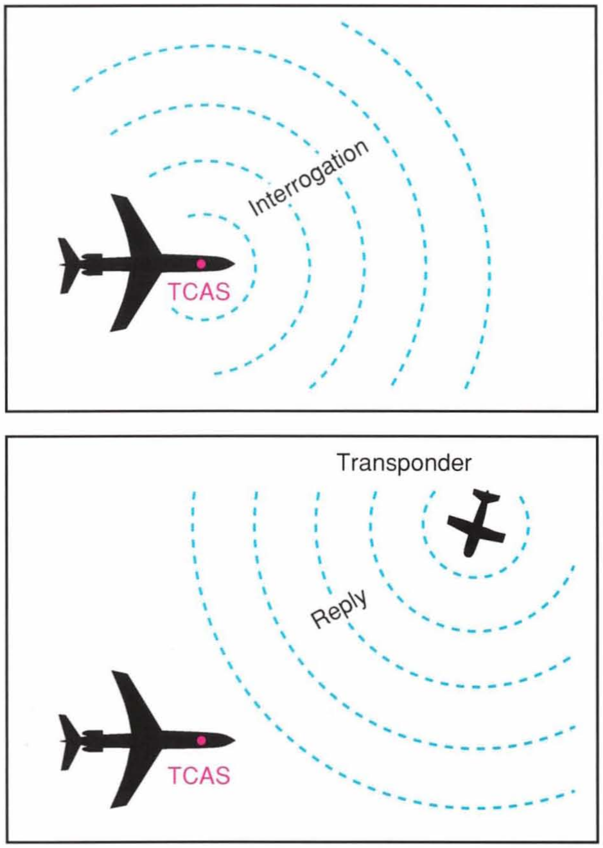

Fig. 1—Air-to-air surveillance obtained through interrogation and reply. Range is determined from the elapsed time between interrogation and reply. Altitude is obtained by reading an altitude code included in the reply. Azimuth angle is measured by a direction-finding antenna on the TCAS aircraft.

tally, in multiples of 100 ft. If the encounter continues to become more dangerous, TCAS will advise the pilot to begin a vertical resolution maneuver, such as climbing. The advisory will be displayed with an accompanying auditory alarm, and a recorded voice will say *climb*. If the other aircraft is also equipped with TCAS, the two TCAS units will exchange coordination messages to insure that the maneuver advisories issued on each aircraft are compatible. This strategy prevents both TCAS aircraft from climbing or descending at the same time.

Air Traffic Control Transponders

TCAS air-to-air surveillance depends on the presence of an air traffic control (ATC) transponder in the other aircraft. Transponders are small receiver-transmitters that, when interrogated with particular radio pulses, transmit a pulsed reply. Currently, transponders are standard ATC equipment, and they form the airborne portion of the ATC surveillance system (in which the ground-based part of the system is a network of radar interrogator-receivers). The FAA requires transponders on all air carrier aircraft, all aircraft under ATC control, and all aircraft flying in certain major terminal areas. Many small aircraft, for additional safety and visibility, also use transponders.

A major advantage of the TCAS concept, compared with other system proposals considered by the FAA, is its ability to interoperate with the standard ATC transponders. Other collision avoidance systems have been proposed that would have required the installation of a special transponder on each aircraft. A collision avoidance system with a special transponder would be easier to design, but the cost and effort to install special transponders on all aircraft was a disadvantage. Given that the FAA has been actively promoting the installation of ATC transponders for many years, the installation of an additional transponder for use only in collision avoidance was viewed as undesirable. Furthermore, the installation of a special transponder would not provide an aircraft with a collision avoidance display. It would make that aircraft visible only to collision-avoidance-equipped aircraft. Because of this limitation, owners of small aircraft would have little inherent motivation to purchase such a transponder. In the TCAS system, however, the ATC transponder serves both purposes. All existing transponder-equipped aircraft will be visible to any TCAS-equipped aircraft, and both the ATC radar surveillance and TCAS will benefit from current efforts to equip all aircraft with transponders.

During the development of TCAS, a system-design ground rule stated that no modifications could be made to existing transponders. TCAS was required to transmit an interrogation to which ATC transponders would reliably reply. It was also required to receive the standard transponder reply without any modifications.

Not all aircraft have ATC transponders at this

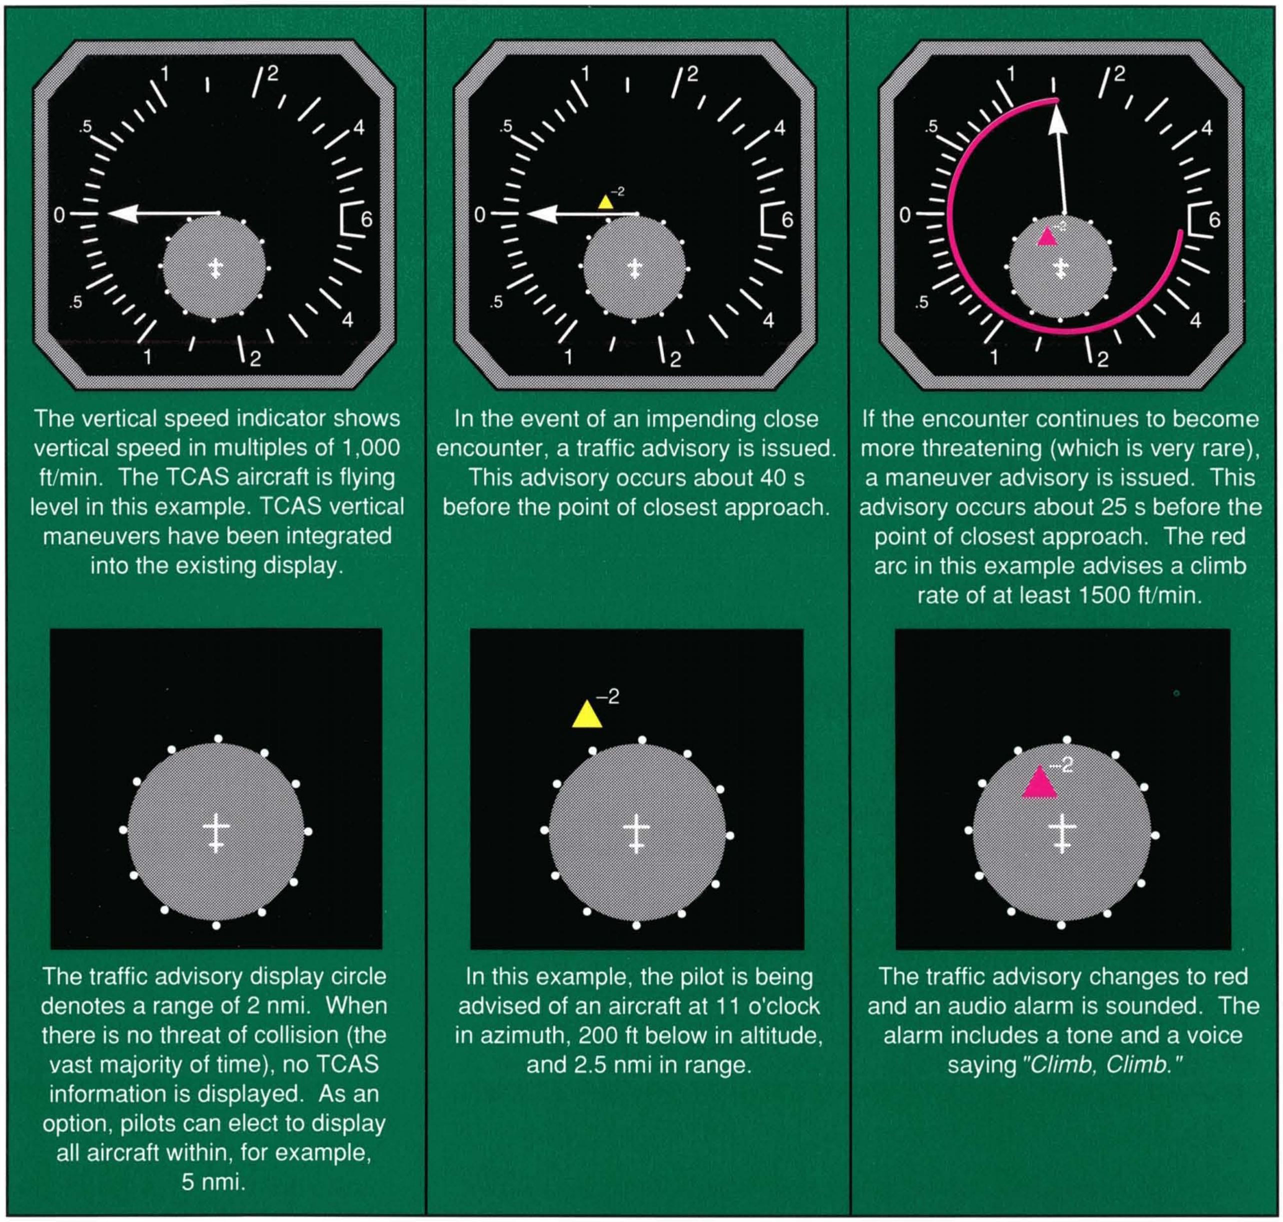

The vertical speed indicator shows vertical speed in multiples of 1,000 ft/min. The TCAS aircraft is flying level in this example. TCAS vertical maneuvers have been integrated into the existing display.

In the event of an impending close encounter, a traffic advisory is issued. This advisory occurs about 40 s before the point of closest approach.

If the encounter continues to become more threatening (which is very rare), a maneuver advisory is issued. This advisory occurs about 25 s before the point of closest approach. The red arc in this example advises a climb rate of at least 1500 ft/min.

The traffic advisory display circle denotes a range of 2 nmi. When there is no threat of collision (the vast majority of time), no TCAS information is displayed. As an option, pilots can elect to display all aircraft within, for example, 5 nmi.

In this example, the pilot is being advised of an aircraft at 11 o'clock in azimuth, 200 ft below in altitude, and 2.5 nmi in range.

The traffic advisory changes to red and an audio alarm is sounded. The alarm includes a tone and a voice saying *"Climb, Climb."*

Fig. 2—TCAS displays and alarms.

time. To assess this limitation on the effectiveness of TCAS, an estimate was made of the percentage of aircraft that are equipped with a transponder. Because aircraft without transponders are primarily small aircraft that spend the majority of the time on the ground, the analysis properly weighted the types of airborne aircraft that a TCAS would encounter. For a TCAS-equipped airliner involved in a close encounter with another aircraft, the percentage of cases in which the other aircraft is transponder equipped was estimated to be 92% in 1983 [1]. This percentage is expected to grow with time.

A similar limitation is the percentage of aircraft that have an ATC transponder without the altitude-reporting capability. The percentage of encounters in which the other aircraft is equipped with an altitude-reporting transponder is 61% [1]. If the other aircraft has a transponder without altitude reporting, TCAS cannot provide a vertical-maneuver advisory, but it can provide an alert that includes a traffic advisory.

This alert would warn the pilot of the presence of the nearby aircraft, and indicate the direction and range, which will help pilots visually locate the other aircraft and thus significantly enhance the effectiveness of visual separation.

Because of the important role of transponders in the ATC system, the percentage of transponder-equipped aircraft is steadily increasing. In 1988, the FAA issued new regulations to increase the regions of airspace in which altitude-reporting transponders are required. The percentages of equipped aircraft given above were derived in 1983. Current percentages are higher and, as TCAS becomes operational, they will continue to increase in the years to come.

Classes of TCAS Equipment

Three classes of TCAS equipment have been identified by the FAA.

(1) *TCAS I*, intended for smaller aircraft, provides traffic advisories but does not provide maneuver advisories.

(2) *TCAS II*, intended for large air carrier aircraft, provides traffic advisories and vertical-maneuver advisories.

(3) *TCAS III* provides horizontal-maneuver advisories in addition to the capabilities of TCAS II.

TCAS I is usually associated with aircraft that have significantly lower airspeeds than air carriers. As a result, the air-to-air surveillance range need not be as great. Thus the power level of the interrogation transmitter can be less, which results in a lower-cost unit. TCAS I is also appropriate for helicopters.

TCAS II, the main subject of this article, was the focus of the TCAS development program. The development of TCAS II is completed at this time; the system now enters the period of operational use. TCAS II standards have been adopted in the United States by the Radio Technical Commission for Aeronautics (RTCA), and adopted internationally by the International Civil Aviation Organization [2, 3]. A TCAS II airborne unit costs approximately $70,000 (not including installation).

Lincoln Laboratory developed TCAS I technology to the point that standardized technical characteristics were adopted by RCTA [4]. Whether TCAS I will now proceed directly to operational use is not clear. If a manufacturer chooses to build a full-power TCAS I (a subset of a TCAS II), then additional TCAS I development or testing is not necessary, since TCAS II has been thoroughly tested under both experimental and operational conditions. Even though a manufacturer could build a reduced-power, lower-cost TCAS I, manufacturers have not developed reduced-power TCAS I products since the RTCA standards were adopted in 1987. To stimulate TCAS I production, the FAA initiated a follow-on program for TCAS I operational testing. Product development of reduced-power TCAS I is expected to be deferred until after the completion of the program.

TCAS III is still under active development. Unlike TCAS II, the horizontal advisories in TCAS III require a much more accurate surveillance in bearing. For example, if another aircraft is passing on the left, the TCAS III surveillance must be sufficiently accurate to indicate the left-right sense of the relative motion. If surveillance inaccuracies cause the track to appear to be passing on the right, an incorrect *turn left* advisory might be issued, which would be a serious error. The TCAS III program developed a monopulse antenna capable of providing the required surveillance accuracy. The goal is an antenna large enough and well located on the TCAS III aircraft (away from reflecting objects) to achieve the necessary bearing accuracy.

Development Challenges

Lincoln Laboratory faced many significant challenges during the development of TCAS, including unwanted radio reflections, interference, and other problems described in this article. At times, particularly in the early years, it did not seem feasible for an airborne sensor to track aircraft reliably with the ATC transponders. Most of the difficulties can be attributed to the transponder radio-signal formats that were standardized just after World War II for ground-based surveillance radar. These signal formats were not intended for use in air-to-air transmissions. For example, echoes occur in air-to-air surveillance because the radio signal reflects from the ground or ocean beneath the two aircraft. The echo is superimposed directly on the received signal, which causes garbling. The garbling affects both interrogations and replies, and often makes radio receptions unusable. During the development of TCAS, Lincoln Laboratory conducted airborne measurements to assess the extent of the degradation due to garbling, and a TCAS design was developed with many provisions to overcome the effects of echoes.

The use of existing transponders in TCAS also led to issues of possible radio interference. Since TCAS transmits interrogations and replies in the same frequency bands as ATC surveillance radar (1030 MHz for interrogations and 1090 MHz for replies), the developmental program had to insure that TCAS would not interfere with ATC radar.

TCAS consists of two major subsystems: (1) air-to-air surveillance and (2) triggering of alarms. Since air-to-air surveillance was the primary activity in TCAS development at Lincoln Laboratory, the following description covers it in greater detail. The Mitre Corporation had the corresponding role in the development of TCAS alarm triggering.

Air-to-Air Surveillance

TCAS conducts air-to-air surveillance in one of two modes, according to the type of transponder in the other aircraft under surveillance. If the other aircraft is equipped with a Mode-S transponder, then TCAS conducts the air-to-air surveillance in Mode S. Otherwise, TCAS conducts the air-to-air surveillance in Mode C. The mode designations distinguish between the newly standardized Mode S and the seven modes used for many years prior to the development of Mode S (see the accompanying article by Vincent Orlando titled "Mode S Beacon Radar System").

A Mode-C interrogation is an all-call, and all aircraft that receive the interrogation transmit a reply. A Mode-C reply contains a digital code that reports the altitude of the replying aircraft. The fact that multiple aircraft will reply to a single interrogation leads to a problem with synchronous garbling of the replies. In Mode S, the interrogations are addressed selectively so that only one aircraft replies to a given interrogation. As a result of selective interrogation, Mode S avoids the synchronous garble that significantly limits Mode C. A TCAS installation includes a Mode-S transponder, and therefore TCAS-to-TCAS surveillance is conducted in Mode S. Because aircraft equipped with Mode-S transponders are tracked by TCAS in Mode S, they need not reply to Mode-C interrogations, as explained below.

The synchronous garbling of replies in Mode C was one of the main challenges in the development of TCAS. Echoes from the ground were another problem area. Other issues that required development effort included

(1) specification of an interrogation power level high enough to provide reliable air-to-air surveillance while low enough not to interfere with ATC radar,
(2) Mode-S surveillance algorithms that provide the interrogation address of each nearby aircraft so that selective interrogation can proceed,
(3) angle-of-arrival antenna development and accuracy assessment, and
(4) surveillance of aircraft equipped with Mode-C transponders without altitude reporting.

Multipath

The reflection of radio signals from the ground or water over which the aircraft are flying (known as *multipath*) was recognized from the beginning as a potential difficulty in TCAS. At the outset of the program in 1975, we made airborne measurements to characterize these reflections. While equipment to measure multipath was being designed and built, efforts were made to assess the phenomenon from existing information. Engineers who had been involved with air-to-air TACAN (Tactical Air Navigation, a system that employs pulsed interrogations and replies similar to TCAS), and with testing of other airborne collision avoidance techniques, suggested that multipath would be a serious

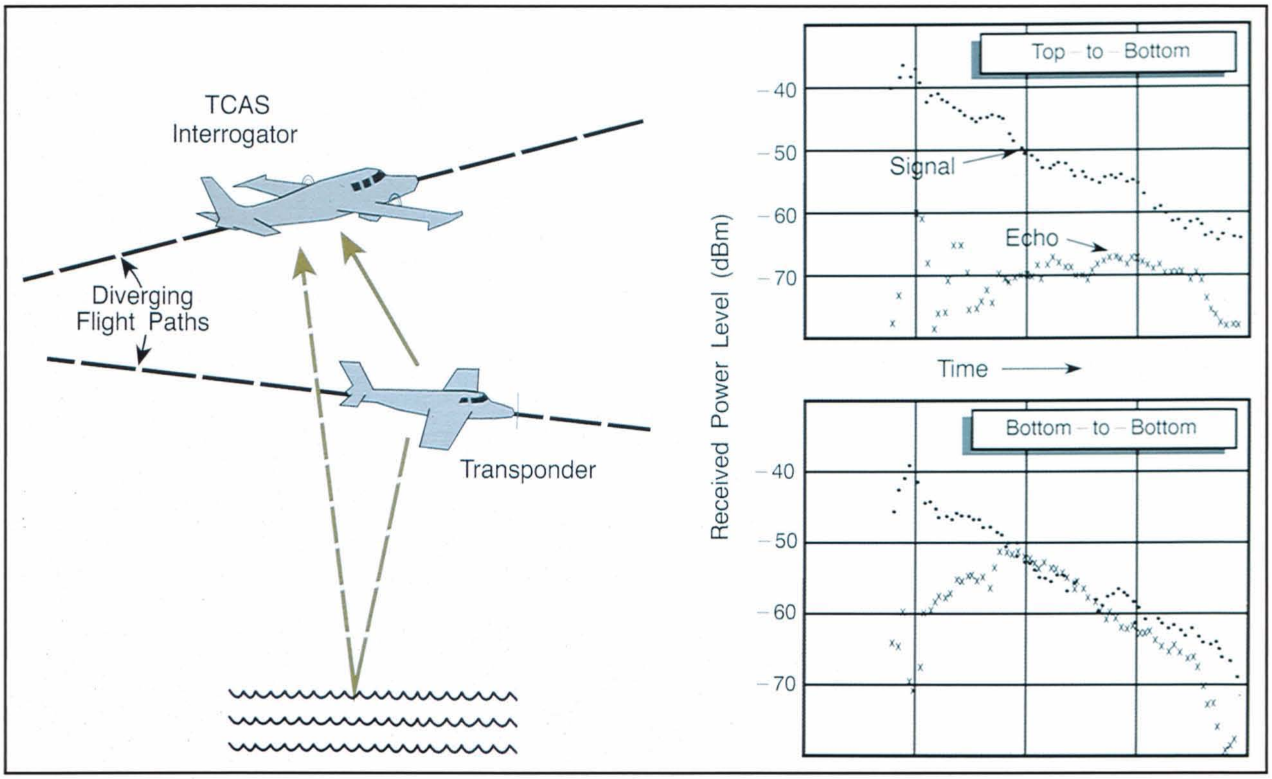

Fig. 3—Measurements by two aircraft of air-to-air multipath. Gradually diverging flight paths were used to determine multipath as a function of range. The action was repeated at different altitudes and over different surfaces. The worst-case echoes shown here were obtained over ocean on a calm day.

difficulty in TCAS, particularly because of the constraint to use existing transponders and existing signal formats. In fact, some knowledgeable people contended that the TCAS concept could not be made into a practical system because of multipath. Multipath was indeed a major disturbance, but a TCAS design was developed to tolerate multipath and provide reliable surveillance in the multipath environment.

Two instrumented aircraft conducted the multipath measurements; one aircraft transmitted a single 0.5-μs pulse and the other aircraft received the pulse along with its echoes. We designed the apparatus to transmit a series of pulses that synchronized the receiving equipment and established a pattern of consistency to distinguish between multipath and interference from other transmitters [5]. To obtain results as a function of range, the two aircraft were flown at the same altitude on paths that slowly diverged. The procedure was then repeated at a number of altitudes to obtain results as a function of altitude. The multipath measurements were conducted over a number of locations, including ocean, cities, rural New England areas, tree-covered mountainous areas, frozen lakes, deserts, and populated areas in the Los Angeles basin.

The multipath test procedure indicated that echoes are detectable in every case. The most common form of echo observed was a delayed replica of the directly received pulse; the delay time would agree with the additional path length expected if the reflecting surface were a flat, level plane. The amplitude of the echo signal varied more than the directly received signal, and the short-term variability extended over 15 dB with a distribution that agreed with a Rayleigh model. The variations were uncorrelated in repeated measurements 20 times per second. The mean value of the echo power also varied slowly as the range between the two aircraft changed [5].

Figure 3 shows the results obtained from a flight over the ocean near Cape Cod Bay on a calm day. Results were obtained for both top and bottom antenna configurations on each aircraft. The power of the direct signal diminished gradually as range increased monotonically, and was approximately the same for the two antenna combinations shown. Figure 3 includes two of the four antenna combinations tested. The power of the multipath signal, on the other hand, demonstrated a different variation as a function of range, and as a result the multipath-to-signal ratio is not constant. The worst multipath-to-signal ratio occurred at a middle value of range with the bottom-to-bottom antenna combination, where the multipath was stronger than the direct signal.

The multipath levels in Fig. 3, obtained over calm ocean, were the highest levels observed among all of the measurement locations. Measurements made over rough sea in this same location indicated that multipath power became consistently lower as sea state increased.

A radio signal that reflects from the ground or ocean necessarily loses power because of both absorption and scattering. Figure 4 illustrates a probable physical mechanism that causes the received multipath power to exceed the direct power. The antenna gain patterns typical for bottom-mounted transponder antenna installations generally have a region of maximum gain pointing in a downward direction by 20° to 30° [6, 7]. The antennas are simple monopoles for which the aircraft fuselage is the ground plane. Therefore, relative to the antenna gain that affects the multipath, the antenna gain affecting the signal transmitted to a co-altitude aircraft is less. The gain difference has an effect both at the transmitting aircraft and at the receiving aircraft. These gain differences probably account for the high-multipath receptions in some conditions.

For a top-mounted antenna the gain pattern is reversed; the maximum gain is in an upward direction. Thus multipath power would be expected to be much less if one of the two antennas is top mounted, and still less if both antennas are top mounted. The measurements agree with these expectations.

Using top antennas is an obvious step in designing TCAS to tolerate multipath. Most existing transponders use bottom antennas, but TCAS was designed to employ both a top and a bottom antenna, and rely mainly on the top antenna. Because of the antenna-gain patterns, signal strength for a top antenna improves when the other aircraft is at a higher altitude, and diminishes when the other aircraft is at a lower altitude. Surveillance reliability thus tends to be a function of the elevation angle between the two aircraft. In other words, reliability is best when TCAS is looking up, and worst when TCAS is looking down. The bottom TCAS antenna fills the gap in the small region of negative elevation angles where the top antenna is at a disadvantage.

Dynamic receiver thresholding is another

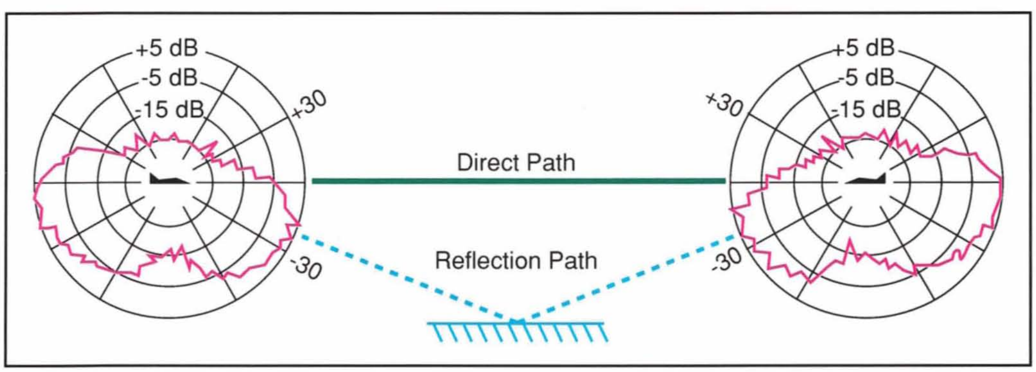

Fig. 4—The effect of antenna patterns on air-to-air multipath by two bottom-mounted monopole antennas. Bottom-mounted monopole antennas unfortunately tend to boost the strength of the unwanted echoes.

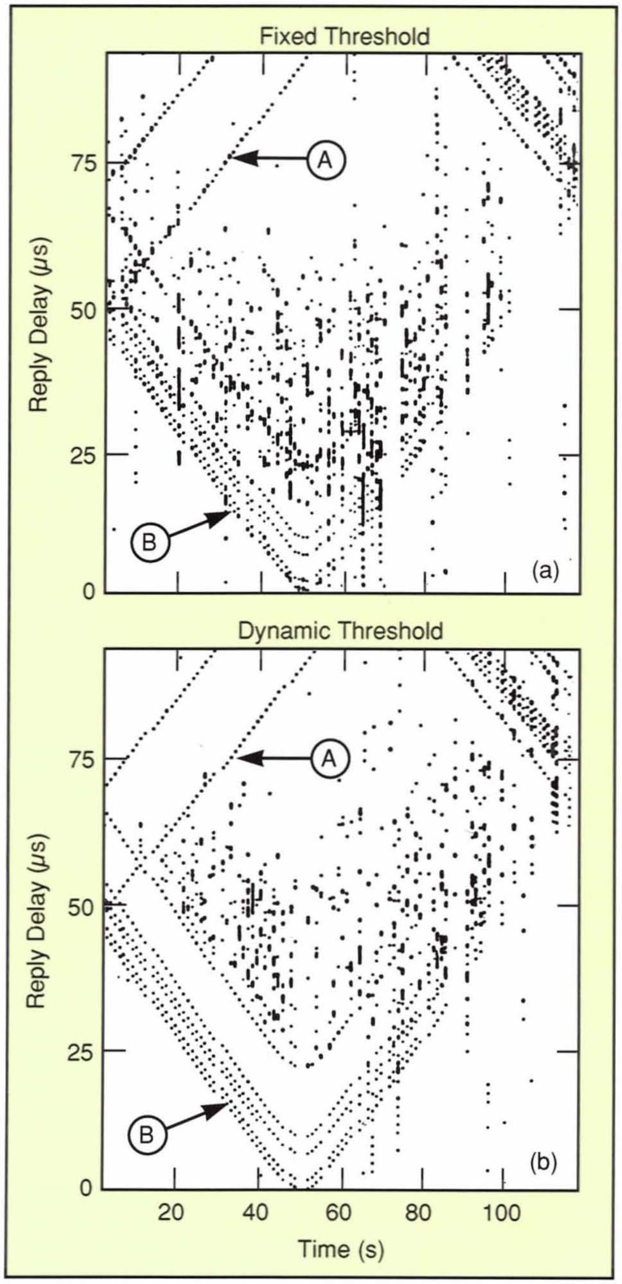

Fig. 5—(a) Reception due to fixed thresholding. (b) Improved reception due to dynamic thresholding. The detection of numerous weak echoes that follow a strong pulse is often eliminated by raising the receiver threshold immediately after the strong pulse is received.

technique used to combat multipath. This simple technique raises receiver threshold immediately after receiving a strong signal. Whenever a reply is received at a power level much higher than the nominal receiver threshold, the first received pulse triggers the dynamic thresholding for a period equal to one reply duration (about 20 μs). The threshold will be raised to a level –9 dB relative to the strength of the pulse that triggered it. Thus any echoes received during this period will be below threshold and will not be declared as replies, except for the rare echoes stronger than the level of –9 dB.

The benefits of dynamic thresholding are, unfortunately, accompanied by some degradation in the form of loss of replies. When a weak signal and a strong signal are received simultaneously, the raised threshold triggered by the strong signal may eliminate the weak signal. For this reason, dynamic thresholding is not ordinarily used in ground-based Mode-C receivers. In TCAS, however, this tradeoff is more favorable because of the action of a technique called whisper-shout, which is described below. The whisper-shout technique groups together replies of approximately the same power levels.

Airborne measurements were conducted to assess this thresholding technique and other techniques. Figure 5 shows a direct comparison of pulse receptions with and without dynamic thresholding. To obtain this data, an experimental TCAS alternated rapidly between two designs for comparison. Thus a nearby aircraft is interrogated and processed in two different ways. The plots show pulse detections as a function of time; the vertical scale gives the arrival time of each received pulse. One nearby aircraft appears as a group of pulses. The first pulse indicates the range of the aircraft and the others give the reported altitude code. The figure shows the pulse structure for two aircraft labeled *A* and *B*. The aircraft designated by *A* has passed the point of closest approach and is now diverging. Since its transponder is not equipped with an encoding altimeter, it replies with just two framing pulses. The aircraft designated by *B* has passed close to the TCAS aircraft at about the middle of the plot. Its replies contain altitude reports, as indicated by the presence of data pulses between the two framing pulses. The extra pulses, particularly evident in Fig. 5(a), are largely due to multipath. Figure 5(b) indicates the use of dynamic thresholding. During most of the encounter with aircraft *B*, the multipath was consistently lower than the threshold of –9 dB and was thus eliminated. When the threshold was restored after each reply, the multipath immediately reappeared. From the series of airborne measurements we concluded that dynamic thresholding provides a large net benefit in TCAS.

The *whisper-shout* technique (described below) is a third means of combating multipath. Although whisper-shout was originally intended to mitigate synchronous garble, it also reduces multipath disturbances on the interrogation link. Whisper-shout causes the interrogation link for the transponder receiver to operate near receiver threshold, which eliminates multipath except for rare instances when the multipath is nearly as strong as the direct interrogation.

Another technique to combat multipath is applied at the track level. The TCAS computer forms tracks from all of the received replies, and each track ideally corresponds to one aircraft. Regardless of the other techniques to eliminate multipath, a track occasionally forms from echoes. The echo track and the nearby valid track typically exist simultaneously. The false track usually has the same altitude as the valid track but longer range. Sometimes two false tracks accompany a valid track. This possibility corresponds to (1) a single delay, in either the interrogation link or the reply link, with a direct transmission in the other link; and (2) a double delay—that is, a reflection in both interrogation and reply. A program in the TCAS computer searches among all the tracks to identify suspicious pairs or triples that have the altitude and ranges consistent with a simple multipath calculation. When such tracks are discovered, TCAS flags the longer-range tracks as suspicious; these tracks are not used in the pilot display. The parameters of this multipath-elimination algorithm have been carefully selected on the basis of airborne measurements. The resulting performance effectively rejects false tracks while it retains valid tracks.

Synchronous Garble

When two or more aircraft under Mode-C surveillance have approximately the same range from the TCAS aircraft, their received replies overlap in time. This phenomenon, called synchronous garble, persists during repeated interrogations until the ranges diverge. Figure 6 illustrates, with respect to a particular aircraft target of interest, how other aircraft that are nearer or farther by about 1.7 nmi will contribute overlapping replies. This range band is substantial and in high-density airspace gives rise to an excessive number of overlaps. For example, if the density of aircraft is 0.1 aircraft per square nmi (a value typical of Los Angeles today [8, 9]), and if the aircraft of interest is at a range of 5 nmi, then the average number of overlapping replies from other aircraft will be 11. This number is too large for reliable reception and decoding. The TCAS receiver-decoder is capable of decoding the reply of interest when overlapped by one or two additional replies. For TCAS to operate in high-density areas such as Los Angeles, synchronous garble must be reduced by an order of magnitude.

The TCAS design includes several techniques for reducing synchronous garble: directional interrogation, whisper-shout, and Mode-C-only interrogation. Figure 6 indicates that directional interrogation will directly reduce synchronous garble. Any single directional interrogation will interrogate only the aircraft that are within the extent of the beamwidth. For example, each of the beams in a four-beam antenna can ideally be as narrow as 90°, which reduces the synchronous garble by a factor of 4. In practice, however, wider beams are necessary, and the achievable improvement factor is approximately 2.4 [9].

Whisper-shout consists of a sequence of interrogations in a small fraction of a second instead of a single interrogation each second (the nominal surveillance-update period). The interrogation sequence begins at low power and increases to the final full power. The objective of the technique is to partition the replies so that only a small subset is received during any one reception period. Of course, at least one reply must be received from each aircraft under surveillance. Whisper-shout accomplishes this result by adding to the standard interrogation a suppression pulse 2 μs earlier. The early pulse

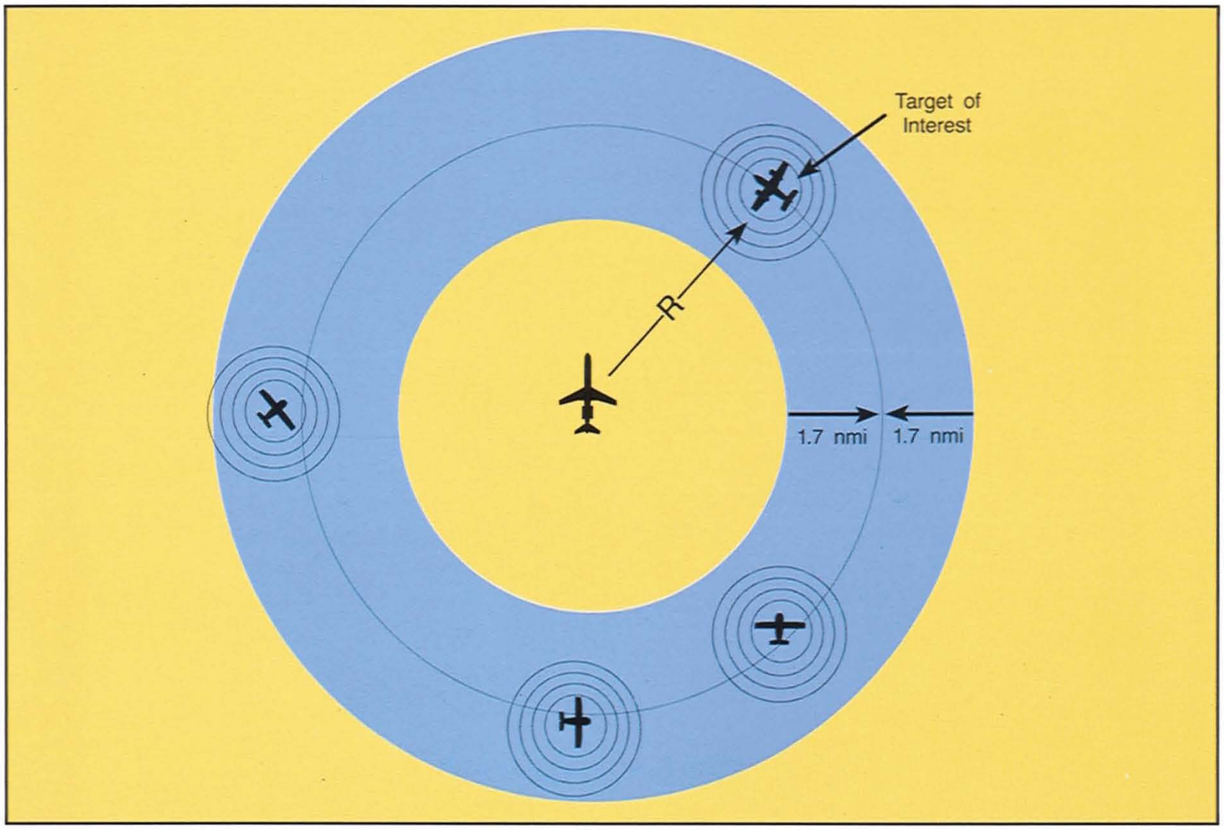

Fig. 6—Synchronous garble due to replies from multiple aircraft. Receptions from a particular target of interest can be overlapped by receptions from other aircraft at nearly the same range. For example, if the target of interest is at a range of 5 nmi, and if the density is 0.1 aircraft per square nmi, then the average number of interfering replies is 11.

suppresses the transponders that have already replied to an interrogation earlier in the sequence. The initial whisper-shout design consisted of four power levels; subsequently this number was increased to 24 power levels to provide surveillance capability for the highest-density airspace.

As we gained more experience with whisper-shout, we realized how effective this technique is. Originally intended to combat synchronous garble, whisper-shout reduces multipath on the interrogation link, as described above. Whisper-shout also benefited the technique of dynamic receiver thresholding, which otherwise might not have been practical because of the problem of blanking some replies in the presence of other stronger replies. Since whisper-shout groups replies at the same power level, it minimizes the blanking that would otherwise occur, with the result that dynamic thresholding has become practical in TCAS.

The other technique for combating synchronous garble, namely Mode-C-only interrogation, is based on the fact that no synchronous garble exists in Mode-S surveillance. For this reason, surveillance of aircraft equipped with Mode-S transponders is done in Mode S, and Mode-C surveillance is not necessary. The standard Mode-C interrogation was modified so that Mode-S transponders do not reply. Specifically, another pulse was added 2 μs after the standard interrogation. Mode-S transponders have a corresponding capability to recognize this additional pulse and not reply. Therefore, this modification removes the population of Mode-S-equipped aircraft from the synchronous-garble environment. It will become increasingly effective in reducing synchronous garble as Mode-S transponders become more widely used.

Power

TCAS interrogations must be powerful enough for reliable air-to-air surveillance yet weak enough not to interfere with ATC radar. To assure that the power is high enough, the path loss associated with the range to the other aircraft must be considered, along with the significant power deviations caused by antenna gain patterns and transponder sensitivity.

The nominal power level for ATC transponders is 250 W radiated, and the nominal receiver sensitivity is –74 dBm. These values are a point of departure for TCAS. If TCAS were to interrogate at 250 W and have a receiver sensitivity of –74 dBm, then the interrogation and reply links would be balanced, and the transmitter and receiver would be about as complex as transponders. While worst-case conditions (e.g., worst possible power deviations due to antenna patterns) would require TCAS to transmit more power than 250 W, a probabilistic analysis indicates that 250 W is sufficient in a very high percentage of cases [10]. Consequently, the 250-W level was adopted initially as the power level for TCAS interrogations, along with the corresponding receiver sensitivity of –74 dBm. It remained to determine whether this power level, operated at a surveillance update rate of once per second, was low enough to assure noninterference with ATC radar.

Initially, TCAS was designed for operation in low to medium densities of aircraft. At that time, the FAA was developing a ground-based collision avoidance system (called ATARS) that was intended for use in the high-density areas. The role of TCAS (then called BCAS) was to provide collision avoidance throughout all of the airspace away from the high-density terminal areas. In 1980, partly on the basis of the successful BCAS program, the FAA made a major change in the overall system concept: ground-based collision avoidance was not pursued, and the scope of airborne collision avoidance was redefined so that the capacity would be increased to cover all airspace (including the highest densities). ATC radar interference from TCAS was one of the major issues in the development program that followed this decision.

The TCAS design includes an interference-limiting function that monitors the interference conditions in the local airspace in which the TCAS aircraft is currently flying. The resulting interference density is used to calculate a maximum interrogation rate-power product that the TCAS transmitter must not exceed. Every TCAS is required to implement this function, by monitoring the environment as specified and by constraining its own interrogation transmitter according to specified formulas. The interference-limiting formulas were derived analytically at Lincoln Laboratory [9], and were subsequently validated by a comprehensive simulation conducted at the Electromagnetic Compatability Analysis Center [11].

The complementary relationship between high aircraft density and the reduced airspeeds in this airspace is one reason that effective surveillance can be provided within this power limiting. When a TCAS-equipped aircraft flies in low-density en route airspace, TCAS operates at full power level, including a full whisper-shout sequence. As the TCAS aircraft flies into airspace of increasing aircraft density, the rate-power product must be reduced at a certain point. This point is reached in terminal airspace in major metropolitan areas, where aircraft speeds are considerably less than in en route airspace. Lower values of closing speed correspond to a reduced range requirement for providing the needed warning time (about 25 s). The reduced range requirement implies a reduced power requirement, according to a square law.

As a result of these provisions, the original 250-W power specification was adopted in TCAS, which is sufficient power for surveillance at the highest closing speeds in en route airspace. As a TCAS aircraft flies into an area of high aircraft density, the interference-limiting function is triggered. TCAS continues to function with a reduced range capability, which is sufficient for effective collision avoidance in that airspace.

Mode-S Surveillance

A major advantage in Mode-S surveillance is absence of synchronous garble, as a result of the individual interrogation of each aircraft. To interrogate a given aircraft, TCAS must know the aircraft's unique address; acquisition of the address is accomplished by a process of squitter reception. A *squitter* is a spontaneous transmission emitted by a Mode-S transponder, in the format of a reply. All Mode-S transponders transmit squitters at a rate of once per second. The requirement to transmit squitters was adopted as a Mode-S standard mainly because of its usefulness in TCAS air-to-air surveillance, although other applications of squitters have been identified. Because the Mode-S development and the TCAS development occurred simultaneously, incorporating this function in Mode S to support the operation of TCAS was possible.

The TCAS computer monitors all Mode-S receptions on 1090 MHz (the reply band) and examines each to determine which are squitters. For each received squitter, its address is examined to determine whether it is a new address or is one of the addresses already in track. When a new address is received, TCAS transmits an addressed interrogation to learn the range of the aircraft.

To minimize unnecessary interference, and to obtain the highest surveillance capacity within interference limiting, TCAS will not transmit interrogations to aircraft that are far away in altitude or far away in range (based on the available information). If the initial interrogation to a particular aircraft reveals that the range is large, which is common, then TCAS will place that track in *dormancy*. That is, instead of interrogating this aircraft at the nominal rate of once per second, TCAS inhibits interrogations for a period of time calculated according to the time it would take for range to become close enough to warrant steady surveillance. After this time has elapsed, TCAS transmits another interrogation and makes another range measurement. If the other aircraft is still far away, the track is again placed in dormancy.

Mode-S surveillance algorithms were developed to minimize interrogations to distant aircraft that do not present an immediate threat, and at the same time to track reliably any aircraft on a collision course. Chapter 4 in Ref. 10 documents the development of this algorithm and gives performance assessments in high-density airspace.

Angle of Arrival

The third dimension of TCAS surveillance is azimuth angle, which is obtained by an angle-of-arrival antenna. Figure 7 shows the configuration of the simple experimental four-element antenna used on a Cessna 421 aircraft at Lincoln Laboratory. The antenna consists of four 2.5-in monopoles mounted in a square with side length equal to 2.7-in, or one-quarter wavelength. A feed network consisting of four L-band hybrids is mounted beneath the antenna and under the surface of the aircraft. One side of the feed network is connected to the four antenna elements; the other side consists of two cables that lead to the receiver-transmitter unit. The receiver measures the phase between the signals in these two cables. Nominally, 1° of change in azimuth corresponds to 1° of change in phase. Reference 12 gives a detailed description of the angle-of-arrival antenna, along with the associated receiving functions and the measured antenna patterns.

The accuracy of this antenna when installed on an aircraft is approximately 8° rms, which is sufficiently accurate for TCAS purposes. Azimuth measurements feed only the traffic advisory display—that is, they indicate to the pilot the location of the other aircraft. TCAS does not use azimuth measurements for triggering alarms or for correlating replies to form tracks. However, when compared with the accuracy of traffic advisories that are now received by pilots via voice radio, the TCAS advisories are significantly more accurate and are provided much more frequently.

A four-element antenna with bare monopoles is appropriate for a low-speed aircraft, but a high-speed air carrier must use a radome or an antenna of lower profile to reduce drag. Several antenna designs for air carriers have been developed by different manufacturers. Most designs are approximately the same horizontal size and

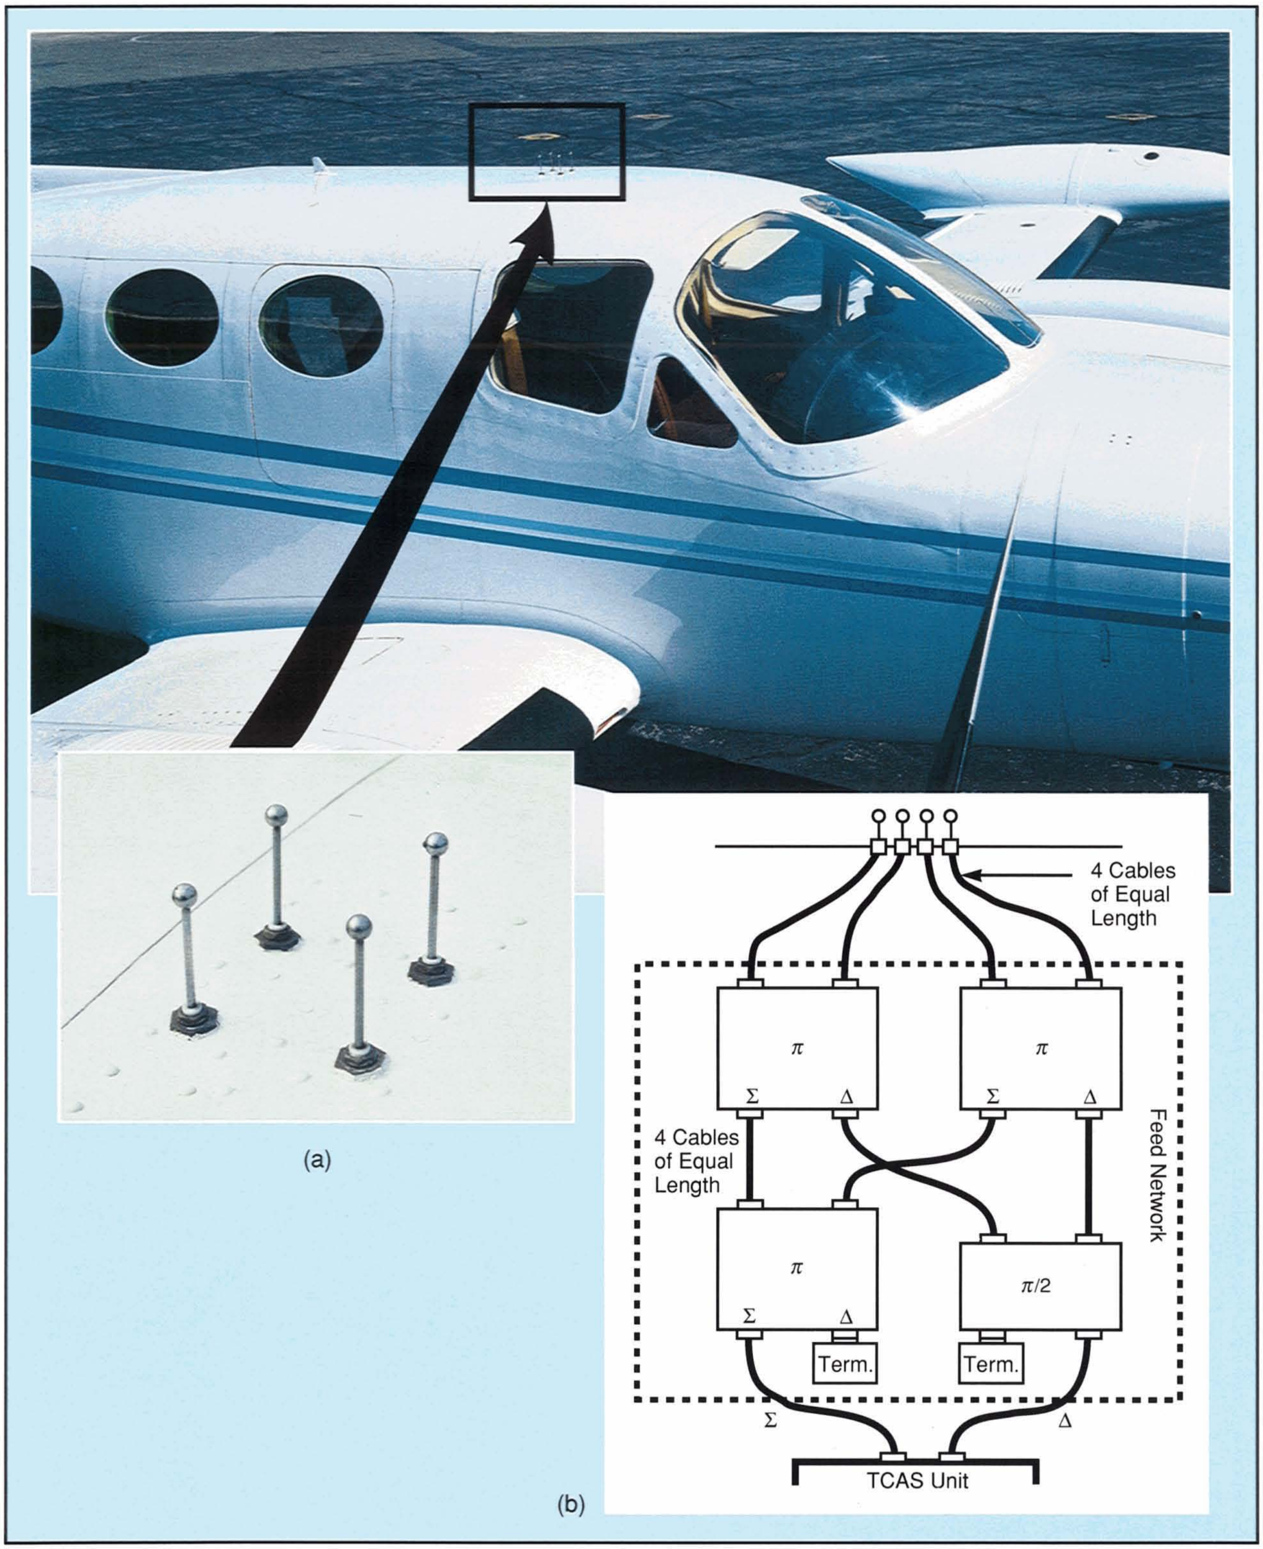

Fig. 7—(a) Angle-of-arrival antenna on the Lincoln Laboratory test aircraft. A 2.7-in square array of four bare monopoles, each 2.5-in in height, provides sufficient azimuth accuracy for TCAS II. For high-speed jet aircraft, the corresponding antenna is one inch high and covered by a radome. (b) Diagram of the antenna feed network.

accuracy as the antenna described here. All the antennas are radome covered and have been made considerably lower through the use of loaded antenna elements.

Non-Altitude-Reporting Transponders

As mentioned above, some aircraft are equipped with ATC transponders but not with reporting altimeters. Although TCAS does not provide a vertical-maneuver advisory for these aircraft, it can measure range and azimuth, and provide the information to the pilot in the form of a traffic advisory. Design of the Mode-C surveillance subsystem for these cases was more difficult because of the absence of altitude information. Altitude information, when available, helps to distinguish among replies from different aircraft. With respect to altitude, all replies from non-altitude-reporting aircraft are indistinguishable and must be sorted according to range and possibly azimuth. The sorting problems are especially difficult during flight in an area of high aircraft density. This problem does not exist in Mode S, because the system identifies each reply with the unique address of the target aircraft.

When the TCAS design for Mode-C surveillance of non-altitude-reporting aircraft was first used in a high-density area, performance was poor. Many false tracks of short duration occurred, along with numerous gaps in the tracks of real aircraft. Performance improvements were developed by adding a third component to the range tracker, to track acceleration as well as range and range rate. Considerable attention was given to the tracker gains (called alpha, beta, and gamma) to make them systematically diminish as the length of a track increases. Special provisions were made to increase gains when the aircraft are coasting through brief but common periods of missed replies. Considerable attention was also given to range-correlation windows. This development was conducted with a detailed data base recorded by an experimental TCAS facility flying over Los Angeles. The resulting tracker design now reliably tracks non-altitude-reporting aircraft, and provides a traffic advisory display of all transponder-equipped aircraft [13].

Fig. 8—Lincoln Laboratory TCAS Experimental Unit. This equipment, installed in a Cessna 421 twin-engine aircraft (left), was operated by Lincoln Laboratory for TCAS experiments, measurements, and the early pilot tests.

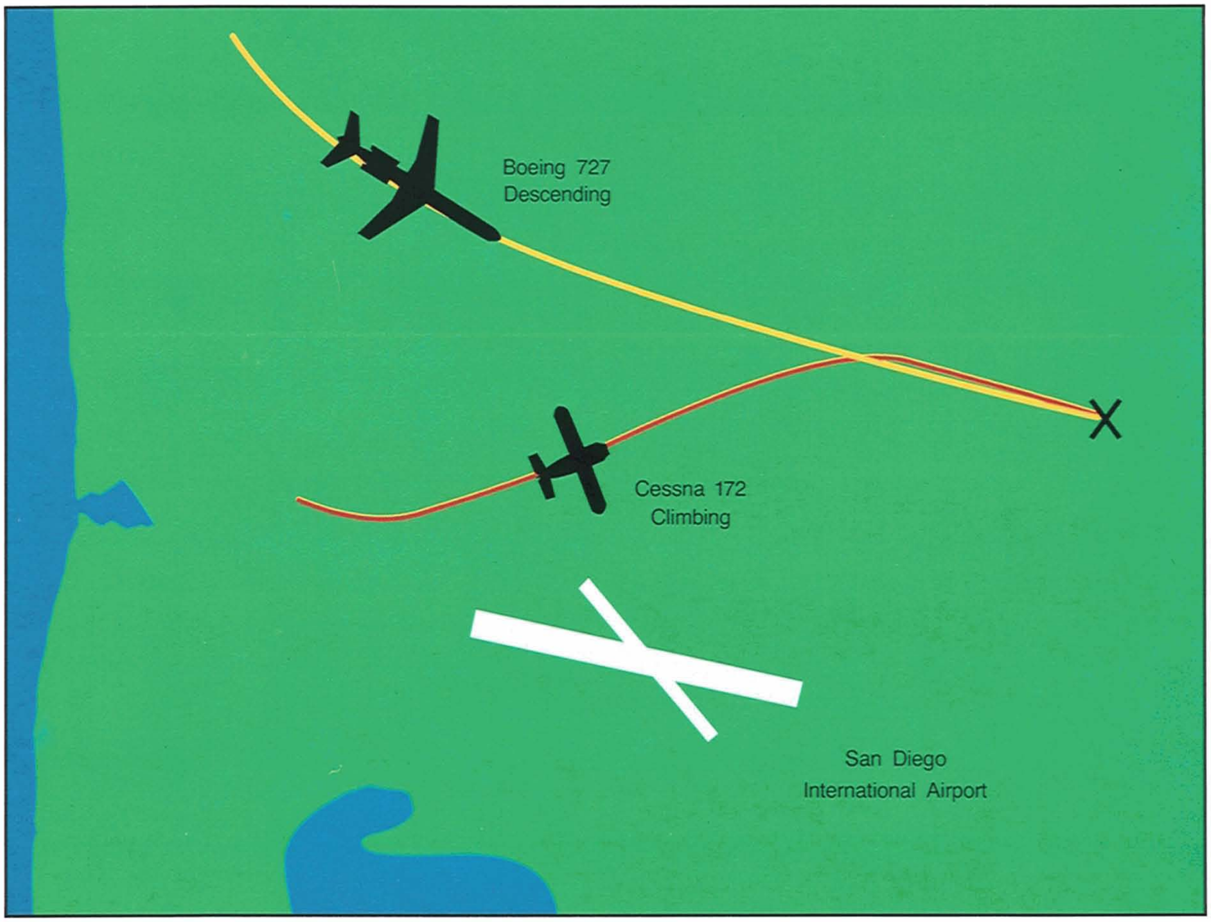

Fig. 9—Geometry of the San Diego collision. The two aircraft were proceeding on the same ground track; the Boeing 727 was descending while the Cessna 172 was climbing.

Airborne Measurements

Airborne measurements such as those described above were central to the development of air-to-air surveillance. Lincoln Laboratory operated a Cessna 421 aircraft and several other aircraft, and built several TCAS experimental units. These units have the capability to record detailed airborne data and to function as real-time TCAS units. Figure 8 shows an experimental TCAS unit installed in a twin-engine Cessna 421 aircraft. Initially, most of the measurements focused on specific issues or phenomena, such as the air-to-air multipath measurements and the improvements brought about by whisper-shout. Later in the program, close encounters were deliberately staged to test the full system and to obtain pilot reactions to the design.

Following the midair collision over San Diego in 1978, the FAA asked whether TCAS could have successfully operated under similar conditions and prevented the collision. On the basis of the geometry of the collision, concerns were expressed about the reliability of TCAS air-to-air surveillance. The collision occurred when a Boeing 727 airliner descended for a landing at San Diego International Airport while a Cessna 172 climbed after departing from the same airport. Figure 9 illustrates how both aircraft were flying east and were on the same ground track. Because the large aircraft was descending while the small aircraft was climbing at a lower speed, an unfavorable relationship existed between the two aircraft for several minutes before the collision. In this geometry, a TCAS installed on the large aircraft would have had a look-down angle

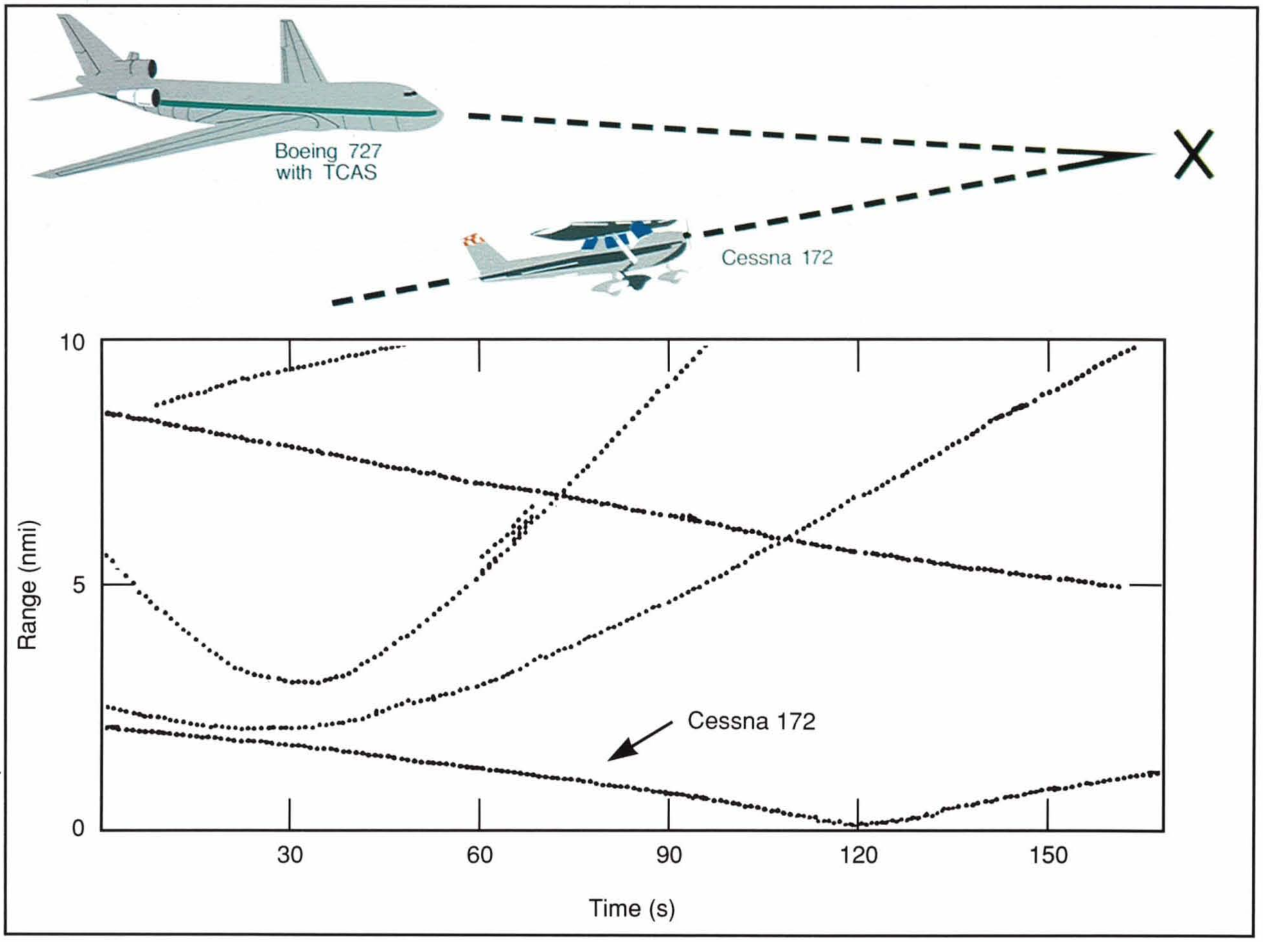

Fig. 10—Results of the reenactment of the San Diego midair collision. Air-to-air surveillance successfully established a track of the Cessna 172 for over two minutes prior to the point of closest approach.

in carrying out surveillance of the other aircraft. The small aircraft was equipped with an altitude-reporting transponder with a bottom-mounted antenna, as is common. Therefore the fuselage of the small aircraft would have shielded the radio signals to and from the transponder. Furthermore, the top-mounted TCAS antenna on the large aircraft (the main antenna for TCAS surveillance) would have had its signals shielded by its own fuselage. The collision also occurred at low altitude, where disturbances caused by ground echoes are maximized.

The above conditions were not qualitatively worse than conditions under which TCAS had previously been routinely tested. For two aircraft on a collision course, an elevation angle of several degrees, positive or negative, is not unusual. Because of the nature of aircraft flight, however, elevation angles beyond 10° are very rare. Among all of the midair collisions in history the values of elevation angle have never exceeded 8° in absolute value [14]. In the San Diego collision, the look-down angle was approximately 6°. Another factor in midair collisions is the closing rate, which tends to be relatively low when elevation angle is either high or low. In the San Diego collision, the closing rate was 75 knots; because of this low value, a few miles of surveillance range would be sufficient to provide adequate warning time. Therefore, it might be argued that the antenna shielding and multipath effects of the San Diego collision would not be unusually challenging for TCAS air-to-air surveillance.

An experiment was undertaken to reenact the

San Diego collision with TCAS. A Boeing 727 operated by the FAA was equipped with one of the Lincoln Laboratory TCAS experimental units, and a Cessna 172 was leased. The Cessna came equipped with an altitude-reporting transponder and a bottom-mounted antenna, which was used without modification in the tests. The encounter was reenacted a number of times in the Boston area, and in all cases TCAS surveillance was successful. Figure 10 shows one example. A track corresponding to the Cessna 172 extended over several minutes prior to the point of closest approach. Throughout these several minutes, surveillance data was available for display to the pilot to indicate the range, altitude, and azimuth angle of the Cessna 172. The data could also have triggered a maneuver advisory with sufficient warning time to prevent the collision. The San Diego collision geometry was also flown with a number of different aircraft types and produced essentially the same results.

The tracks plotted in Fig. 10 include four other aircraft in addition to the Cessna 172. Some surveillance imperfections that appear in the figure are worth examining because they indicate types of imperfections that occur in TCAS surveillance. At $t = 70$ s a track at 6 nmi gives rise to two additional tracks, presumably because of multipath. At $t = 160$ s a track at 5 nmi is dropped, although presumably the aircraft still exists. A fade in signal strength can cause such a track drop, possibly because one of the aircraft banked, or possibly because of a large look-down angle. Track reliability typically degrades as look-down angle increases. Fortunately, the co-altitude and near-altitude aircraft constitute the main threat of midair collisions, and in these cases surveillance is reliable. A reliability analysis of airborne measurements was undertaken to estimate the overall reliability of TCAS air-to-air surveillance. The overall result is 96%, which applies to a full population of encounters with a realistic mix of closing speeds and traffic densities [14].

Triggering of Alarms

The TCAS computer examines all of the surveillance tracks to determine whether any of them indicates an impending collision. If so, a maneuver advisory is displayed to the pilot along with an audible alarm. A traffic advisory is also issued before the maneuver advisory to aid in visual acquisition of the other aircraft and alert the pilot to respond to the maneuver advisory as soon as it appears.

An effective maneuver advisory must be generated with sufficient advanced warning time. Time is required for the pilot to react, for the aircraft to react and develop a climb rate (if the recommended maneuver is a climb), and for the climb to generate the needed displacement. The total time period required for a maneuver is approximately 25 s.

A *tau* alarm boundary is the basic technique used in TCAS to trigger alarms that, without producing an excessive alarm rate, provide the needed warning time for possible collisions. The range measurements provided by air-to-air surveillance are tracked to estimate range rate, and the resulting range and range rate are used to construct a linear extrapolation forward in time to determine the time of zero range. The resulting time, called tau, is

$$\text{tau} = \frac{\text{range}}{-\text{range rate}} .$$

This value would be the time remaining before collision if the two aircraft were on a collision course and were flying at constant velocities. By comparing tau with a threshold of 25 s, an alarm can be generated to provide the desired warning time.

To allow for aircraft accelerations and inaccuracies in the estimate of range rate, the alarm boundary is extended slightly by adding an offset to range. The resulting alarm boundary is

$$\text{alarm when } \frac{(\text{range} + D)}{(-\text{range rate})} < \text{threshold}.$$

The two parameters are assigned different values according to altitude; the most common assignments are a threshold of 25 s and $D = 0.3$ nmi. The altitude measurements are treated similarly to avoid alarms for aircraft that are safely separated in altitude. Conceivably, azimuth information could also be used to avoid alarms for encounters with horizontal separation. This capability is under development in TCAS III.

The alarm boundary that has been developed successfully provides advisories when needed while it keeps the total alarm rate acceptably low. The total alarm rate is approximately one maneuver advisory per 30 flight hours [15, 16].

An alarm boundary also triggers traffic advisories. The formulas are similar, with an increase of 15 s in tau threshold. As a result of the larger boundary, traffic advisories are generated at a higher rate, approximately one traffic advisory per two flight hours.

In addition to considerations of physics, the development of the TCAS alarm subsystem addressed a broad range of issues, including compatibility of TCAS alarms with air traffic control and other existing means of aircraft separation, integration of the TCAS displays and alarms into cockpits, and human factors.

Should TCAS Be an Executive System?

The Air Transport Association initially proposed that TCAS be an *executive system*. In other words, TCAS should issue maneuver advisories but not traffic advisories, and pilots should be instructed to follow the TCAS advisories without exception. Pilots, on the other hand, generally indicated a preference for the inclusion of traffic advisories and for pilot instruction that would not require rigid adherence to the TCAS maneuver advisories. This issue remained unresolved for many years.

As we gained experience testing TCAS in operational environments, we increasingly appreciated the value of the combination of the traffic advisories and a degree of pilot discretion. By the time of the RTCA's 1983 publication of TCAS standards [2], a traffic-advisory function was considered optional. Because the angle-of-arrival antenna for azimuth measurements had been developed by then, the traffic advisory display could indicate the direction of the other aircraft. Pilots who had TCAS experience indicated that azimuth information was the most useful component of the traffic-advisory display.

The traffic-advisory display is now a required function in the final TCAS standards. Furthermore, pilots can optionally extend the range of the traffic-advisory display beyond the minimum range needed for precursors to maneuver advisories. Pilots now typically enable the traffic display at all times, not only for close encounters.

The issue of pilot discretion in whether to follow a TCAS maneuver advisory has been extensively discussed. In TCAS operational testing, incidents occurred in which a pilot elected not to follow a TCAS maneuver advisory, on the basis of information available to the pilot that was unavailable to the TCAS computer. For example, encounters took place in which the approaching aircraft was below and climbing so that the linearly projected track indicated a possible collision. In this case the climbing aircraft was under instructions from ATC to level off at a lower altitude, and the pilot of the TCAS aircraft was aware of this intention through normal monitoring of the ATC radio. In such a case, the separation is provided by voice radio and a TCAS alarm is not necessary. These incidents are consistent with the basic principle that TCAS is intended as a backup system and is not intended to override the existing means of separation.

The aviation community ultimately decided that TCAS is not an executive system. According to the adopted standards, when a pilot has additional information, such as a visual sighting of the other aircraft, the TCAS maneuver advisory need not be executed. In the absence of additional information, however, the standards require a pilot to execute the maneuver that TCAS advises.

The Domino Effect

From the beginning the FAA considered the possibility that aircraft following TCAS advisories might significantly deviate from the ATC instructions and disrupt the ATC system. Conceivably, a TCAS aircraft maneuvering to avoid one aircraft might enter a conflict with a third aircraft. Then if that third aircraft is TCAS equipped it might maneuver and cause a conflict with a fourth aircraft. This scenario has been called a domino effect.

This issue was initially studied by means of an ATC simulation. The deviation in flight path that resulted from a TCAS maneuver was discovered to be small relative to the ATC-determined separations between aircraft. Because of the low alarm rate of TCAS and the normal spacings between controlled aircraft, no domino effect would occur [17]. This issue continued to receive attention as the TCAS design evolved in detail and as more information about TCAS performance under operational conditions became available. The more extensive results confirm that TCAS does not cause a domino effect and does not disrupt the ATC system [18].

Operational Testing

TCAS airborne tests were initially conducted under experimental conditions with aircraft operated by Lincoln Laboratory and the FAA and with pilots who were part of the development program. Subsequent steps were taken to obtain experience with TCAS under operational conditions. In 1981 a Boeing 727 operated by Piedmont Airlines was equipped with a special TCAS unit that recorded data but did not provide a pilot display. The TCAS equipment was built by Dalmo Victor Corp., and it recorded data while the aircraft proceeded in its normal passenger-carrying service. In this way TCAS performance data were obtained for actual operational conditions in every respect except for pilot response. A number of TCAS alarms were triggered during approximately 900 flight hours of data. The results measured the TCAS alarm rate and provided insights into the types of encounters that trigger alarms [19]. An air carrier operated by Air France also carried out a similar program in European airspace [20].

The next major step was to install TCAS with a pilot display as well as a data recorder in an operational air carrier. This installation was done in 1987 in another Boeing 727 operated by Piedmont Airlines. The TCAS equipment, built by Dalmo Victor, was a newer generation of equipment than that used in 1981. Figure 11 shows the TCAS equipment installed in the cockpit of the Piedmont aircraft. The program successfully accumulated over 600 flight hours of TCAS operation, and produced three kinds of data: (1) the TCAS data recordings giving surveillance tracks and alarms, (2) comments from trained observers regarding the observable conditions during the flights and the apparent usefulness of TCAS, and (3) systematically accumulated comments provided by the pilots. These results were again helpful in gaining experience with TCAS and understanding how it links with existing systems.

The author had the opportunity to be one of the cockpit observers in this program. Little TCAS activity occurred during that particular 10-hour period (no maneuver advisories and just one traffic advisory), but it was a satisfying experience to see the equipment installed and playing an integrated and appropriate role in the airspace system.

Additional operational testing of TCAS was carried out in 1988 and 1989. In a limited installation program, two manufacturers each designed and built several TCAS units and installed them on operational airliners. TCAS equipment built by Bendix/King was installed on a Boeing 737 aircraft and a DC-8 aircraft, both operated by United Airlines. TCAS equipment built by Sperry/Dalmo-Victor was installed on two MD-80 aircraft operated by Northwest Airlines. The program was organized by the FAA, and sponsored jointly by the FAA and the companies involved. Altogether the program added over 4,000 flight hours to the base of TCAS operational experience.

Conclusions

The TCAS design that resulted from this process achieves an effective balance among several considerations. Air-to-air surveillance is made possible through interoperability with ATC transponders that are in widespread use today. Even the first aircraft that are equipped with TCAS are able to carry out surveillance on all of the transponder-equipped aircraft. In spite of significant multipath disturbances, high densities of synchronous garble, and antenna shielding by aircraft fuselages, air-to-air sur-

Fig. 11—Operational testing in a Boeing 727 operated by Piedmont Airlines. The cockpit photograph shows the traffic advisory display.

veillance has been made reliable by a number of special techniques. The radio signals that TCAS transmits to carry out surveillance and maneuver coordination are accomplished at sufficiently low rates and powers so that TCAS does not interfere with ATC equipment operating in the same radio frequency bands. Alarm boundaries are set to provide sufficient warning time to prevent collisions, while they keep the total alarm rate low enough to be acceptable to pilots and ATC controllers. Pilots who have flown with TCAS are consistently enthusiastic about it.

Airborne tests and measurements played a principal role in the development program. While some aspects of TCAS could have been resolved by computer modeling and simulation, most could not have been understood without actual airborne measurements. In a number of experiences the airborne measurements yielded surprises that later seemed entirely reasonable once the phenomena were understood.

Multipath was the most difficult of the various challenges encountered in TCAS development. Even though a given surveillance track appears continuous and smooth to a pilot, in many cases an underlying density of multipath disturbances exists in both the interrogation link and the reply link. Several functions in the TCAS design sort out and clean up these disturbances.

"Why did it take so long?" is a question often asked about the TCAS development, which began in 1975. Part of the answer relates to the need for a number of program activities beyond the purely technical issues of air-to-air surveillance and timely triggering of alarms. For example, the support of pilots, controllers, airlines, and avionics manufacturers was necessary, and international standardization was beneficial. These processes cannot be carried out in a short time.

Another answer relates to the perceived threat of midair collisions. The adoption of a safety system like TCAS depends partly on technical developments and partly on the perceived need for the system. Following the 1978 midair collision in San Diego, an increased interest in the TCAS program focused on ways of minimizing the time to achieve operational status, because of a fear that the rate of midair collisions was increasing. If the rate of midair collisions had actually increased since 1978, then TCAS conceivably could have been called upon to solve the problem many years ago. Fortunately, the collision rate has not increased but has actually decreased. In the intervening years, further TCAS development has resulted in a number of design refinements and in a better understanding of TCAS behavior when integrated into the operational environment.

References

1. "System Safety Study of Minimum TCAS II," MITRE Corp., p. 3-5 (Dec. 1983), MTR-83W241.
2. "Minimum Operational Performance Standards for Traffic Alert and Collision Avoidance System (TCAS) Airborne Equipment," Radio Technical Commission for Aeronautics (Sept. 1983), RTCA/DO-185.
3. "Interim Standards and Recommended Practices (SARPs), Airborne Collision Avoidance Systems (ACAS)," *Report of the International Civil Aviation Organization (ICAO) Surveillance Improvements and Collision Avoidance Systems Panel (SICASP) Meeting, Montreal, 10–21 Apr. 1989.*
4. "Minimum Operational Performance Standards for an Active Traffic Alert and Collision Avoidance System I (Active TCAS I)," Radio Technical Commission for Aeronautics (20 Mar. 1987), RTCA/DO-197.
5. A.R. Paradis, "L-Band Air-to-Air Multipath Measurements," *Project Report ATC-77*, Lincoln Laboratory (6 Sept. 1977).
6. K.J. Keeping and J.C. Sureau, "Scale Model Pattern Measurements of Aircraft L-Band Beacon Antennas," *Project Report ATC-47*, Lincoln Laboratory (4 Apr. 1975).
7. D.W. Mayweather, "Model Aircraft L-Band Beacon Antenna Pattern Gain Maps," *Project Report ATC-44*, Lincoln Laboratory (16 May 1975).
8. W.H. Harman, "Air Traffic Density and Distribution Measurements," *Project Report ATC-80*, Lincoln Laboratory (3 May 1979).
9. W.H. Harman and R.S. Kennedy, "TCAS II: Design and Validation of the High-Traffic-Density Surveillance Subsystem," *Project Report ATC-126*, Lincoln Laboratory (12 Feb. 1985).
10. W.H. Harman, "Effects of RF Power Deviations on BCAS Link Reliability," *Project Report ATC-76*, Lincoln Laboratory (7 June 1977).
11. G. Patrick, J. Ludlam, and C. Gilchrist, "The Impact of a Traffic Alert and Collision Avoidance System on the Air Traffic Control Radar Beacon System and Mode S System in the Los Angeles Basin," Electromagnetic Compatibility Analysis Center (May 1985), DOT/FAA PM-84/30.
12. D.A. Spencer, R.R. LaFrey, J. DiBartolo, and W.H. Harman, "TCAS Experimental Unit (TEU) Hardware Description," *Project Report ATC-133*, Lincoln Laboratory (6 June 1986) p. 77.
13. M.L. Wood, "TCAS II ATCRBS Surveillance Algorithms," *Project Report ATC-131*, Lincoln Laboratory (28 Jan. 1985).
14. W.H. Harman, R.R. LaFrey, J.D. Welch, and M.L. Wood, "Active BCAS: Design and Validation of the Surveillance Subsystem," *Project Report ATC-103*, Lincoln Laboratory (17 Dec. 1980) p. 39.
15. D.H. Tillotson, L.D. Bolstridge, and G.P. Gamborani, "In-Service Evaluation of the Traffic Alert and Collision Avoidance System (TCAS) Industry Prototype," ARINC Research Corp., DOT/FAA/SA-88/2.
16. G.K. Schwind, W.B. Cotton, U.R.C. Gustafsson, and W.E. Nylen, "Summary User Evaluation Report on the Traffic Alert and Collision Avoidance System (TCAS II) Limited Installation Program," United Airlines (Oct. 1988).
17. J.E. Lebron and S.J. Mulder, "System Safety Study of Minimum TCAS II for Instrument Weather Conditions," *MITRE Project Report MTR-85W28* (1985).
18. N.A. Spencer, "An Update to the TCAS II System Safety Study," *MITRE Project Report MTR-88W115* (Dec. 1988).
19. T.P. Berry and R.D. Brock, "In-Service Evaluation of the Dalmo Victor Active Beacon Collision Avoidance System (BCAS/TCAS)," ARINC Research Corp. (Oct. 1982), DOT/FAA/RD-82/90.
20. V.P. Boykin, "Analysis of the TCAS II Logic Performance during the Air France Evaluation," *MITRE Technical Report MTR-86-W97* (Sept. 1986).

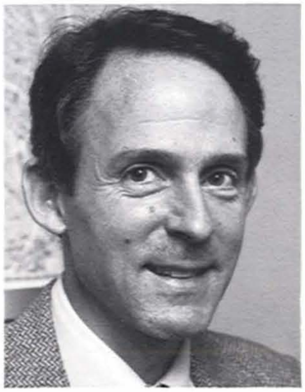

WILLIAM H. HARMAN is a staff member in the Air Traffic Surveillance Group. His research speciality is radar and communications. Bill received a B.S. degree from the University of Delaware, an M.S. degree from the University of Michigan, and a Ph.D. from MIT, all in electrical engineering. He has been at Lincoln Laboratory since 1971, and he has worked on the TCAS program since it began in 1975.