

TCAS-II: Design and Validation of the High-Traffic-Density Surveillance Subsystem

**W. H. Harman
R. S. Kennedy**

12 February 1985

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



Prepared for the Federal Aviation Administration,
Washington, D.C. 20591

This document is available to the public through
the National Technical Information Service,
Springfield, VA 22161

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.

1. Report No. DOT/FAA/PM-84/5	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle TCAS II: Design and Validation of the High-Traffic-Density Surveillance Subsystem		5. Report Date 12 February 1985	
		6. Performing Organization Code	
7. Author(s) William H. Harman and Robert S. Kennedy		8. Performing Organization Report No. ATC-126	
9. Performing Organization Name and Address Massachusetts Institute of Technology Lincoln Laboratory P.O. Box 73 Lexington, MA 02173-0073		10. Work Unit No. (TRAIS)	
		11. Contract or Grant No. DOT-FA77-WAL-817	
12. Sponsoring Agency Name and Address Program Engineering and Maintenance Service Department of Transportation Federal Aviation Administration Washington, D.C. 20591		13. Type of Report and Period Covered Project Report	
		14. Sponsoring Agency Code	
15. Supplementary Notes The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology, under Air Force Contract F19628-85-C-0002.			
16. Abstract Lincoln Laboratory, under FAA sponsorship, is developing an airborne collision avoidance system (TCAS II), concentrating primarily on the air-to-air surveillance subsystem. The surveillance functions required are to detect the presence of nearby transponder equipped aircraft, and then generate a surveillance track on each aircraft, issuing range and altitude reports once per second. The development effort from mid-1981 to the present has focused on the surveillance problems associated with high aircraft density. A number of surveillance techniques to deal with the high density environment have been identified and evaluated mainly through airborne measurements. A TCAS II design was synthesized, and this design was subjected to in-flight testing in the Los Angeles Basin using a Boeing 727. Results indicate that the performance objectives have been met.			
17. Key Words collision avoidance ATC beacon systems air-to-air surveillance TCAS BCAS Mode S DABS multipath interference fading		18. Distribution Statement Document is available to the public through the National Technical Information Service, Springfield, Virginia 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 162	22. Price

CONTENTS

EXECUTIVE SUMMARY	1
1. INTRODUCTION	1-1
1.1 Concept of TCAS	1-1
1.2 Air-to-Air Surveillance	1-1
1.3 Surveillance in High Aircraft Densities	1-3
1.4 Purpose and Overview of This Report	1-5
2. OVERVIEW OF DESIGN ISSUES	2-1
2.1 Mode C Synchronous Garble	2-1
2.1.1 Directional Interrogation and Whisper-Shout	2-1
2.1.2 Interference Limiting	2-1
2.1.3 Surveillance Processing Improvements	2-1
2.1.4 Increased Number of Reply Decoders	2-3
2.2 Fruit	2-3
2.2.1 Keying MTL to Whisper-Shout	2-4
2.2.2 Surveillance Processing Improvements	2-4
2.3 False Squitter Detections	2-4
2.3.1 Squitter Format Changed	2-5
2.4 Omnidirectional Squitter Reception Limited by Fruit	2-5
2.4.1 Multiple Beams and Multiple Receivers	2-5
2.4.2 Error Correction	2-6
2.5 Interference to Other Systems	2-6
2.5.1 Limiting Standard Revised	2-6
2.5.2 Adaptive Power Reduction	2-7
2.5.3 Beam Limiting	2-8
2.5.4 Mode S Algorithm Optimization	2-8
2.5.5 Reduction of Bottom Antenna Role	2-8
2.5.6 Keying Self Suppression Time to Antenna and/or Power	2-8
2.5.7 Improved Interrogation Decoding	2-8
2.5.8 Reduction of Scan Rate	2-9
2.6 False Tracks	2-9
2.6.1 Reduction of Bottom Antenna Role	2-9
2.6.2 Surveillance Processing Improvements	2-10
3. SURVEILLANCE IN MODE C	3-1
3.1 Whisper-Shout	3-1
3.1.1 More Capable Forms of Whisper-Shout	3-1
3.1.2 Baseline Whisper-Shout Design	3-4
3.2 Directional Interrogation	3-8
3.2.1 Beam Limiting	3-8
3.2.2 Airborne Experimentation	3-8
3.2.3 Phantom Mode A Interrogation	3-11
3.2.4 Beam Limiting Near Threshold	3-13
3.2.5 Late Mode C Replies	3-15
3.2.6 Example from Airborne Data	3-18
3.2.7 Summary	3-20

CONTENTS (CONT'D)

3.3	Role of the Bottom Antenna	3-20
3.4	Power Reduction	3-21
3.5	Airborne Measurements in the Los Angeles Basin	3-24
3.5.1	Truth	3-24
3.5.2	Flight Path	3-28
3.5.3	Aircraft Density	3-28
3.5.4	Advisory Rate	3-28
3.5.5	Antenna Problem	3-31
3.5.6	Case Studies	3-31
3.5.7	False Tracks	3-35
3.5.8	Statistical Analysis	3-38
3.5.9	Effect of Elevation Angle	3-38
3.5.10	Whisper-Shout Profiles	3-40
3.5.10.1	Fruit Rate	3-43
3.5.10.2	Synchronous Replies	3-43
3.5.10.3	Results	4-1
4.	SURVEILLANCE IN MODE S	4-1
4.1	Introduction	4-1
4.1.1	Functional Requirements	4-3
4.1.2	System Structure	4-6
4.1.3	Design Constraints	4-10
4.2	False Address Problem	4-16
4.3	Monitor State	4-16
4.3.1	Design Considerations	4-19
4.3.2	Parameters of the Running Sum Algorithm	4-24
4.3.3	Altitude Processing	4-26
4.4	Acquisition State	4-26
4.4.1	Functions	4-26
4.4.2	Threat Assessment	4-27
4.4.3	Interrogation Parameters	4-30
4.5	Track State	4-31
4.6	Dormant State	4-33
4.7	Performance Evaluation	4-33
4.7.1	Performance Goals	4-34
4.7.2	Models	4-35
4.7.3	Results	4-36
4.7.3.1	Low Altitude Encounters	4-43
4.7.3.2	High Altitude Encounters	5-1
5.	INTERFERENCE LIMITING	5-1
5.1	Interference Limiting Standard	5-1
5.1.1	Derivations	5-1
5.1.1.1	Reply Ratio	5-8
5.1.1.2	Suppression of Own Transponder	5-8
5.1.1.3	Fruit	5-12
5.1.2	Interference Simulation	

CONTENTS (CONT'D)

5.2	Interrogation Limiting Algorithm	5-14
5.2.1	Structure	5-14
5.2.2	Parameter Estimates	5-17
5.2.3	Performance Evaluation	5-19
5.2.3.1	Operating Environments	5-19
5.2.3.2	Simulation of the Environment	5-20
5.2.3.3	Low-Altitude Results	5-22
5.2.3.4	High Altitude Results	5-28
5.3	Transponder Suppression	5-30
5.3.1	Need for Re-examination of Mutual Suppression	5-30
REFERENCES		6-1
APPENDIX A Measured Reliability of Reply Decoding		A-1

ILLUSTRATIONS

1-1	Concept of TCAS	1-2
1-2	Synchronous garble	1-4
3-1	Simpliest form of whisper-shout	3-2
3-2	The form of whisper-shout in the BCAS design	3-3
3-3	Airborne measurements with an experimental form of whisper-shout	3-5
3-4	Comparison among five whisper-shout sequences	3-6
3-5	Baseline whisper-shout design	3-7
3-6	Airborne measurements using the baseline whisper-shout design	3-9
3-7	Effect of sidelobe suppression on directional transmit	3-10
3-8	Phantom Mode A interrogation	3-12
3-9	Single-pulse suppression	3-14
3-10	Beam limiting near threshold	3-16
3-11	Late Mode C replies	3-17
3-12	Directional interrogation - airborne data	3-19
3-13	Role of bottom antenna - airborne data	3-23
3-14	Performance using reduced power	3-25
3-15	Baseline TCAS II design for testing in LA	3-26
3-16	Flight path in the LA Basin	3-27
3-17	Aircraft density in Los Angeles	3-29
3-18	Measured antenna patterns	3-30
3-19	Nineteen close encounters in LA - directional performance	3-32
3-20	Nineteen close encounters in LA - omnidirectional performance	3-33
3-21	Performance measured in LA - directional design	3-36
3-22	Performance measured in LA - omnidirectional design	3-37
3-23	Elevation angles of TCAS targets	3-39
3-24	Whisper-shout profiles for omnidirectional interrogator	3-41
3-25	Whisper-shout profiles for directional interrogator	3-42
4-1	Levels of threat	4-2
4-2	Structure of the Mode S surveillance processor	4-5
4-3	Histogram of single-bit errors	4-11
4-4	Interrogations to repeating false addresses	4-13
4-5	Single-bit error rate observed in high density airspace	4-14
4-6	Mode S All-Call format	4-15
4-7	Threat assessment for monitor state processing	4-17
4-8	Monitor state processing algorithms	4-20
4-9	Operation of the Mode S surveillance processor	4-28
4-10	Track state interrogation power	4-32
4-11	Low altitude collision protection for normal operation with MTLs of 6, 9, and 12 dB	4-37
4-12	Low altitude collision protection with the bottom antenna disabled for MTLs of 6, 9, and 12 dB	4-38

ILLUSTRATIONS

4-13	Low altitude collision protection inferred from seven flight tests over land	4-39
4-14	Low altitude collision protection inferred from six flight tests over water	4-40
4-15	High altitude collision protection for normal operation with MTLs of 0, 3, and 6 dB	4-44
4-16	High altitude protection with the bottom antenna disabled for MTLs of 0, 3, and 6 dB	4-45
5-1	Interference limiting standard	5-2
5-2	Beam sharpening Factor, B	5-10
5-3	Interrogation limiting flow diagram	5-15
5-4	Altitude density function	5-21
5-5	Density of aircraft at three points in time	5-23
5-6	Variation of normalized interference limits and MTL in simulated high density environment	5-24
5-7	Variation of normalized interference limits during a rapid descent	5-27
5-8	Variation of normalized interference limits and MTL during a high altitude descent	5-29
5-9	Interrogation backscatter measurements (Mode C interrogation)	5-33
5-10	Interrogation backscatter measurements (Mode S interrogation)	5-34
A-1	Reply overlap scenarios tested	A-2

TABLES

2-1	Possible Design Changes	2-2
3-1	Surveillance False Alarms	3-22
3-2	False Track Rate, LA Basin	3-34
4-1	Design Constraints	4-7
4-2	Area in Which TCAS Differs from BCAS	4-9
4-3	Constraints on the Values of C, Z and z	4-23
4-4	Characteristics of Recorded Flight Encounters	4-42
5-1	Interim Results from Interference Simulation	5-13
5-2	Variation of MTL for a Low Altitude Environment	5-26
5-3	Variation of MTL for a High Altitude Environment	5-31
A-1	Reply Processor Test Conditions	A-3
A-2	Test Results	A-4

EXECUTIVE SUMMARY

This report summarizes a program conducted to design and evaluate TCAS II avionics, focusing on the air-to-air surveillance subsystem.

Concept of TCAS

The Traffic Alert and Collision Avoidance System (TCAS) is a system of airborne equipment being developed by the FAA for the purpose of preventing mid-air collisions. TCAS is intended as a collision avoidance backup to the existing system of air traffic control.

In one mode of operation, illustrated in Fig. 1-1, TCAS would prevent a collision between two aircraft, each equipped with a unit called TCAS II. Each TCAS II would sense the presence of the other aircraft, measure its location (in range, altitude, and bearing), detect a hazardous situation if one develops, and then display a resolution advisory (such as "climb" or "descend") to the pilot, after first carrying out an automatic coordination between the two aircraft to assure that the action taken by one aircraft will complement the action taken by the other aircraft.

As illustrated in Fig. 1-1, the TCAS II also affords protection against aircraft equipped with either Mode S or existing Secondary Surveillance Radar (SSR) transponders. For Mode S transponders, air-to-air surveillance is carried out in Mode S. For existing transponders, air-to-air surveillance is carried out in Mode C* (using Mode C-only interrogations, to which Mode S transponders do not reply). Mode S is used for surveillance of other TCAS II-equipped.

The TCAS II also affords protection against aircraft equipped with TCAS I which is a simpler form of TCAS. In these cases, there is no automatic coordination between the two aircraft; when necessary, the TCAS II generates a resolution advisory unilaterally, and in all respects behaves as if the other aircraft were equipped with just a transponder.

A TCAS II installation can conceptually be divided into two subsystems: (1) surveillance and (2) control logic. The former is the subject of this report.

Air-to-Air Surveillance

Air-to-air surveillance is accomplished by transmitting interrogations and receiving replies. The range between the two aircraft is determined from the time elapsed between interrogation transmission and reply reception. The altitude of the target aircraft is obtained from the altitude code, which is contained in the reply. Bearing relative to the nose of own aircraft is obtained by a direction finding antenna which is part of the TCAS II installation. Bearing measurements are coarse (standard deviation of about 10°), and are used in a traffic display but not in the control logic.

* The distinction between Mode C and Mode S is explained in Ref. 1.

The FAA is also developing, separately from the work documented here, an "Enhanced TCAS II" which uses a more accurate direction finding antenna (standard deviation of about 1°). The goal of that development is to achieve the capability for including horizontal resolution advisories in the control logic.

The altitude of the target aircraft is required by the TCAS II unit in order to generate vertical resolution advisories. Thus transponder-equipped aircraft that are not altitude reporting cannot participate in TCAS in this sense. For such aircraft, however, TCAS II can provide a measure of protection in the form of traffic advisories. Here the display indicates to the pilot the range and relative bearing of the target aircraft. Mode C surveillance of such non-altitude reporting aircraft is more challenging than in the altitude reporting case; the absence of an altitude measurement along with each range measurement makes it more difficult to form tracks from the set of received replies. This difficulty has been addressed in the TCAS development program, and a special form of surveillance processing, tailored to this mode, has been developed. This work is being documented separately.

Surveillance in High Aircraft Densities

The design of the air-to-air surveillance function of TCAS II builds on the previous development of BCAS (Beacon Collision Avoidance System, Ref. 2), by the addition of a number of improvements to accommodate higher aircraft densities. The BCAS design was intended for operation in low to moderate densities up to 0.02 aircraft/nmi². This value of density is not exceeded throughout most of the airspace in the United States. But it is exceeded locally in major metropolitan areas. Currently in parts of the Los Angeles Basin, the density averages about 0.1 aircraft/nmi². In 1981, the FAA adopted a change in the airborne collision avoidance concept, signified by the change in name from BCAS to TCAS. The design goal for aircraft density was changed to include the major metropolitan areas plus an allowance for future growth in air traffic. A density of 0.3 aircraft/nmi² was adopted as the specific goal.

In changing the BCAS design to accommodate this higher density, a number of issues had to be considered. Primary among these is the issue of synchronous garble in Mode C, illustrated in Fig. 1-2. Here, TCAS is performing surveillance using omnidirectional Mode C interrogations. When received, the replies from a particular aircraft-of-interest will be overlapped by replies from other aircraft at approximately the same range. This is called synchronous garble because the desired reply and the interfering replies are triggered by the same interrogation. If, for example, the aircraft-of-interest is at a range of 5 nmi and the aircraft density is 0.1 aircraft/sq. nmi, then the average number of other aircraft near enough in range to cause synchronous garble is 11 . It is impossible to reliably detect a reply in the presence of 11 overlapping replies.

Design Issues Addressed and Main Results

A conceptually straightforward technique for reducing synchronous garble is directional interrogation. A 4-beam antenna can be used, for example, and this is the design addressed in detail in this report. A directional interrogation eliciting a reply from the aircraft-of-interest (Fig. 1-2) will not elicit replies from other aircraft far away in azimuth, so synchronous garble is reduced. Additional interrogations transmitted in the other 3 beams make it possible to track these other aircraft as well.

Another technique that has been investigated for reducing synchronous garble is an increase in the number of whisper-shout interrogation levels. Whisper-shout is a multiple interrogation technique that was developed during the BCAS program (Ref. 3).

The methodology of the TCAS II design program can be described in terms of a number of improvements applied to BCAS, such as directional interrogation and extended whisper-shout, to make it capable of operating in high aircraft densities. Initially, the physical mechanisms (synchronous garble for example) that would cause performance degradation as density increases were identified. For each mechanism, several possible design changes were considered and evaluated by analysis, simulation, and airborne testing.

The TCAS II design that was developed has the following main characteristics:

Directional interrogation using a 4-beam antenna, with 90° beams, pointing forward, aft, left, and right, and including transmit sidelobe suppression. The antenna used in airborne testing is about 1/2 inch high and about 8 inches in diameter.

24-level whisper-shout, which is considerably more capable than the 4-level design in BCAS.

Role of bottom antenna. The bottom antenna plays a relatively minor role in this design. It is an omnidirectional monopole, whereas the top is directional. The bottom interrogations have lower peak power than the top by 18 dB, and a shorter whisper-shout sequence, 4 interrogations as compared with 24 for the top-forward beam. The role of the bottom antenna was reduced for two reasons. One is the reduction of false tracks (arising from multipath). The other reason involves the efficient use of the limited number of interrogations permitted in high density regions.

Changed squitter format. The Mode S squitter (which is the spontaneous transmission emitted by Mode S transponders, used in TCAS for detection of discrete addresses) was changed in message content. In its current form, the 24 parity bits appear in the clear, that is, not overlaid by the address as had been the case previously. This change was instituted primarily so that error

detection can be applied upon squitter reception. Error detection essentially eliminates the possibility of deriving false addresses from squitter receptions, which could otherwise become a major problem in high density airspace.

Improved Mode S surveillance processing. Mode S interrogations are transmitted individually to each target aircraft, and thus have to be carefully managed to prevent their becoming excessively numerous in high density airspace. This managing is done by the Mode S surveillance processor, which was redesigned extensively during the TCAS development program.

Revision of interference limiting standard. The interference limiting standard developed in the BCAS program placed limits on interrogation rate and power for the purpose of keeping all interference effects caused by BCAS to a negligibly low level. In transitioning to TCAS, the interference limiting standard had to be revised for several reasons. One concerns self suppression of own transponder (sometimes called "mutual suppression"). Because of directional interrogation and the expanded form of whisper-shout, a TCAS II unit will transmit interrogations at a considerably higher rate than that of BCAS. This could lead to a problem in the form of excessive self suppression. To manage this, a second inequality has been added to the interference limiting standard. In addition, the replies triggered by TCAS will constitute interference to other systems. Operation in high density airspace makes this effect potentially much more significant in TCAS than it was in BCAS. Accordingly a third inequality has been added to the interference limiting standard to limit the maximum amount of fruit generated by TCAS.

Performance

TCAS II performance was assessed in a number of ways including airborne measurements focusing on individual techniques and simulation of the Mode S surveillance processor. A primary step in the performance assessment process was a series of airborne measurements in the Los Angeles Basin aimed at evaluating the Mode C surveillance design as a whole. The LA Basin is known to have the highest density of aircraft in the United States. These tests were conducted in a Boeing 727 equipped with an experimental TCAS II unit having a 4-beam directional interrogator as well as the other TCAS II design characteristics listed above.

Performance was assessed by analyzing the data in several ways. One study focused on aircraft targets-of-opportunity that by chance passed by in a relatively close encounter. Surveillance reliability was good. In such cases the percentage of time during which the target aircraft was in track was about 97% (during the 50 second period prior to the point of closest approach in each encounter).

In a second study the detailed pattern of replies was analyzed to derive a quantitative estimate of the effectiveness of 4-beam directional interrogation in alleviating synchronous garble. These results show an improvement factor of 2.4, which is in agreement with the amount predicted according to the geometry of directional interrogation.

A third study was statistical, based on all of the aircraft that passed within 5 nmi in range while being within $\pm 10^\circ$ in elevation angle. The purpose of this study was to determine the functional dependence of surveillance reliability on aircraft density. The results indicate that there was not a significant degradation in performance as a function of density. The density values experienced in the LA Basin during these tests, although very high in an absolute sense, were not high enough to significantly degrade surveillance performance.

Conclusion

A TCAS II design which incorporates a top-mounted directional antenna and a bottom-mounted omnidirectional antenna and which employs a 24-level whisper-shout sequence and proven Mode S surveillance algorithms is capable of excellent surveillance reliability in today's high-density Los Angeles Basin environment and is predicted to continue to provide excellent performance in similar environments through the end of the century without detectable degradation to the performance of the ground-based beacon surveillance system.

1. INTRODUCTION

1.1 Concept of TCAS

The Traffic Alert and Collision Avoidance System (TCAS) is a system of airborne equipment, being developed by the FAA, for the purpose of preventing mid-air collisions. TCAS is intended as a collision avoidance backup to the existing system of air traffic control.

In one mode of operation, illustrated in Fig. 1-1, TCAS would prevent a collision between two aircraft, each equipped with a unit called TCAS II. Each TCAS II would sense the presence of the other aircraft, measure its location (in range, altitude, and bearing), detect a hazardous situation if one develops, and then display a resolution advisory (such as "climb" or "descend") to the pilot, after first carrying out an automatic coordination between the two aircraft to assure that the action taken by one aircraft will complement the action taken by the other aircraft.

As illustrated in Fig. 1-1, the TCAS II also affords protection against aircraft equipped with either Mode S or existing Secondary Surveillance Radar (SSR) transponders. For Mode S transponders, air-to-air surveillance is carried out in Mode S. For existing transponders, air-to-air surveillance is carried out in Mode C* (using Mode C-only interrogations, to which Mode S transponders do not reply). Mode S is used for surveillance of other TCAS II-equipped.

The TCAS II also affords protection against aircraft equipped with TCAS I which is a simpler form of TCAS. In these cases, there is no automatic coordination between the two aircraft; when necessary, the TCAS II generates a resolution advisory unilaterally, and in all respects behaves as if the other aircraft were equipped with just a transponder.

A TCAS II installation can conceptually be divided into two subsystems: (1) surveillance and (2) control logic. The former is the subject of this report.

1.2 Air-to-Air Surveillance

Air-to-air surveillance is accomplished by transmitting interrogations and receiving replies. The range between the two aircraft is determined from the time elapsed between interrogation transmission and reply reception. The altitude of the target aircraft is obtained from the altitude code, which is contained in the reply. Azimuth relative to the nose of own aircraft is obtained by a direction finding antenna which is part of the TCAS II installation. Azimuth measurements are coarse (standard deviation of about 10°), and are used in a traffic display but not in the control logic.

* The distinction between Mode C and Mode S is explained in Ref. 1.




Fig. 1-1. Concept of TCAS

The FAA is also developing, separately from the work documented here, an "Enhanced TCAS II" which uses a more accurate direction finding antenna (standard deviation of about 1°). The goal of that development is to achieve the capability for including horizontal resolution advisories in the control logic.


The altitude of the target aircraft is required by the TCAS II unit in order to generate vertical resolution advisories. Thus transponder-equipped aircraft that are not altitude reporting cannot participate in TCAS in this sense. For such aircraft, however, TCAS II can provide a measure of protection in the form of traffic advisories. Here the display indicates to the pilot the range and relative bearing of the target aircraft. Mode C surveillance of such non-altitude reporting aircraft is more challenging than in the altitude reporting case; the absence of an altitude measurement along with each range measurement makes it more difficult to form tracks from the set of received replies. This difficulty has been addressed in the TCAS development program, and a special form of surveillance processing, tailored to this mode, has been developed. This work is being documented separately.

1.3 Surveillance in High Aircraft Densities

The design of the air-to-air surveillance function of TCAS II builds on the previous development of BCAS (Beacon Collision Avoidance System, Ref. 2), by the addition of a number of improvements to accommodate higher aircraft densities. The BCAS design was intended for operation in low to moderate densities up to 0.02 aircraft/nmi². This value of density is not exceeded throughout most of the airspace in the United States. But it is exceeded locally in major metropolitan areas. Currently in parts of the Los Angeles Basin, the density averages about 0.1 aircraft/nmi². In 1981, the FAA adopted a change in the airborne collision avoidance concept, signified by the change in name from BCAS to TCAS. The design goal for aircraft density was changed to include the major metropolitan areas plus an allowance for future growth in air traffic. A density of 0.3 aircraft/nmi² was adopted as the specific goal.

In changing the BCAS design to accommodate this higher density, a number of issues had to be considered. Primary among these is the issue of synchronous garble in Mode C, illustrated in Fig. 1-2. Here, TCAS is performing surveillance using omnidirectional Mode C interrogations. When received, the replies from a particular aircraft-of-interest will be overlapped by replies from other aircraft at approximately the same range. This is called synchronous garble because the desired reply and the interfering replies are triggered by the same interrogation. If, for example, the aircraft of interest is at a range of 5 nmi and the aircraft density is 0.1 aircraft/sq. nmi, then the average number of other aircraft near enough in range to cause synchronous garble is 11. It is impossible to reliably detect a reply in the presence of 11 overlapping replies.

A conceptually straightforward technique for reducing synchronous garble is directional interrogation. A 4-beam antenna can be used, for example, and this is the design addressed in detail in this report. A directional



$N = 4 \pi R \Delta D$, where:

N = number of aircraft having range between $R - \Delta$ and $R + \Delta$.

R = range of target of interest

$\Delta \approx 1.7$ nmi (which is the reply length expressed as a distance) and

D = aircraft density.

Fig. 1-2. Synchronous garble.

interrogation eliciting a reply from the aircraft-of-interest (Fig. 1-2) will not elicit replies from other aircraft far away in azimuth, so synchronous garble is reduced. Additional interrogations transmitted in the other 3 beams make it possible to track these other aircraft as well.

Another technique that has been investigated for reducing synchronous garble is an increase in the number of whisper-shout interrogations. The whisper-shout technique is described in depth in Sec. 3.1.

1.4 Purpose and Overview of This Report

The purpose of this report is to document the results of the TCAS II surveillance development program. Chapter 2 outlines the issues that were addressed and the surveillance techniques that were considered. The other chapters describe the individual investigations and their results.

2. OVERVIEW OF DESIGN ISSUES

The TCAS II design program can be described in terms of a number of improvements applied to BCAS to make it capable of operating in high aircraft densities. The physical mechanisms (such as synchronous garble) that would cause performance degradation as density increases are listed in Table 2-1. For each mechanism, several design changes were considered. These are also listed in the table. The entries in Table 2-1 are described in the paragraphs that follow.

2.1. Mode C Synchronous Garble

Synchronous garble is a problem inherent in Mode C surveillance attributable to the all-call nature of the Mode C interrogation. Synchronous garble results in incorrect demodulation of altitude codes or complete inability to detect replies. These effects reduce the probability of tracking aircraft and produce false tracks.

2.1.1 Directional Interrogation and Whisper-Shout

The two main techniques identified for alleviating synchronous garble are directional interrogation and a more capable form of whisper-shout. These are both intended to partition the set of target aircraft into smaller sets of aircraft that reply to a single interrogation. Chapter 3 describes the development work on this subject that led to a particular design and describes the validation of this design through airborne measurements.

2.1.2 Interference Limiting

The introduction of directional interrogation in TCAS II required that changes be made in the interference limiting standard. Interference limiting provides bounds on permissible combinations of interrogation rates and powers for the purpose of assuring that any interference effects on other systems (such as SSR) are small enough to be negligible. In BCAS, interference limiting consisted of a condition, involving an interrogator's rate and power, that had to be satisfied by each BCAS interrogator. The condition was based on a criterion limiting the reduction in transponder reply ratio to 2 percent or less. Omnidirectional interrogation was a standard condition in BCAS, and this condition was used in deriving the interference limiting inequality. To provide for the possibility of directional interrogation in TCAS, it was necessary to re-examine the interference limiting issue. The work done in revising the interference limiting standard and in validating the results is presented in Chapter 5.

2.1.3 Surveillance Processing Improvements

Several additional techniques were considered for improving the ability to track aircraft in a synchronous garble environment without actually reducing the synchronous garble itself. Such techniques include the use of relative bearing angle, and whisper-shout index, in forming and extending tracks.

TABLE 2-1.

POSSIBLE DESIGN CHANGES

MECHANISMS THAT MAY LIMIT PERFORMANCE	POSSIBLE CHANGES IN TCAS II DESIGN	POSSIBLE CHANGES IN SYSTEM DESIGN
synchronous garble	<ul style="list-style-type: none"> • interrogate directionally • increase whisper-shout resolution • improve surveillance processing • increase number of reply decoders 	<ul style="list-style-type: none"> • revise interference limiting standard
fruit	<ul style="list-style-type: none"> • key MTL to whisper-shout • improve surveillance processing 	
false squitter detections	<ul style="list-style-type: none"> • test confidence • test relative bearing and/or amplitude • reduced use of bottom antenna 	<ul style="list-style-type: none"> • change squitter format
omni squitter reception limited by fruit	<ul style="list-style-type: none"> • use multiple beams and receivers • use error correction 	<ul style="list-style-type: none"> • change squitter format
interference to other systems	<ul style="list-style-type: none"> • adaptively reduce power • limit beam • optimize Mode S algorithms • reduce use of bottom antenna • key suppression time to antenna and/or power • improve interrogation decoder • reduce scan rate 	<ul style="list-style-type: none"> • revise interference limiting standard
false tracks	<ul style="list-style-type: none"> • reduce role of bottom antenna • improve surveillance processing 	

Another technique is the optimization of the values of tracking parameters, such as the number of coasts permitted before a track is dropped. These techniques were not explored simply because it was possible to achieve acceptable performance without them.

2.1.4 Increased Number of Reply Decoders

Another idea considered was an increase in the number of reply decoders. Four decoders were used in the BCAS equipment built by Lincoln Laboratory compared to three decoders in the BCAS equipment built by Dalmo Victor (Ref. 3, p. 82-3). Conceivably the large number of replies received in high density airspace could overload the bank of reply decoders and real replies would be lost simply because of insufficient space in which to save them. On the other hand, an increase in the number of reply decoders would not be expected to yield a major improvement in tolerance to aircraft density, since the additional replies to be saved in the added decoders would have been received in a severe overlap condition and would in most cases be corrupted by synchronous garble. Based on this reasoning, it was decided to not pursue this possible improvement in favor of the more promising improvements that directly reduce synchronous garble.

Appendix A gives the results of measurements of the reliability of correctly decoding a reply in the presence of interfering replies.

2.2. Fruit

Asynchronous replies received by a TCAS unit are called "fruit." These are replies triggered by other interrogators, and they appear in all reply modes. When a Mode C fruit reply is received during the listening period following a TCAS II interrogation in Mode C, then by itself it is indistinguishable from a desired synchronous reply. It is the function of the surveillance processing algorithms to distinguish between fruit and synchronous replies in establishing tracks.

In the BCAS program it was found that distinguishing fruit and synchronous replies is readily accomplished, with the result that fruit effects did not significantly degrade either the reliability of tracking real aircraft or the false track rate. These BCAS results apply in the low to medium density airspace for which BCAS was intended.

The transition from BCAS to TCAS changed the fruit conditions considerably. The higher aircraft densities into which TCAS can operate increase fruit rates proportionately. Furthermore, both the use of directional interrogation and the increase in the number of whisper-shout interrogations increase the number of reply listening periods, and thus increase the number of received fruit replies for a given fruit rate.

The overall increase in fruit can be estimated quantitatively as follows. An increase in aircraft density from 0.02 to 0.3 aircraft/nmi² is a fifteen-fold increase. The particular directional whisper-shout design that

was developed in this program uses a 4-beam antenna and a total of 83 interrogations per scan (where a scan is the surveillance update period, nominally 1 sec.). Since BCAS used 8 interrogations per scan, the number of reply listening periods is increased by a factor of about 10. Thus in each scan, the TCAS II unit may have to contend with 150 times as many fruit replies as BCAS.

2.2.1 Keying MTL to Whisper-Shout

One way of reducing the number of fruit replies received is by keying or matching the receiver MTL (minimum triggering level) in each listening period to the power level of that whisper-shout interrogation. Many of the whisper-shout interrogations are transmitted at very low power levels. In such a case, the aircraft that reply are for the most part those for which the antenna gains are high. For example, these aircraft may be at high elevation angles, where their bottom-mounted transponder antenna is transmitting in a favorable direction, and where the top-mounted TCAS II antenna is receiving in a favorable direction. It is also to be expected that for some targets the antenna pattern ripples will by chance line-up in such a way that the combined antenna gain is substantially greater than nominal. For these reasons the desired replies following a whisper-shout interrogation of low power are typically received at relatively high power levels. Thus a raised value of MTL is appropriate in eliminating fruit while still allowing the desired replies to be received. This technique was adopted for use in the experimental equipment tested and was found to operate successfully as shown in Sec. 3.5, which presents the results of airborne testing with this equipment.

2.2.2 Surveillance Processing Improvements

If the greatly increased fruit background were to cause the false track rate to become unacceptable, it would be appropriate to modify the surveillance processing algorithms to create a more favorable balance between false track rate and probability of tracking real aircraft. These possible improvements have not been explored because the false track rates experienced in airborne tests have remained at acceptable levels, as reported in Chapter 3.

2.3 False Squitter Detections

A squitter is a self-identifying message transmitted spontaneously by a Mode S transponder. When received by a TCAS II unit, a squitter indicates the presence of that aircraft and its discrete address, which can then be used in interrogating the aircraft in Mode S. In the BCAS development program, it was realized that there was some possibility of receiving false squitter information. That is, the process of receiving squitters and declaring the presence and address of an aircraft would occasionally be incorrect; an aircraft would be declared with the wrong address. How this could happen is described in some detail in Sec. 4.2. As a consequence of false squitter declarations, unnecessary interrogations would be transmitted based on these incorrect addresses, and these interrogations would use up part of the allowable total interrogation rate, thus reducing the number of real aircraft that could be tracked.

In BCAS several design features were adopted to minimize the rate of false squitter declarations. One was simply a duplicate test that was satisfied only when at least 2 squitter receptions indicated exactly the same address. Another BCAS feature was a confidence test in which the Mode S reply detector circuit declared a confidence bit along with each data bit in a detected reply. The reception was used in squitter declaration only if 21 or more of the 56 bits were flagged as high-confidence (Ref. 3, p. 29-32). An assessment of the final design of BCAS indicated that false squitter detections, while possible, were infrequent enough that no significant problem would result.

The transition from BCAS to TCAS opened this issue again. The very much higher aircraft densities into which TCAS is intended to operate will increase the rate of false squitter declarations substantially. One reason for an increase is the larger number of Mode S aircraft transmitting squitters, each of which can potentially become a false squitter detection. Another reason is the higher fruit environment.

2.3.1 Squitter Format Changed

The design change that was adopted was a change in the squitter message format to include error protection coding. Section 4.2 explains how this was accomplished. This change essentially eliminates the false squitter problem altogether. The other techniques that were considered (as listed in Table 2-1) thus became unnecessary and are not included in the TCAS II design.

2.4 Omnidirectional Squitter Reception Limited by Fruit

It is appropriate to use omnidirectional reception for squitters since their bearing angles of arrival are not known in advance. In the BCAS development program it was recognized that the fruit rates received by omnidirectional BCAS equipment are substantially greater than fruit rates that are typical for SSR ground stations. This difference is attributable to the omnidirectional reception in BCAS as compared to narrow-beam reception in SSR ground stations. Furthermore, the omnidirectional fruit rates in medium and high density airspace are high enough that they may significantly impact reception of Mode S replies and squitters. This impact can be described as a deterioration of receiver sensitivity, an effect described quantitatively in Sec. 4. Study of these effects during the BCAS program showed that no significant degradation in performance would result in the aircraft densities for which BCAS was designed.

The adequacy of omnidirectional squitter reception in high-density airspace was investigated as part of the transition from BCAS to TCAS.

2.4.1 Multiple Beams and Multiple Receivers

Directional reception would reduce the fruit rate during squitter listening periods. A single receiver could be used with a multi-beam directional antenna, in which case the receiver would have to be time-shared among the different beam positions as is typical in SSR. For squitter

reception, however, this may lead to a problem since many squitters would arrive at the TCAS II aircraft from directions other than the one currently being received. One solution to this problem would be to increase the standard squitter rate above the value 1/sec. adopted in BCAS. But such a change would have an undesirable impact on the interference aspects of TCAS design. A more costly approach would be to use multiple receivers, one for each antenna beam.

2.4.2 Error Correction

The change in squitter format discussed in Sec. 2.3 (which adds error protection coding to the squitter format) brings about an improvement in the performance of omnidirectional squitter reception, if an error correction function is added in the TCAS II design. The error correction capability is useful in several respects and has been adopted in the TCAS II design. As a result, the omnidirectional squitter reception (Sec. 4.7) performance is satisfactory, and it is not necessary to invoke directional reception.

2.5 Interference to Other Systems

Since TCAS interrogations and replies will be transmitted in frequency bands already in use, the possibility that TCAS might interfere with and degrade the performance of existing equipment was considered. It is necessary for the TCAS development program to limit its interference effects and to assure that such electromagnetic compatibility will in fact be achieved. In BCAS, this interference issue was addressed by the interference limiting function (described in Sec. 2.1), and by a comprehensive computer simulation performed by the Electromagnetic Compatibility Analysis Center (ECAC). But the fact that TCAS is intended for use in high density airspace made this interference issue much more challenging than it was in BCAS.

2.5.1 Limiting Standard Revised

Because of directional interrogation and an expanded form of whisper shout, a TCAS II unit will transmit interrogations at a considerably higher rate than that of BCAS. This could lead to a problem in the form of excessive self suppression of own transponder (sometimes called "mutual suppression"). To manage this, another inequality has been added to the interference limiting standard. This is described in Sec. 5.1.

Another effect is that the replies triggered by TCAS will constitute fruit interference to other systems. Operation in high density airspace makes this effect potentially much more significant in TCAS than it was in BCAS. Accordingly, as is described in Sec. 5.1, another inequality has been added to the interference limiting standard to limit the maximum amount of fruit generated by TCAS.

2.5.2 Adaptive Power Reduction

There is a fundamental difference between BCAS and TCAS regarding the conditions under which interference limiting is invoked. BCAS could operate in the low to medium density airspace for which it was designed without reaching the limiting point of the standard; thus the limiting standard served mainly as an overload control. In TCAS, however, the interference limit may be reached at a density considerably less than the maximum design density. Thus when TCAS operates in an area of maximum density, it will be functioning with reduced interrogation rate or power or both. The reduced power is still sufficient to achieve acceptable performance because of the natural correlation between density and closing speed. The reasoning for this statement is as follows.

Closing speeds in high density airspace are significantly less than values typical in low density airspace, as confirmed in airborne measurements (Ref. 3, p. 100-102). The goals for TCAS II design have been selected accordingly. In low density airspace, TCAS II will be capable of handling closing speeds up to 1200 knots. In the highest density airspace, TCAS II will be capable of handling closing speeds up to 500 knots. Lower closing speeds imply shorter range surveillance because sufficient time is available for the pilot and aircraft to react to a resolution advisory. A shorter range requirement implies, in turn, a lower interrogator power. Thus, if interference limiting in high density causes the interrogator power to be reduced, it is still possible to achieve satisfactory performance.

This qualitative reasoning provided the guidelines for the TCAS II development effort. Several things remained to be worked out quantitatively:

(1) An interference limiting algorithm, which is a part of a TCAS II unit. The algorithm performs power reduction as necessary to keep within the interference limiting standard, but does not reduce power more than necessary and sacrifice long range performance. The development of this algorithm is described in Sec. 5.2.

(2) Estimation of the amount of power reduction that will occur in high density. This has been estimated through simulation to be about 3 to 6 dB at low altitudes in the high densities for which TCAS II is being designed. The detailed result is described in Sec. 5.2.

(3) Assessment of surveillance reliability when operating at the reduced power. This has been addressed in several ways: airborne measurements in Mode C using targets of opportunity (described in Sec. 3), reprocessing of Mode S airborne data recorded, using a simulation of high density effects (described in Sec. 4.7).

2.5.3 Beam Limiting

Directional interrogations in Mode C can be beam limited by the sidelobe suppression action that results from the use of P2 pulses. This improves the ability to reduce synchronous garble, as discussed in Sec. 3.2. It also reduces the level of fruit interference generated by TCAS.

2.5.4 Mode S Algorithm Optimization

Mode S interrogations are controlled by algorithms that decide such things as: when to begin interrogating an aircraft whose squitters have been received, and when to stop interrogating an aircraft after it reaches long range or its replies become unreliable. The TCAS II design was more challenging in this respect than the BCAS design because of the needed capability for high density operation. Accordingly, a Mode S design study was undertaken, with the goals of assessing the need for improving the algorithms and then specifying improvements as necessary to make high density operation possible. This work is described in Chapter 4.

2.5.5 Reduction of Bottom Antenna Role

In BCAS the top and bottom antennas were used equally: the same number of interrogations were transmitted from each and with the same power levels. However, the bottom antenna was found to be significantly inferior to the top for purposes of air-to-air surveillance. This observation suggested that a more efficient design would be achieved by reducing the role of the bottom antenna relative to the top, and that such an improvement would be particularly significant in the context of TCAS II where interference limiting places a constraint on interrogation rate and power. Work on this issue is described in Sec. 3.3.

2.5.6 Keying Self Suppression Time To Antenna and/or Power

In BCAS the self suppression time (suppression of own transponder functions at the time of each interrogation transmission) was 200 μ sec. This constant value was used regardless of which antenna was being used for the interrogation and regardless of the interrogation power. The interrogation itself has a duration of about 20 μ sec, but the suppression was made longer because of multipath effects (Ref. 3, p. 20-23). Since the multipath effects may be expected to be more severe for bottom antenna transmissions and more severe for high power transmissions, the design could be made more efficient by keying the suppression time to antenna and/or power. This issue has been addressed in the TCAS program through airborne measurements of the duration of multipath backscatter. This work is described in Sec. 5.3.

2.5.7 Improved Interrogation Decoding

The long self suppression used after a TCAS interrogation is intended to prevent own transponder from decoding an interrogation when multipath backscatter is received immediately following the transmission of an interrogation. Part of the problem is due to the fact that own transponder's

interrogation decoder may accept an erratic multipath reception that has little resemblance to a valid interrogation. Stricter standards could be written for the interrogation decoding function of own transponder, so that real interrogations would be decoded with essentially the same reliability and yet the frequency with which multipath would qualify as an interrogation would be greatly reduced. This would make it possible to reduce the duration of self suppression, which in turn would increase the allowable interrogation rate permitted within the interference limiting standard. It was found that varying the self-suppression time was sufficient to avoid over-suppression of the on-board transponder. Thus, transponder design changes are not necessary.

2.5.8 Reduction of Scan Rate

BCAS was designed with a scan rate of 1/second, which means that each track of an aircraft would be updated with a new position measurement nominally once each second. An obvious change that might be considered in transitioning to TCAS is to reduce the scan rate, which would make it possible to conduct surveillance on a larger number of aircraft within the same interference limits. It was determined that after other improvements had been made, the capability of the resulting TCAS II design was sufficient to meet the interference limiting goals with a one-second scan rate.

2.6 False Tracks

A false track is a surveillance track that is delivered to the control logic subsystem but that does not correspond to a real aircraft. In TCAS II as in BCAS, there are no false tracks in Mode S, but in Mode C false tracks do occur. The mechanism that prevents Mode S false tracks is the selectively addressed interrogation function; unless a received interrogation agrees exactly in all 24 bits with a transponder's unique address, the transponder will not reply.

False tracks in Mode C are of concern because of the possibility that a resolution advisory (RA) may be triggered by a false track, or that an RA that was triggered by a real aircraft may be modified by a false track. Such "false RAs" were very rare in BCAS. At the time of the BCAS Conference in January 1981, not a single false RA had occurred in all of the airborne experience which consisted of several hundred flight hours. But in the context of TCAS the false track rate is expected to be higher for several reasons: one is the higher density of aircraft and higher fruit rate, and another is the increase in the number of fruit replies that results simply from the increased number of interrogations. Thus design changes aimed at false track reduction were needed.

2.6.1 Reduction of Bottom Antenna Role

Since many of the false tracks observed are due to multipath, and since multipath effects are consistently more severe when using the bottom antenna, a reduction in the role of the bottom antenna is a straightforward way of reducing false tracks. This technique has been addressed by means of airborne measurements, as described in Sec. 3.3.

2.6.2 Surveillance Processing Improvements

The false track rate can be affected by changes in the surveillance algorithms. For example the fundamental tradeoff between false track rate and probability of miss is affected by tracking parameters, such as the number of scans in which a reply must be received before a track is established. As described in Sec. 3.5.7, changes were made in the handling of multipath tracks and provisions were added to filter out-of-beam replies. These reduced the false track rate sufficiently in high density so that other tracking parameter changes, which would have reduced the probability of track, were not required.

3. SURVEILLANCE IN MODE C

This chapter describes the investigations of whisper-shout, directional interrogation and the other surveillance improvement techniques outlined in the preceding chapter. Results of experiments are given, followed by a definition of the TCAS II design that resulted. The chapter concludes with quantitative performance results obtained from airborne measurements in the Los Angeles Basin.

3.1 Whisper Shout

The purpose of whisper-shout is to partition or subdivide the set of synchronously garbling aircraft so that fewer will reply to any one interrogation.

The simplest form of whisper-shout is illustrated in Fig. 3-1. In this 2-level whisper-shout, the purpose is to divide the synchronous garble population into two approximately equal subsets. The first interrogation* is transmitted at a relatively low power level so that approximately half of the aircraft in the synchronous garble range band will receive it above threshold. Thus only these will reply to the first interrogation, and the synchronous garble problem will be reduced by a factor of about 2 in this first reply listening period. The second interrogation is transmitted at full power so as to be detectable by all of the aircraft. But this interrogation is preceded by an additional pulse, denoted S1, of power level nearly equal to that of the first interrogation. The purpose of S1 is to trigger the suppression function in those transponders that replied to the first interrogation. Thus this first set of aircraft will not reply again, and so in the second listening period, the synchronous garble problem will again be reduced by a factor of about 2. To make sure that each aircraft replies to either the first or the second interrogation, the power of S1 is made slightly less than that of the first interrogation, thus overlapping the two reply bands.

In the BCAS design, a 4-level form of whisper-shout was used, illustrated in Fig. 3-2. It may also be noted in this figure that there are two suppression pulses instead of the one (S1) shown in the preceding figure. This alternative way of accomplishing the whisper-shout suppression was used in BCAS because it allowed more time to change the transmitter power level. It will be shown in Sec. 3.2 that when directional interrogation is used, the single pulse suppression is preferable.

3.1.1 More Capable Forms of Whisper-Shout

To handle the very much higher aircraft densities associated with TCAS II, higher resolution whisper-shout sequences were investigated. It is to be expected that increasing the number of interrogations in the whisper-shout sequence will further reduce the number of aircraft that reply to a single interrogation.

* P1 and P3 constitute a Mode C interrogation. The purpose of P4 is to cause Mode S transponders to not reply; these aircraft are tracked separately in Mode S.




Fig. 3-1. Simplest form of whisper-shout.




Fig. 3-2. The form of whisper-shout in the BCAS design.

To verify this expectation, airborne measurements were undertaken comparing the BCAS form of whisper-shout, in which the interrogation spacing is 6 dB, to a higher resolution form of whisper-shout in which the interrogation spacing is 2 dB. Whereas the "overlap" was 3 dB in BCAS (that is, the suppression was 3 dB lower in power than the preceding interrogation), a 1 dB overlap was used in the higher-resolution whisper-shout sequence. The experiment was conducted by alternating between the two sequences so that data of both kinds were recorded in each 1-sec. scan. Results from these airborne measurements are shown in Fig. 3-3. The plot shows the average number of replies per interrogation for each of the interrogations in the sequence. The results indicate that the higher-resolution sequence was successful in reducing the reply counts and thus would significantly alleviate synchronous garble effects.

In a further experiment of this kind, five forms of whisper-shout were compared. A description of this experiment is best stated in terms of the whisper-shout "bin width," which is the difference in dB between an interrogation and the associated suppression. In the original BCAS design, for example, the bin width was 9 dB, and in the higher-resolution sequence represented in Fig. 3-3, the bin width was 3 dB. This experiment was intended to determine if the number of replies to a whisper-shout interrogation would be roughly proportional to bin width.

Airborne measurements were conducted alternating each second between five sets of whisper-shout interrogations. The BCAS design was included as one of the sets, and the others all were of smaller bin widths, namely 4 dB, 3 dB, 2 dB, and 1 dB. The results are shown in Fig. 3-4, where the average number of replies per interrogation are plotted as a function of bin width. These results confirm that a reduction in bin width causes a significant reduction in the number of replies per interrogation. This relationship holds consistently in all of the points plotted in Fig. 3-4.

3.1.2 Baseline Whisper-Shout Design

Based on these favorable results, a specific whisper-shout design for use in TCAS II was selected. This baseline design is defined in Fig. 3-5, where it is compared with the 4-level BCAS design. The new design has 24 levels, and alternates between bin widths of 2 dB and 3 dB. In selecting this baseline design, it was necessary to consider interference limiting (which is the subject of Chapter 5). When a TCAS II aircraft using this whisper-shout sequence flies into an area of aircraft density so high that some modification in transmitted interrogation rate or power is required, the procedure will be simply to truncate the sequence beginning at the top. This will reduce the number of interrogations per second, the peak interrogation power, and the rate-power product, while still maintaining an effective surveillance capability for most of the aircraft in the vicinity.

To verify this expectation, airborne measurements were undertaken comparing the BCAS form of whisper-shout, in which the interrogation spacing is 6 dB, to a higher resolution form of whisper-shout in which the interrogation spacing is 2 dB. Whereas the "overlap" was 3 dB in BCAS (that is, the suppression was 3 dB lower in power than the preceding interrogation), a 1 dB overlap was used in the higher-resolution whisper-shout sequence. The experiment was conducted by alternating between the two sequences so that data of both kinds were recorded in each 1-sec. scan. Results from these airborne measurements are shown in Fig. 3-3. The plot shows the average number of replies per interrogation for each of the interrogations in the sequence. The results indicate that the higher-resolution sequence was successful in reducing the reply counts and thus would significantly alleviate synchronous garble effects.

In a further experiment of this kind, five forms of whisper-shout were compared. A description of this experiment is best stated in terms of the whisper-shout "bin width," which is the difference in dB between an interrogation and the associated suppression. In the original BCAS design, for example, the bin width was 9 dB, and in the higher-resolution sequence represented in Fig. 3-3, the bin width was 3 dB. This experiment was intended to determine if the number of replies to a whisper-shout interrogation would be roughly proportional to bin width.

Airborne measurements were conducted alternating each second between five sets of whisper-shout interrogations. The BCAS design was included as one of the sets, and the others all were of smaller bin widths, namely 4 dB, 3 dB, 2 dB, and 1 dB. The results are shown in Fig. 3-4, where the average number of replies per interrogation are plotted as a function of bin width. These results confirm that a reduction in bin width causes a significant reduction in the number of replies per interrogation. This relationship holds consistently in all of the points plotted in Fig. 3-4.

3.1.2 Baseline Whisper-Shout Design

Based on these favorable results, a specific whisper-shout design for use in TCAS II was selected. This baseline design is defined in Fig. 3-5, where it is compared with the 4-level BCAS design. The new design has 24 levels, and alternates between bin widths of 2 dB and 3 dB. In selecting this baseline design, it was necessary to consider interference limiting (which is the subject of Chapter 5). When a TCAS II aircraft using this whisper-shout sequence flies into an area of aircraft density so high that some modification in transmitted interrogation rate or power is required, the procedure will be simply to truncate the sequence beginning at the top. This will reduce the number of interrogations per second, the peak interrogation power, and the rate-power product, while still maintaining an effective surveillance capability for most of the aircraft in the vicinity.




Fig. 3-3. Airborne measurements with an experimental form of whisper-shout.




Fig. 3-4. Comparison among five whisper-shout sequences.

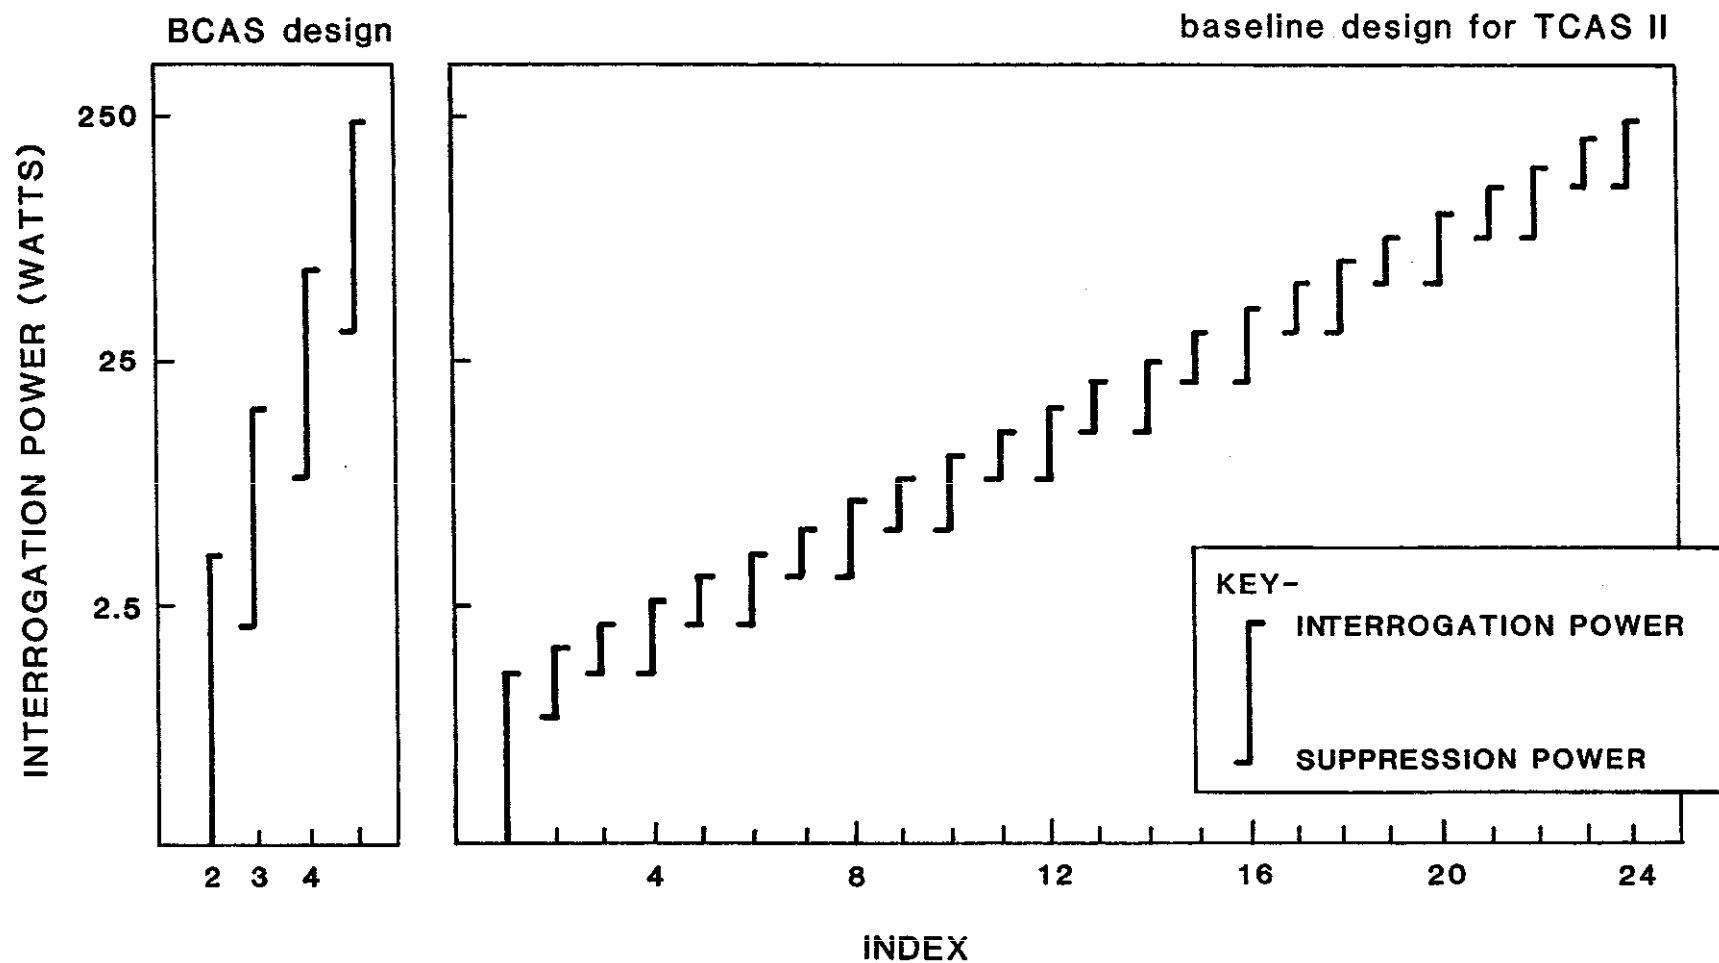


Fig. 3-5. Baseline whisper-shout design.

Airborne measurements have been carried out using this baseline design. The data in Fig. 3-6 are typical of the results of these measurements. As before, the measurements were set up in the form of a comparison with the BCAS design. This figure shows range tracks as a function of time. It is seen that there are numerous cases in which the 24-level whisper-shout achieved significantly better performance than BCAS.

Section 3.5 below describes further airborne measurements using this 24-level whisper-shout sequence, in this case in the LA Basin. Flights were conducted in LA in order to experience very high traffic density conditions. Performance was found to be good, and the results support the conclusion that the baseline whisper-shout design of Fig. 3-5 is suitable for TCAS II.

3.2. Directional Interrogation

Directional interrogation is a conceptually straightforward technique for combatting synchronous garble. A directional interrogator elicits replies from aircraft in one sector at a time, thus significantly reducing the number of replies per interrogation.

3.2.1 Beam Limiting

In developing a specific design, an initial issue to consider is whether or not to use sidelobe suppression (SLS) for beam limiting. SLS can be implemented by incorporating P2 pulses in the interrogations, as is normal for ground based interrogators (Fig. 3-7). When a received interrogation is accompanied by a P2 pulse of power greater than the interrogation, the transponder does not reply. If the TCAS II interrogator transmits P2 pulses on a notched pattern, the relative powers in space of P1 and P2 will serve to limit the region of replying aircraft to just the mainbeam.


If sidelobe suppression is not used with a directional antenna, the antenna will interrogate to some extent in all directions (Fig. 3-7). Considering the modest front-to-back ratios that will be achievable with airborne antennas of reasonable size, it is concluded that directional transmissions without SLS will not achieve the sector-by-sector separation normally associated with directional interrogation.

Based on these considerations, transmit sidelobe suppression has been adopted in the TCAS design.


3.2.2 Airborne Experimentation

A 4-beam directional interrogator was built by Dalmo Victor and installed in an FAA Boeing 727. This aircraft was also equipped with an omnidirectional TCAS Experimental Unit (TEU, built by Lincoln Laboratory) so that comparisons could be made to help show the degree of improvement derived through directional interrogation. Interrogations from the two units were interleaved in each 1 sec. scan.


Fig. 3-6. Airborne measurements using the baseline whisper-shout design.



OMNIDIRECTIONAL



DIRECTIONAL
WITHOUT SLS



DIRECTIONAL
WITH SLS




Fig. 3-7. Effect of Sidelobe Suppression on Directional Transmitt.

The 3-dB beamwidth of the Dalmo Victor directional antenna is about 90° in each of the four directions, which are aimed forward, right, aft, and left. The antenna also provides a notched SLS control pattern corresponding to each of the four directional beams. An omnidirectional transmitting pattern is also provided.

Both units were capable of transmitting the baseline whisper-shout sequence (shown in Fig. 3-5) so that directional interrogation and high-resolution whisper-shout, could be tested together to reveal any interactions between them. In fact there were some significant interactions as described below.

Both units were configured to record data at the reply level. That is, surveillance tracks and control logic products were not recorded. Surveillance processing was carried out after the flights. This was done so that the limited tape recorder capacity could be used most effectively, and so surveillance processing could be kept flexible by recording data that did not depend on surveillance processing.

The flight plans included provisions for a mission to the LA Basin in order to experience the very high aircraft density known to exist there. Initial airborne experimentation was performed in the Boston to Washington, D.C. area, to validate the experimental equipment, and to gain experience with the equipment and data formats. This local experimentation also yielded qualitative performance results, which were supplemented later by the flights in LA.

3.2.3 Phantom Mode A Interrogation

As airborne data began to be collected, one of the first things noted was a problem of unwanted replies appearing at shorter range relative to the expected replies from certain aircraft. After examining such data in detail, it was concluded that the mechanism causing these unwanted replies is as follows (see Fig. 3-8). The interrogation transmitted by the directional unit consisted of 6 pulses, as shown. The interrogation is identical to the BCAS interrogation (Fig. 3-2) with the addition of a P2 pulse for sidelobe suppression. Note that for these experiments the whisper-shout suppression was transmitted as a pair of pulses.

Consider a scenario in which there is a particular target aircraft and an interrogation being transmitted in some other direction. The desired reaction is for the aircraft to not reply, because of sidelobe suppression. But if the interrogation is received at the transponder near threshold as illustrated, it becomes possible for a Mode A reply to be triggered by the combination of S1 and P2. Such replies would not occur if the transponder threshold transition were abrupt, such that a pulse is detected with probability zero when below threshold and probability one when above. If the threshold were abrupt and




Fig. 3-8. Phantom Mode A interrogation.

- a) S1 were received below threshold, then it would not be detected, and could not contribute to an interrogation detection,
- b) S1 were received above threshold, then S2 would also be above threshold, and the pair would put the transponder into suppression.

Either way, there would be no reply. In reality, however, the threshold behavior is not abrupt. There is a band caused by receiver noise, typically 2 to 3 dB wide, over which pulse detection varies from zero to one. Thus when S1 is received in this band, as illustrated in Fig. 3-8, it is possible for S1 to be detected and S2 to be not detected. When this occurs, the transponder will not go into suppression, and a subsequent pulse reception may combine with S1 to form an accepted interrogation. If, as in this scenario, P2 is received exactly 8 μ sec after S1, the transponder will reply in Mode A.

The conditions that allow these undesired replies are present only when directional interrogations are combined with high resolution whisper-shout. The occurrence of the 8 μ sec pulse spacing is a result of the particular timing in this implementation of S1 relative to P1. Furthermore, because of the high resolution whisper-shout sequence being used, it is likely that several of the interrogations will be received with S1 in the threshold region.

This problem can be cured in a straightforward manner by changing the timing of the whisper-shout suppression. In considering other values of the time between S1 and P1, it was necessary to check all of the defined interrogation modes to be sure that another similar problem did not appear in place of this one. Among the candidates considered were the single-pulse suppression, illustrated in Fig. 3-9. Here the first two pulses, S1 and P1, act together to suppress transponders whenever S1 is detectable.

The single-pulse suppression was first tested at Lincoln Laboratory using a rooftop antenna driven by a TEU. This test employed both the single-pulse suppression and the two-pulse suppression, interleaved in each 1-second scan. The two techniques were compared against the same aircraft targets and, there was no difference in surveillance performance. The directional interrogator unit was then modified by Dalmo Victor to use the single-pulse suppression. In all of the airborne testing that has followed, no unforeseen problems have appeared, and the unwanted early Mode A replies have been eliminated.

3.2.4 Beam Limiting Near Threshold

Another observation that was made when airborne data first became available involves the mechanism of beam limiting near transponder threshold. SLS limits the beamwidth over which transponders reply to any one interrogation. In choosing the P2 power level, it is necessary to ensure that, for every transponder, the beamwidth is sufficiently large to prevent gaps between beams. Because the National Standard permits a 9 dB tolerance in the




Fig. 3-9. Single-pulse suppression.

P1-to-P2 power test*, it was originally concluded that the transmitted P2 power would have to be quite low, and that as a result the reply beamwidth for typical transponders would be much larger than 90°. The end result might be a relatively small amount of improvement attributable to directional interrogation.

As airborne data became available it was realized that there is an important relationship between SLS and whisper-shout that affects the uniformity of beam limiting. This is illustrated in Fig. 3-10. The transponder will reply only when S1 is just below threshold and P1 is just above threshold. In this scenario, P2 is received slightly above P1 - 9 dB. According to the National Standard, reply is optional. But in actuality, since P2 is well below threshold the transponder will reply.

Based on this realization, the power level of P2 transmissions was increased from a level 4 dB below P1 to the same level as P1. Furthermore it was concluded that reply beamwidths will be more uniform from transponder to transponder, and that the performance improvement attributable to directional interrogation will be somewhat better than was originally expected.

The degarbling performance of directional interrogation can be estimated quantitatively as follows. Based on antenna patterns measured in an anechoic chamber prior to installation, and for an interrogator transmitting with P1/P2 = 0 dB,

reply beamwidth = 125°	if THR = 0 dB
122°	if THR = 1 dB
118°	if THR = 2 dB
115°	if THR = 3 dB
111°	if THR = 4 dB

where THR is the transponder P1/P2 reply threshold. Because of the whisper-shout action, THR is at most a few dB for the interrogations eliciting replies. An average value of THR is about 1 dB, and the corresponding value of beamwidth can be taken as an estimate of the effective average;


effective average beamwidth = 122°

$$\text{improvement factor} = \frac{360^\circ}{122^\circ} = 2.9$$

3.2.5 Late Mode C Replies

The first airborne data also revealed another problem. The set of received replies was seen to contain unwanted replies appearing at longer range (by about 1/6 mile) relative to the desired replies from certain aircraft. It was determined that these unwanted replies were caused by the mechanism illustrated in Fig. 3-11. The combination of P2 and P4 acts like a Mode C interrogation, producing a Mode C reply that is late by 2 μ sec.


* Reply is required when $P2 < P1 - 9 \text{ dB}$. Reply is prohibited when $P2 > P1$.



SPECIFICATION: REPLY OPTIONAL.

IN ACTUALITY: REPLY LIKELY.

Fig. 3-10. Beam limiting near threshold.



SCENARIO: TARGET OUT OF MAINBEAM

DESIRED REACTION: NO REPLY

ACTUAL REACTION: LATE MODE C REPLY (INTERMITTENT)

Fig. 3-11. Late Mode C replies.

Such replies had not been anticipated, based on an abrupt threshold model: If P1 is received above threshold, and P2 exceeds P1 as in the illustration, then these two pulses would both be detected and they would put the transponder into suppression (which is the normal SLS mechanism). If P1 is received below threshold, then P4 would also be below threshold and would not contribute to the late Mode C interrogation.

To understand how these unwanted replies can be triggered, it is necessary to, once more, view the transponder threshold as a band rather than an abrupt transition. When P1 is received in the threshold band, it is possible for P1 to be missed and yet P4 to be detected. Whenever this happens, an unwanted late Mode C reply will result.

Such a mechanism will of course be intermittent, and this is consistent with the observed airborne data: The number of late Mode C replies is approximately 15% of the number of desired replies received.

Two cures were considered. First, one might transmit P1 at a higher power level than P3 and P4, perhaps by 1 dB. Alternatively, the unwanted replies could be removed in surveillance processing, using the 2 μ sec spacing as a means of identifying them. It was found that the experimental interrogator being used could not readily be modified to change P1 power relative to the other pulses, and for this reason it was decided to remove the replies in surveillance processing. However, the unwanted replies will still be present in the set of received replies and will constitute additional synchronous garble.

Since it is the purpose of directional interrogation to reduce synchronous garble, these late Mode C replies will slightly reduce the effectiveness of the technique. The improvement factor, estimated to be 2.9 in the preceding section, may be expected to be reduced to approximately

$$\text{net improvement factor} = 2.9/1.15 = 2.5$$

3.2.6 Example from Airborne Data

The initial airborne data was also examined for reasonableness in regard to directional interrogation. It was expected that examples could readily be found in which an encounter with a target aircraft produced replies first to one beam, then another, and then possibly a third. In fact, such examples were immediately apparant, one of which is shown in Fig. 3-12. This target aircraft first appears ahead and slightly to the right (judging from the azimuth values recorded). It passed to the left, coming as close as about 1.5 nmi. Replies are shown in the figure as range vs. time, with replies to interrogation in different beams plotted separately. This reply data shows that initially only the front beam elicited replies from this aircraft. Later the left beam did, and finally the back beam did, consistent with the flight path. There are no gaps at the beam transitions.


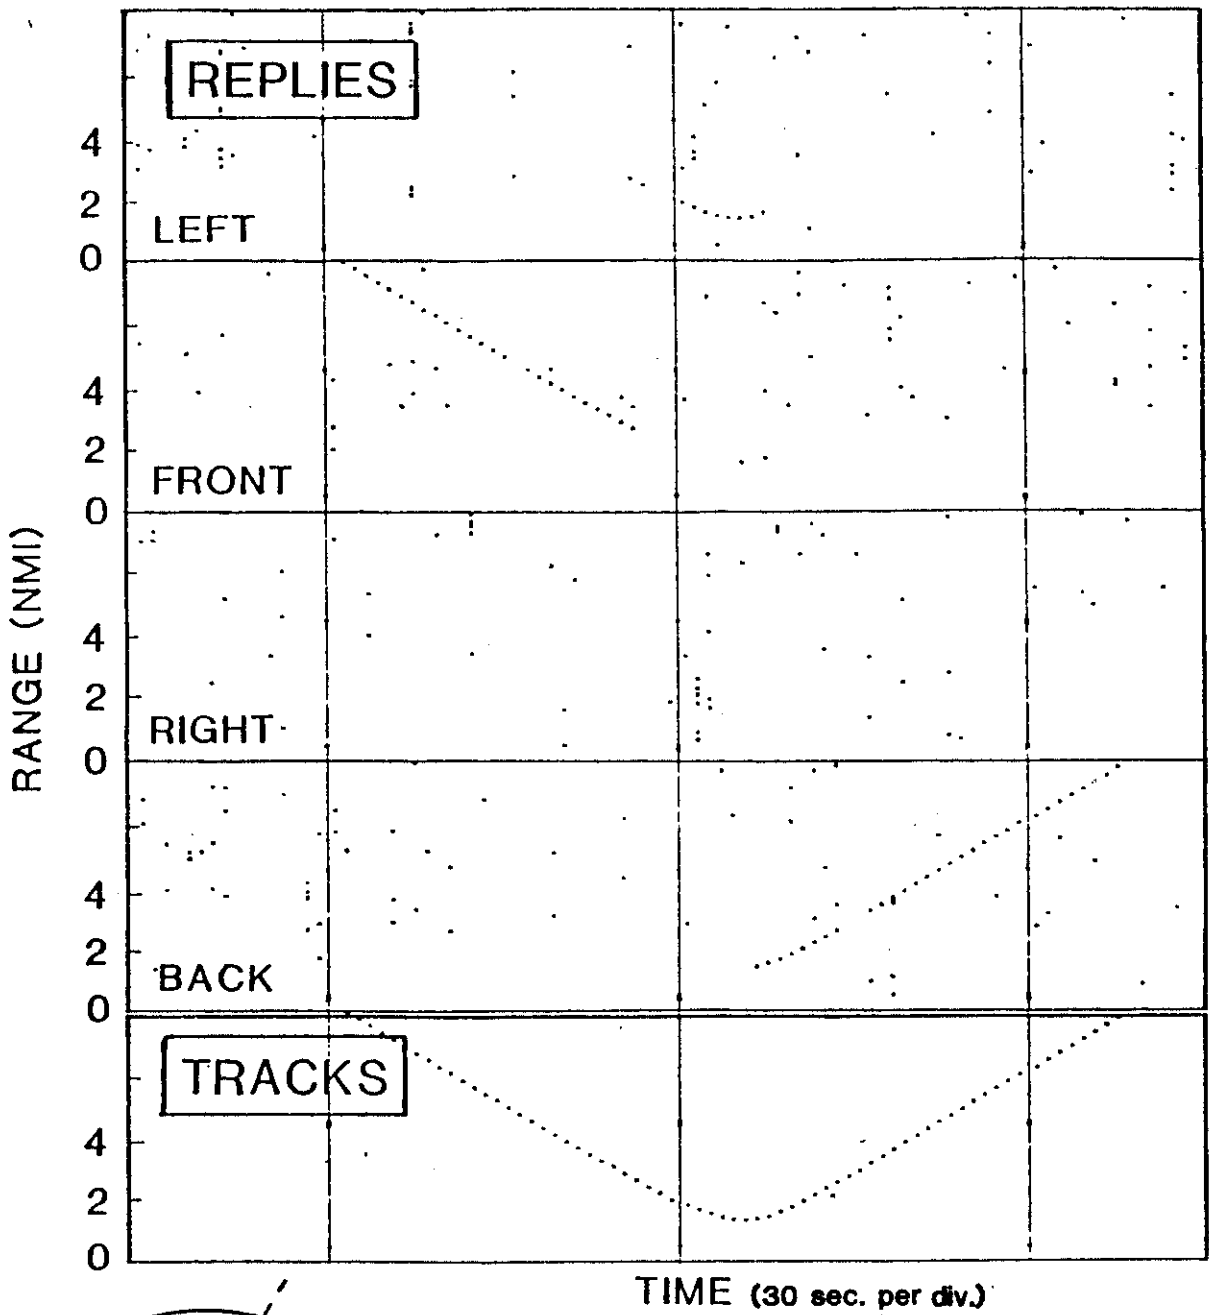


Fig. 3-12. Directional interrogation - airborne data.

3.2.7 Summary

These initial airborne experiments proved to be very worthwhile. They revealed three new mechanisms:

- early Mode A replies
- beam limiting near threshold
- late Mode C replies

all of which relate to the combined use of directional interrogation with high resolution whisper-shout, and whose understanding is important to successful use of directional interrogation. Understanding these mechanisms led immediately to several changes in design:

- change to 1-pulse suppression
- increase P2 power
- add filter in surveillance processing to eliminate late Mode C replies

With these changes in place, and with the assurance provided by examples as in Fig. 3-12 that the behavior of directional interrogation is reasonable, the next step was to conduct further airborne measurements in high density airspace. This was carried out by flights in the LA Basin as described in Sec. 3.5.

3.3 Role of the Bottom Antenna

As of January 1981, there had been no instances in which a false track caused a false alarm or modified a real alarm. This was encouraging since the airborne testing had amounted to several hundred hours of experience by that time. Even so, it was realized that Mode C false tracks do occur and that therefore some false and modified alarms would eventually occur. During the next two years, the airborne experience increased by many more hundreds of hours, and in that time, several instances of false and modified alarms have been observed. The data recorded in Piedmont* aircraft, for example, includes about 900 hours, and in this data there is one instance of a modified alarm and no instances of isolated false alarms. In addition, a considerable amount of testing has been done by the FAA Technical Center on the East Coast and in the Chicago area, and by Lincoln Laboratory in the Boston area. In this additional data there have been 8 instances of false alarms.

* In the Piedmont Phase I operational evaluation a TCAS II unit was installed on two Boeing 727 aircraft and carried during normal operations. The TCAS II advisories were not displayed to the pilots.

These false alarms have been studied individually and categorized according to the mechanisms causing the false tracks. The results are given in Table 3-1. The results show that the largest single source of false alarms was multipath. That is, for a real aircraft that is being tracked, reflection from the ground or water gave rise to a second track.

Since multipath-induced false tracks are mainly associated with the TCAS-II bottom antenna, it became appropriate to consider reducing the role of the bottom antenna. By reprocessing the recorded data from all of the instances of multipath false alarms, it was found that 4 of the 5 occurrences would have been eliminated by deleting the 3 highest-power bottom interrogations (that is, by reducing the bottom antenna interrogation power by 18 dB).

In considering a reduction of the role of the bottom antenna to reduce false tracks, it is necessary to know what the effect would be on the reliability of tracking real aircraft. An experiment was set up to gather airborne data for a performance comparison between a design using top and bottom antennas equally and a design that reduces the role of the bottom antenna. The interrogation sequences to be compared were selected to have the same total number of interrogations and the same power-sum*, both of which are quantities constrained by interference limiting (Sec. 5.1). The results of several measurements of this type showed that reduced-bottom designs perform nearly as well as the equal-use design, having surveillance reliability that is less by only about 2 or 3 percent while reducing false track rate by a large factor. In one of the experiments (Fig. 3-13), the reduced-bottom design is the whisper-shout sequence being adopted for TCAS II (see Fig. 3-15), and here the performance reduction is just 2.3 percent (of track-seconds for aircraft within $\pm 10^\circ$ in elevation angle).

Since the reduced-bottom design achieves a reduction of about 5:1 in false tracks with less than a 3% reduction in real tracks, it has been included in TCAS II.

3.4 Power Reduction

In very high density airspace, closing speeds are reduced and thus the range requirements of TCAS II are reduced. Under these conditions it should be possible to reduce the interrogation power level. Indeed, to conform with the interference limiting standards, it will be necessary in some cases to reduce power by as much as 6 dB. It was important to determine the amount of degradation in surveillance reliability that will result.

This has been addressed by both analysis and airborne measurements. The analysis uses the method documented in Ref. 4. The airborne data was obtained by reprocessing whisper-shout data already recorded, omitting the higher power

* Power sum is the sum over a 1 second period of the interrogation powers.

TABLE 3-1
SURVEILLANCE FALSE ALARMS

	Piedmont data (900 hours)	other airborne data
Isolated false alarms	0	1 - synch. garble 3 - other
Modified real alarms	1 - multipath	4 - multipath

COMPARISON:

TOP	14	VS.	24
BOTTOM	14		4

DATA:

New York area, 40 min., 13 August 1982

RESULT OF DECREASING BOTTOM ANTENNA ROLE:

PROBABILITY OF TRACK decrease by 2.3%

FALSE TRACKS decrease by a large factor . *

*5:1 reduction of false alarms in Piedmont Phase I data.

Fig. 3-13. Role of bottom antenna - airborne data.

levels. The airborne results for a 6-dB power reduction are summarized in Fig. 3-14 together with the analytical results. The quantity plotted is the amount of decrease in the percentage of aircraft in track. The agreement between calculation and measurement is reasonably good considering the variability in the data points. The data shows that when interrogation power is reduced by 6 dB, it is still possible to achieve effective surveillance at ranges up to about 5 nmi.

3.5 Airborne Measurements in the Los Angeles Basin

After investigation of high density surveillance techniques individually, the next step was to assemble these techniques into a baseline design and test the design by flying in the Los Angeles Basin.

The measurements were conducted as described in Sec. 3.2.2. The baseline directional design for surveillance in Mode C has the characteristics listed in Fig. 3-15, with the exception that it was not possible to key MTL to whisper-shout using this directional equipment (see Sec. 2.2). The baseline omnidirectional design is the same except for:

- full power = 54 dBm
- full sensitivity = -74 dBm
- whisper-shout, top - 24 levels (see Fig. 3-15, top-forward)
bottom - as in Fig. 3-15
- MTL keyed to whisper-shout, as in Fig. 3-15

3.5.1 Truth

The measurements were based on targets of opportunity. Use of data from ground based sensors for establishing a data base of truth was considered. However, in view of the poor surveillance reliability of such ground based equipment relative to the reliability of the experimental airborne equipment, and the fact that the test aircraft had two independent operating sensors using two pairs of antennas, it was decided that truth would best be derived from the data tapes recorded by the two TCAS interrogators. This was done using a manual process involving a number of computer-generated plots of replies and tracks.

3.5.2 Flight Path

The flight path through the LA Basin is shown in Fig. 3-16. It passed directly through the Long Beach area which, based on earlier data (Ref. 5), was expected to be the location of highest aircraft density. The flight path also passed over LA International Airport (LAX), and through the San Fernando Valley, passing between the general aviation airports at Van Nuys and Burbank, which are well known for high density of general aviation traffic.




Fig. 3-14. Performance using reduced power.

- Top antenna
- 4 beams (forward, right, left, aft)
 - 90° beamwidth
 - transmit SLS, P1-P2 crossover at approx. $\pm 60^\circ$
 - angle-of-arrival on reception
- Bottom antenna
- omnidirectional monopole
- Interrogation power
- top-forward radiated power at azimuth peak: +55 dB relative to a 0 dBm monopole
 - bottom: 54 dBm radiated
- Receiver MTL
- top-forward, at azimuth peak: -75 dB relative to a 0 dBm monopole
 - bottom: -74 dB relative to a 0 dBm monopole
- Whisper-shout
- top-forward, 24 levels (0 dB, see table)
 - top-right, 20 levels (-4 dB, table minus first 4 entries)
 - top-left, 20 levels (-4 dB, table minus first 4 entries)
 - top-aft, 15 levels (-9 dB, table minus first 9 entries)
 - bottom, 4 levels (-18 dB, see table)

power levels in dB relative to full power, full sensitivity			
index	interrogation power	suppression power	receiver MTL*
1 (top)	0	-3	0
2 (top)	-1	-3	0
3 (top)	-2	-5	0
4 (top)	-3	-5	0
5 (top)	-4	-7	0
6 (top)	-5	-7	0
7 (top)	-6	-9	0
8 (top)	-7	-9	-1
9 (top)	-8	-11	-2
10 (top)	-9	-11	-3
11 (top)	-10	-13	-4
12 (top)	-11	-13	-5
13 (top)	-12	-15	-6
14 (top)	-13	-15	-7
15 (top)	-14	-17	-8
16 (top)	-15	-17	-9
17 (top)	-16	-19	-10
18 (top)	-17	-19	-11
19 (top)	-18	-21	-12
20 (top)	-19	-21	-13
21 (top)	-20	-23	-14
22 (top)	-21	-23	-15
23 (top)	-22	-25	-16
24 (top)	-23	none	-17
1 (bot.)	-18	-21	-12
2 (bot.)	-20	-23	-14
3 (bot.)	-22	-25	-16
4 (bot.)	-24	none	-18

*not actually implemented in the Dec. 1982 tests

Fig. 3-15. Baseline TCAS II design for testing in LA.




Fig. 3-16. Flight path in the LA basin.

Altitude was constant at 5500 feet for about 50% of the data and 8500 feet for the remainder. In addition, several takeoffs and landings were included in the mission; each day's flying included two takeoffs from LAX (where the Boeing 727 was based), two landings at LAX, and a low approach at Long Beach.

3.5.3 Aircraft Density

The bulk of the flying was on a weekend (4-5 December 1982) so as to experience the highest aircraft density. Fortunately, the weather was relatively clear due to a severe storm that had passed through the region several days before. It was good flying weather, conducive to a high density of aircraft.

The data tapes show that the aircraft density was in fact quite high. Figure 3-17 shows density values observed during one pass of the route from north to south. The average density (including all transponder equipped aircraft) is seen to be about 0.1 per nmi². Peaks over 0.2 were observed occasionally. About half of these aircraft are altitude reporting. These values are generally consistent with density measurements made previously (Ref. 5).

3.5.4 Advisory Rate

A number of instances were observed in which an aircraft passed close by. In many of these cases, the aircraft came close enough to trigger a traffic advisory (TA) or resolution advisory (RA). The test aircraft did not respond to these RAs. Four such instances occurred during the time period plotted in Fig. 3-17, and these are marked in the figure.

The overall rate of RAs was 2.2 per hour, which is very high relative to the rate that would be experienced during an operational flight. For example, in the Piedmont Phase I flights, the RA rate was 1/37 hours. This difference is largely a consequence of the flight path adopted for these experiments; the aircraft remained in the high density airspace and at low altitude all the time, whereas an operational flight is in such airspace only a small fraction of the time.

3.5.5 Antenna Problem

Several months after the mission, it was discovered that a problem had developed in the directional antenna subsystem. The problem was a leakage of water into both top and bottom antenna units. As a result, the antenna patterns were distorted and may have also changed with time to some extent. An estimate of the top antenna patterns as they existed during the LA mission is shown in Fig. 3-18. These patterns were obtained by an indirect technique that makes use of detailed whisper-shout data. Figure 3-18 should be regarded as an approximation since azimuth estimates made by the same antenna were used in constructing these figures. The front beam is seen to be much higher in gain than the other three beams, whereas, by design, all four were to be




Fig. 3-17. Aircraft density in Los Angeles.




Fig. 3-18, Measured antenna patterns.

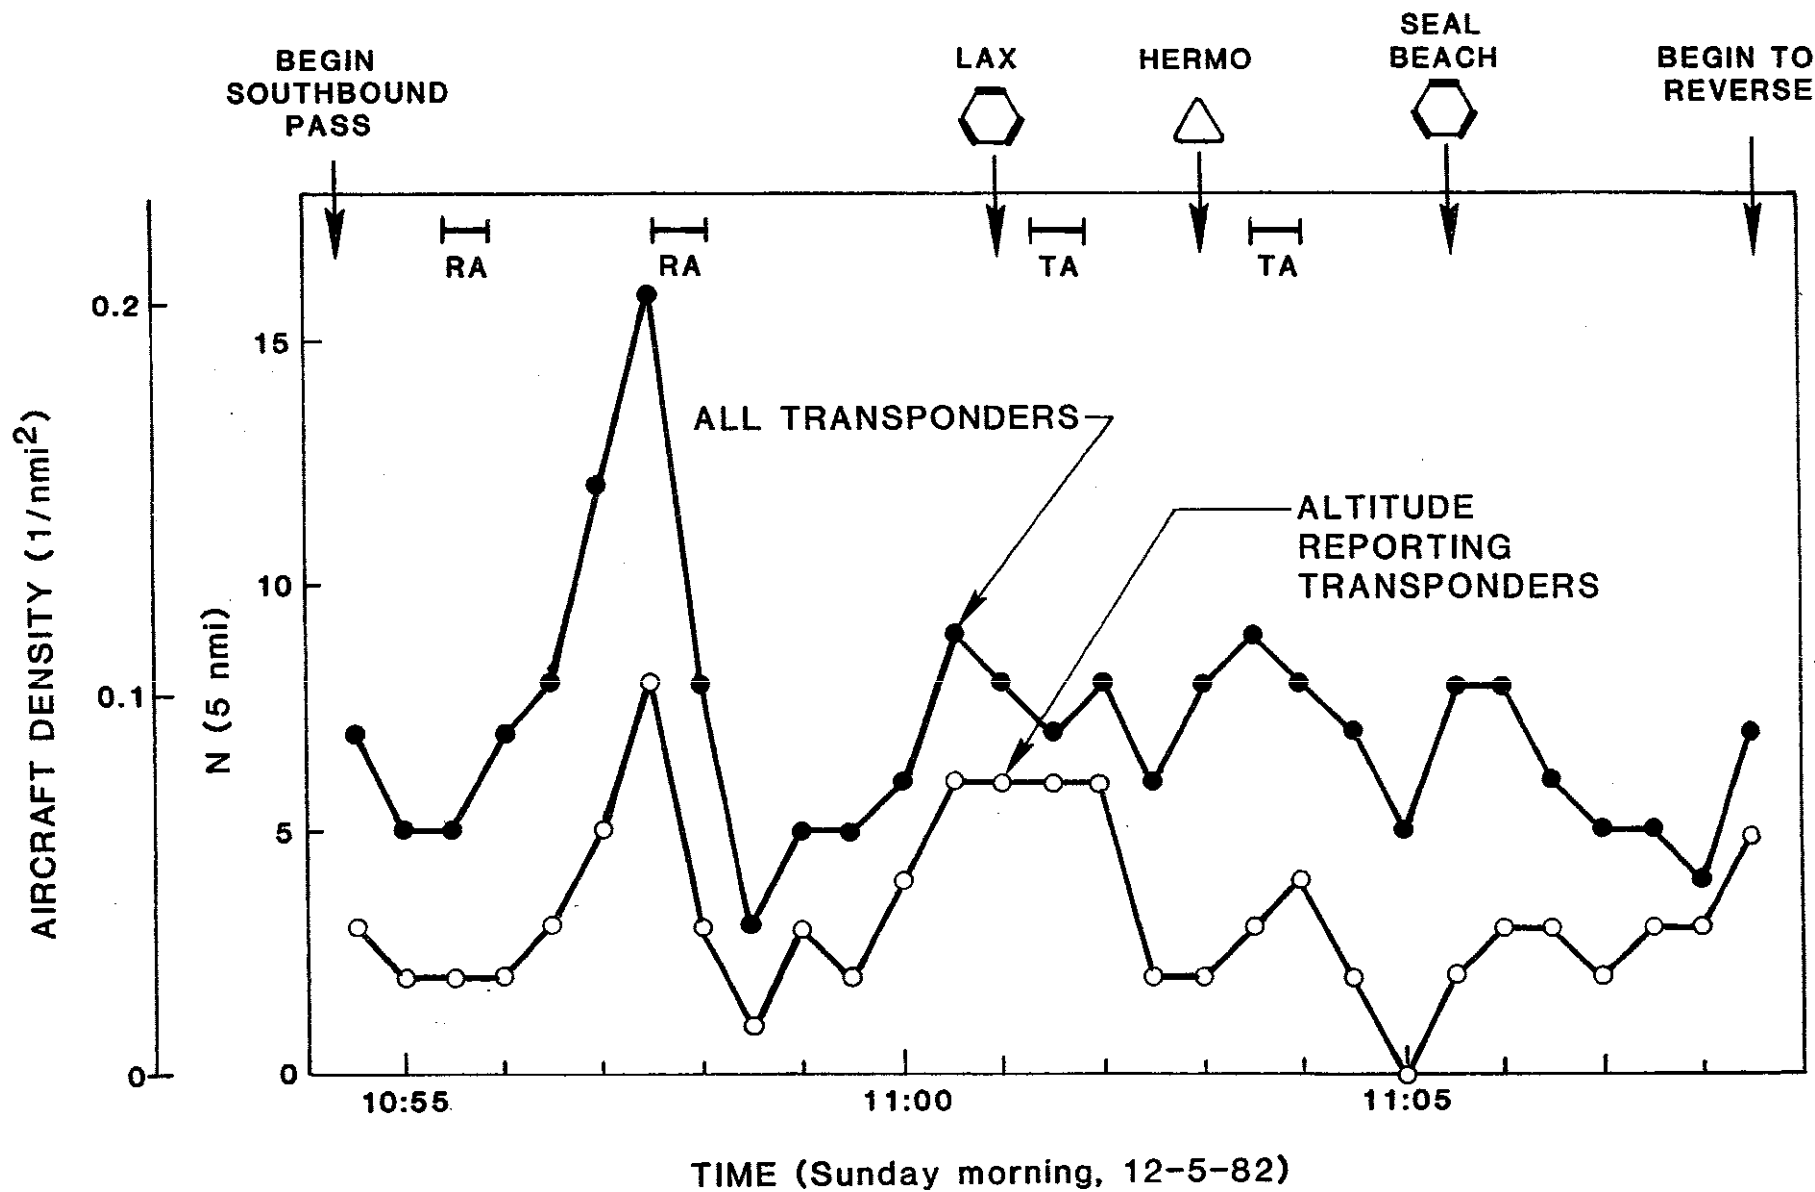


Fig. 3-17. Aircraft density in Los Angeles.

identical. It is also seen that the front beam is narrower than expected. Nevertheless the antenna did succeed in directionally interrogating target aircraft and in producing azimuth measurements that appear to be serviceable in spite of the water leak.

3.5.6 Case Studies

Contained in the data recorded in LA are a number of close encounters that occurred by chance. A set of 19 close encounters that occurred in a 2-hour period was analyzed in detail, where the criterion for being a close encounter was that the aircraft came within 2 nmi in range while being within 1200 ft in altitude.

Performance in tracking these aircraft, each for the 50-second period leading up to the point of closest approach, is shown in Figs. 3-19 and 3-20. In Fig. 3-19 each D signifies the event that the target aircraft is in track by the directional unit for one scan (one second). The figure also lists the aircraft density within 5 nmi during the encounter. In most of the encounters the target was in track continually throughout the 50 sec. period. There were a few instances of gaps or late track initiations. The overall percentage of time during which the target was in track in this data set is 97%. In Fig. 3-20 each O represents the condition of being in track by the omnidirectional unit for one scan. Qualitatively, the results here are the same, and here too the overall reliability is 97%. In both cases the performance is very good.

3.5.7 False Tracks

There were no false alarms in the LA data set. That is, at no time did a false track satisfy the conditions for generating a resolution advisory or a traffic advisory. There were, however, some false tracks. These were studied to determine the false track rate for tracks within $\pm 10^\circ$ in elevation angle and between 3 and 5 nmi in range. Results from 84 minutes of data are given in Table 3-2 (in the row marked "original design"). As a percentage, the false track rates for both systems are much higher than the values seen in the bulk of earlier data. In particular, the omnidirectional system percentage is larger by 30:1 relative to the BCAS performance during the 1980 Eastern tour. There are a number of factors that would be expected to cause this percentage to be different.

Factors that would increase false track percentage: (1) Higher fruit environment, and flight path that stays constantly in high density. (2) More severe multipath environment in LA, and flight path that remains constantly at low altitude. (3) More whisper-shout interrogations, and as a result, more fruit replies for a given fruit environment. (4) Relatively high proportion of non-altitude-reporting aircraft in LA. These contribute to the false track rate* (numerator) but not to the number of

*Non-altitude-reporting aircraft contribute to false tracks, both with and without altitude. The effect of interest here is the contribution to tracks with altitude.

OVERALL RELIABILITY = 97%




Fig. 3-19. Nineteen close encounters in LA ~ directional performance.

OVERALL RELIABILITY = 97%

CASE DENSITY

1	0.09
2	0.06
3	0.11
4	0.11
5	0.12
6	0.18
7	0.15
8	0.18
9	0.18
10	0.17
11	0.08
12	0.11
13	0.23
14	0.12
15	0.08
16	0.09
17	0.03
18	0.17
19	0.15




Fig. 3-20. Nineteen close encounters in LA - omnidirectional performances.

TABLE 3-2
FALSE TRACK RATE, LA BASIN

	DIRECTIONAL	OMNIDIRECTIONAL
ORIGINAL DESIGN	487 track sec. 6.7%	214 track sec. 2.9%
IMPROVED DESIGN	79 track sec. 1.1%	139 track sec. 1.9%

Notes:

- traffic = 7350 aircraft seconds
- For comparison, in the 1980 Eastern Tour, the false track rate was 0.1%.

altitude-reporting aircraft (demoninator), and so tend to increase the percentage.

Factors that would decrease false track percentage: (1) Reduced use of bottom antenna.

Two changes to the surveillance algorithms appeared to be warranted and were tried. One was a change in the multipath elimination algorithm to permit it to work with non-sea-level reflectors. Another change (applicable only to the directional unit) is azimuth filtering. This filtering discards any reply whose azimuth is inconsistent with the interrogation direction. Together, these changes reduced the false track rate considerably, to the values given in the second row of Table 3-2.

Such changes would be expected to degrade detection performance to some extent. However, it was found that the effects on surveillance reliability were insignificant, and in fact the excellent performance shown in Figs. 3-19 and 3-20 was obtained after these changes were made. Thus these changes have been adopted into the baseline TCAS II design.

3.5.8 Statistical Analysis

A statistical analysis was undertaken with the goal of assessing surveillance reliability as a function of traffic density and estimating the degree of improvement attributable to the directional antenna. The data set was divided into one-minute segments, and for each the maximum traffic density was determined. For this purpose, traffic density was computed as the number of aircraft between 2 and 5 nmi divided by $2l\pi$.^{*} The aircraft count included all transponder equipped aircraft, whether or not they were altitude reporting. The counting involved a detailed manual procedure based on computer plots of replies and tracks from both experimental systems. Probability of track, $P(T)$, was estimated as the percentage of aircraft-seconds during which the aircraft was in track, limiting attention to aircraft within $\pm 10^\circ$ in elevation, angle between 3 and 5 nmi in range, and for which both own aircraft and the target aircraft were at least 600 feet above ground level.

This study was performed omnidirectionally (that is, without noting the azimuths of the targets), and for this reason the same peak power was used in each of the four beams. The baseline TCAS II design, on the other hand, uses different powers in the four beams: highest in the front, less to the sides, and still less aft (Fig. 3-15). Thus relative to the baseline design, additional interrogations in the back and side beams were added for this study.

The results are given in Figs. 3-21 and 3-22 along with a curve showing measured BCAS performance for comparison (Ref. 3). These results were obtained prior to the algorithmic changes associated with false tracks and prior to a discovery that the lowest power omnidirectional interrogation had

^{*} $2l\pi$ is the area of the annular ring between 2 and 5 nmi.




Fig. 3-21. Performance measured in LA - directional design.




Fig. 3-22. Performance measured in LA - omnidirectional design.

inadvertently been omitted. When the data were reprocessed using a -18 dB interrogation as a replacement for the missing interrogation and using the revised algorithms, the overall average value of P(T) for the omnidirectional design rose from 90% (as marked in the figure) to 92%. For the directional design the average remained at 90%.

The data in Figs. 3-21 and 3-22 suggest the following observations: (1) for both omnidirectional and directional units, performance is significantly better than that of the original BCAS design. (2) Because of the scatter of data points, the rate of degradation vs. traffic density is not evident in either case. It would take an environment considerably more dense before a significant degradation would become apparent. (3) The results for the directional unit do not indicate an improvement relative to the omnidirectional unit. The degraded antenna performance together with insufficient aircraft density may account for this. A more detailed examination of directional performance is described in the "whisper-shout profiles" section below.

3.5.9 Effect of Elevation Angle

In the course of the statistical analysis of probability of track, P(T), it became evident that many of the "targets-of-interest" ($\pm 10^\circ$) were at very low altitude, near the cutoff at -10° . A quantitative study (Fig. 3-23) confirmed that, in fact, about one half of all targets-of-interest were in the band -5 to -10° . This observation suggested that the $\pm 10^\circ$ definition may lead to a pessimistic assessment of TCAS II, relative to its performance in an operational environment.




An elevation angle comparison was made of this data vs. the elevation angles experienced in case studies of real mid-air collisions*, and vs. the resolution advisory encounters in the Piedmont Phase I data. The comparison indicates (Fig. 3-23) that indeed the $\pm 10^\circ$ analysis is pessimistic; an analysis based on a $\pm 5^\circ$ definition would be more representative of operational performance.

The P(T) analysis was repeated using a $\pm 5^\circ$ elevation angle definition for targets-of-interest, and a significant increase in the values of P(T) resulted. The overall average, which was 89% for $\pm 10^\circ$, rose to 95% for $\pm 5^\circ$. This result is more consistent with the excellent performance seen above in the 19 case studies.

3.5.10 Whisper-Shout Profiles

One of the main objectives of the airborne measurements in Los Angeles was to assess TCAS performance using directional interrogation, and in particular to assess the degree of improvement relative to use of omnidirectional whisper-shout. The statistical study of P(T) vs. density did not, however, reveal any significant improvement achieved by the directional

*From a set of 15 actual mid-air collisions, Ref. 6., pp. C-1 through C-3.



* Note: Of the 15 actual mid-air collisions identified in Ref. 6, two were in this range band, 3 to 5 nmi, at 25 sec. prior to collision: ③ which occurred at Duarte, CA, and ⑬ which occurred at Carmel, N.Y. Since ⑬ involved two air carrier aircraft, it is shown here as two points, corresponding to the use of TCAS by first one aircraft and then the other. Since ③ involved just one air carrier aircraft, it is shown here as one point.

Fig. 3-23. Elevation angles of TCAS targets.

design. This result was partially true because both designs performed well in the LA environment. Measurements in a higher density environment (if one existed) might have revealed a performance improvement. The hoped-for improvement was explored further by means of an indirect method based on an analysis of whisper-shout characteristics. This method makes use of the whisper-shout profiles shown in Figs. 3-24 and 3-25. These figures display the number of replies per interrogation as a function of whisper-shout index.

3.5.10.1 Fruit Rate

The first step was to try to distinguish between fruit and synchronous replies since their effects are very different; it is only the number of synchronous replies that may be expected to be reduced through the use of directional interrogation. To estimate fruit rate, a whisper-shout profile was formed for the range band 0.1 to 1.1 nmi, a close-in region where few synchronous replies would be received. The results plotted in Fig. 3-24 have characteristics that would be expected: less fruit during sweeps in which MTL was elevated (Fig. 3-15). Quantitatively, the relationship agrees with a uniform-in-range model of aircraft traffic.

The fruit rate received by the directional unit, 3200/sec. (Fig. 3-25), was considerably less than that received by the omnidirectional unit, 11200/sec. (Fig. 3-24). This implies a reduced sensitivity, which is probably a result of the degradation in antenna performance (due to water) described above. The amount of the degradation can be estimated as follows. According to antenna measurements made by Dalmo Victor prior to installation, the peak gain of the directional antenna was +2 dB relative to an ideal monopole. Thus the azimuth-average gain was about +1 dB relative to a monopole. Cable losses were 3 dB for both systems. MTL values were measured as:

MTL, directional unit = -75 dBm

MTL, omnidirectional unit = -79 dBm

Together, the differences add up to:

Antenna gain	+1 dB
Cables	0 dB
Receiver MTL	-4 dB
<hr/>	
Total	-3 dB

That is, the measurements of the equipment prior to airborne testing indicated that the directional unit would be less sensitive to fruit by 3 dB. The airborne results in Figs. 3-24 and 3-25 imply, however, that the directional unit was actually less sensitive to fruit by about 10 dB (this value obtained by noting in Fig. 3-24 the omnidirectional MTL shift such that the fruit rates are equal). The 7 dB difference between the prediction (3 dB) and the measurement (10 dB) is an estimate of the degradation attributable to the water in the antenna.




Fig. 3-24. Whisper -shout profiles for omnidirectional interrogator.




Fig. 3-25. Whisper - shout profiles for directional interrogator.

It may also be noted from Fig. 3-24 that under nominal conditions (omnidirectional MTL = -77 dBm, cable = 3 dB), the fruit rate would be about 9000 replies/sec.

3.5.10.2 Synchronous Replies

The lowest of the three curves plotted in Fig. 3-24 and in Fig. 3-25 can be considered to indicate fruit replies, and the differences between the other data points and the lowest curve can be considered to indicate synchronous replies. Synchronous reply data are shown for two range bands, 1 to 3 nmi and 3 to 5 nmi.

3.5.10.3 Results

Examination of the plotted data leads to the following observations.

- (1) The directional data resembles the bell-shaped curve seen previously in similar data (Fig. 3-3) except that the fall-off on the right side is not apparent. This is probably due to the sensitivity degradation caused by the antenna.
- (2) Both units exhibit an alternating high-low characteristic, which is to be expected as a result of the alternation between 2-dB and 3-dB whisper-shout bins (Fig. 3-15). This provides additional evidence that a change in bin size from 3 dB to 2 dB produces a significant reduction in number of replies per interrogation.
- (3) A dip is evident in the omnidirectional data around the region where interrogation attenuation is 14 dB. This has been explained by consideration of previous measurements showing the accuracy of the whisper-shout attenuator. This data shows a discontinuity in the attenuator characteristic, occurring between 15 dB and 16 dB (presumably because of the switching of all 5 bits at that transition). Of all the whisper-shout interrogations, only those at 13 dB, 14 dB, and 15 dB span this discontinuity, and because they do span it, they would be expected to have bin sizes smaller than nominal. The dip seen in Fig. 3-24 agrees with this expectation.
- (4) For the omnidirectional design, the whisper-shout sequence does not extend sufficiently low in power to reach a point where reply density is small. The lowest power interrogation, at -23 dB (inadvertently omitted in these measurements) would gather an undesirably large number of replies. It may be concluded that the sequence should be extended at the low end to approximately -30 dB.
- (5) The average number of replies to one interrogation has in fact been reduced by the introduction of directional interrogation. The reduction factor, based on the region of highest reply density, and calculated separately for the two range bands, is:

Reduction factor = 2.4 for range = 1 to 3 nmi

2.4 for range = 3 to 5 nmi

This is close to the anticipated improvement factor of 2.5 (Sec. 3.2).

In summary, examination of these whisper-shout profiles has yielded several useful results: a measure of the fruit environment in the LA Basin; an estimate of the degradation in receiving sensitivity resulting from water in the directional antenna; additional evidence of the effectiveness of whisper-shout; a conclusion that the whisper-shout sequence should be extended at the low end; and an estimate of the degarbling effectiveness of directional interrogation.

4. SURVEILLANCE IN MODE S

4.1 Introduction

4.1.1 Functional Requirements

The function of the Mode S surveillance processor is to identify and track Mode S-equipped aircraft. The implementation of this function is constrained by the requirement that the TCAS transmissions not cause undue interference to other services in the 1030/1090 MHz bands.

In Section 5.1, the above constraint is translated into limits upon the interrogation power and rate of the system. When the normal operation of the surveillance processor would cause these limits to be violated, the interrogation limiting algorithm described in Section 5.2 exercises pre-emptive control to ensure that they continue to be satisfied. Since the primary purpose of this control is to protect other aviation-related activities it does not ensure that the desired level of collision protection is maintained. Thus it is important that the design of the surveillance processor provide satisfactory collision protection in the required operating environment when this control is present. Since each Mode S aircraft is individually addressed, this becomes more difficult as the density of aircraft increases. The design used for BCAS, which emphasized early interrogation of all detected aircraft, cannot provide satisfactory collision protection for the aircraft densities in which TCAS is required to operate.

To satisfy the constraints and provide adequate collision protection in high aircraft densities it is necessary to restrict interrogations to only those aircraft that might pose a collision threat. The opportunity to distinguish between threatening and non-threatening aircraft without interrogating them is provided by the reception of Mode S transmissions that are either replies to other interrogations or are spontaneously emitted. The former are called fruit, the latter are called squitters. In particular, a crude measure of an aircraft's range is provided by the frequency with which the transmissions received from it exceed a given rate threshold. Also, aircraft altitude is contained within replies to surveillance interrogations. Thus an aircraft need be interrogated only if these parameters indicate that it could be a collision threat within the time interval that is required for planning and executing evasive maneuvers.

To see how this information can be used by the surveillance processor, it is helpful to think of each Mode S-equipped aircraft as falling into one of three categories as depicted in Fig. 4-1. Category I contains those aircraft that could become collision threats to the TCAS-equipped aircraft if evasive action is not taken. The immediacy of this possible threat dictates that aircraft in this category be interrogated regularly and tracked so that evasive maneuvers can be taken. At the other extreme, Category III contains

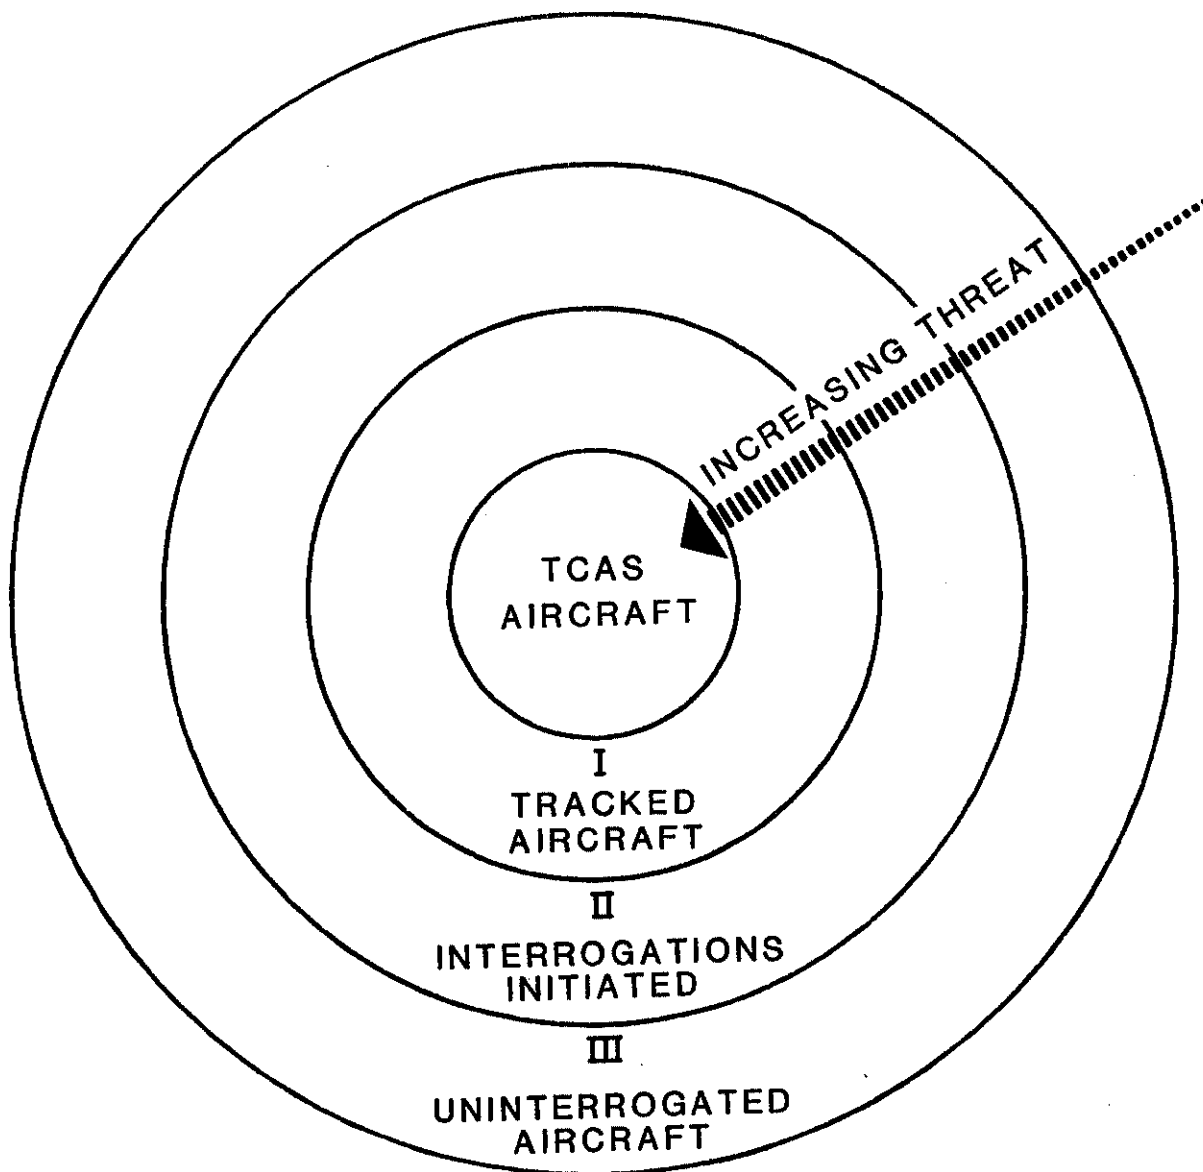


Fig.4-1. Levels of threat.

those aircraft that are rarely, if ever, interrogated. This may occur either because the unsolicited transmissions received from them indicate that they cannot become collision threats for some considerable time, or because little, if anything, is known of their presence. For the surveillance processor to provide acceptable collision protection it must only rarely, if ever, allow aircraft that are near-term threats to be assigned to Category III.

Finally, Category II contains those aircraft that were previously in Category III but whose threat potential, as assessed from their unsolicited transmissions, has increased to the point where more information concerning their trajectory must be obtained by interrogating them. This is a transient category. Aircraft are reassigned to either Category I or III after the interrogation has been made.

To obtain good collision protection the unsolicited information received from an aircraft must be processed so that the transition from Category III to II to I is accomplished in time to allow evasive maneuvers to be taken. However, to limit interference, as many aircraft as possible should be kept in Category III. If this is not done, the collision protection provided by the system may itself be seriously degraded by the interrogation limiting algorithm. Finally, the number of aircraft that are assigned to Category I should be as small as possible while ensuring that all collision threats are included in that category.

The algorithms that cause aircraft to be assigned to the three categories must strike a balance between these conflicting goals. Equally important are the interrogation patterns used in Categories I and II. A reduction in the power or rate of the TCAS interrogations will reduce interference to other services, but will also reduce the collision protection provided.

In the sections that follow, the design approach that led to a satisfactory balance is described. First, the broad structure of the system is specified. Then the design of the blocks within that structure is discussed in more detail. Most of the system parameters were determined either by application of design ground rules or by simulation studies. Finally, the performance of this design is verified by simulation and by using data from airborne encounters as inputs to a software implementation of the system.

4.1.2 System Structure

The categories described above correspond to a structure for the surveillance processor that involves four states* to which a detected aircraft

*The term state, rather than category, is used to differentiate the system's assessment of the threat posed by an aircraft from the actual threat.

can be assigned. The states are:

1. Monitor state
2. Acquisition state
3. Track state
4. Dormant state

The acquisition and track states correspond, respectively, to Categories II and I in Fig. 4-1. The two remaining states correspond to Category III. The monitor state is for aircraft that are judged to be non-threats based only upon the information gained from the reception of their unsolicited transmissions. The dormant state is for aircraft that have been judged to be non-threats after their range has been determined by interrogation.

The structure of the surveillance processor is related to the four states as shown in Fig. 4-2. Detected aircraft are initially assigned to the monitor state upon the detection of their unique ID. They remain in this state until the rate of reception of their unsolicited transmissions indicates either that:

1. They are so far removed from the TCAS aircraft that they are not an immediate threat to it:
- or
2. They may be a threat and the altitude information received from them re-enforces this conclusion.

In the first instance, the aircraft ID is removed from the system files and any further receptions of it are treated as though it had not previously existed. In the second instance, the aircraft is assigned to the acquisition state.

Aircraft that have been assigned to the acquisition state are interrogated until either:

1. An acceptable reply is obtained; or
2. It appears that such a reply will not be forthcoming.

In the first instance, the additional information obtained from the reply is used to more accurately assess the threat posed by the aircraft. The aircraft is then assigned to the track state if the threat is significant and the aircraft is assigned to the dormant state if it cannot become a threat for some considerable time*.

In the second instance, the aircraft is reassigned to the monitor state since continued interrogations may cause the interrogation limiting algorithm to degrade the collision protection against all aircraft. The reassignment

*As discussed in Sections 4.4 and 4.5, the monitor state is sometimes assigned.

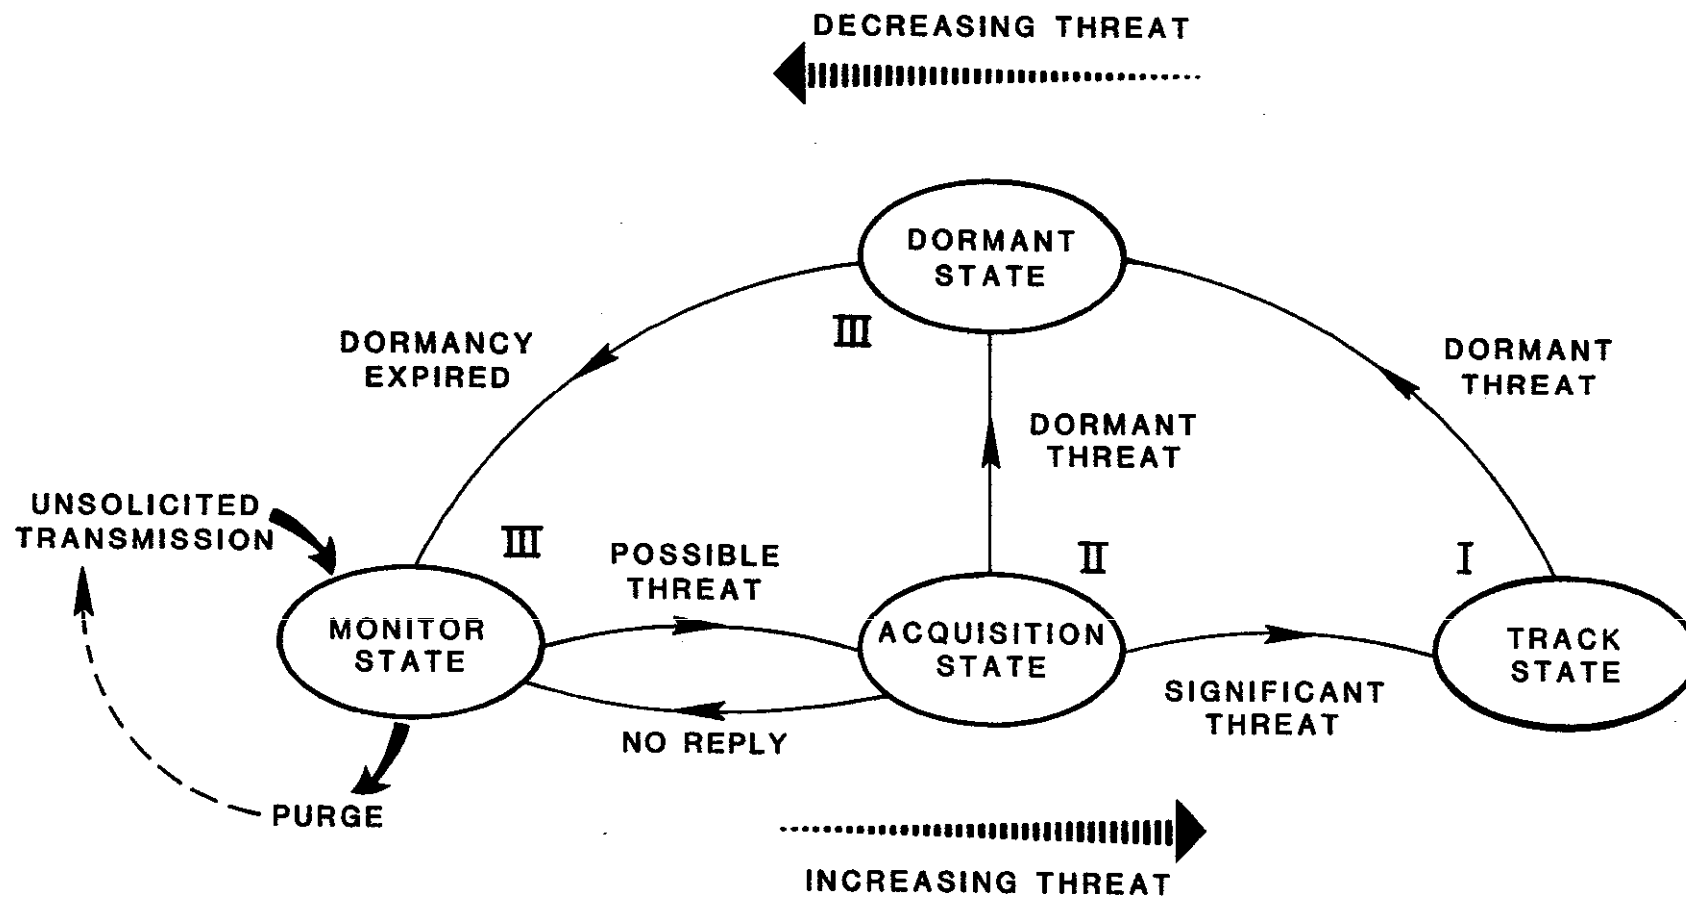


Fig.4-2. Structure of the Mode S surveillance processor.

of the aircraft to the monitor state would be of marginal value if it were unaccompanied by a change in the conditions under which the aircraft would, once again, be assigned to the acquisition state. However, for a properly operating system, the fact that a reply was not obtained from the aircraft implies that its range was greater than had been thought. Thus, it should not have been assigned to the acquisition state in the first place and should not be reassigned to the state until the reception of its transmissions indicate that its range has decreased significantly. Thus, the conditions under which an aircraft is changed from the monitor state to the acquisition state should depend upon the number of times it has been re-assigned to the monitor state after an unsuccessful interrogation. Similarly the number of unsuccessful interrogations for which the state is changed to monitor from acquisition should vary according to the number of times that change has recently been made.

When an aircraft has been assigned to the track state it is interrogated regularly and tracked. This process continues until it is certain that a collision with that aircraft cannot occur for some considerable time. The aircraft is then assigned to the dormant state*.

Targets assigned to the dormant state are not interrogated since they cannot become collision threats for some considerable time. This assignment is changed to the monitor state when there is any possibility that the aircraft has become a near-term threat, as indicated in Fig. 4-2.

4.1.3 Design Constraints

Given the system structure shown in Fig. 4-2, it remains to specify:

1. The algorithms that are used to determine when the aircraft state should be changed, and,
2. The operations performed for aircraft in each of the four states.

Both of these specifications are strongly influenced by the information that the system is allowed to use concerning the position, motion and capabilities of aircraft and the system parameters that can be varied dynamically.

To draw upon the experience obtained from flight tests of BCAS it was decided to constrain this study of minimum TCAS II design in a number of respects. These constraints are listed in Table 4-1 and are discussed below.

A major impact of the first group of constraints is to exclude TCAS designs that 1) measure received power levels to estimate aircraft range, 2) utilize on-board information concerning the TCAS aircraft that is not available either from the TCAS equipment itself or from the associated Mode S transponder and 3) measure aircraft bearing.

*As discussed in Sections 4.4 and 4.5, the monitor state is sometimes assigned.

of the aircraft to the monitor state would be of marginal value if it were unaccompanied by a change in the conditions under which the aircraft would, once again, be assigned to the acquisition state. However, for a properly operating system, the fact that a reply was not obtained from the aircraft implies that its range was greater than had been thought. Thus, it should not have been assigned to the acquisition state in the first place and should not be reassigned to the state until the reception of its transmissions indicate that its range has decreased significantly. Thus, the conditions under which an aircraft is changed from the monitor state to the acquisition state should depend upon the number of times it has been re-assigned to the monitor state after an unsuccessful interrogation. Similarly the number of unsuccessful interrogations for which the state is changed to monitor from acquisition should vary according to the number of times that change has recently been made.

When an aircraft has been assigned to the track state it is interrogated regularly and tracked. This process continues until it is certain that a collision with that aircraft cannot occur for some considerable time. The aircraft is then assigned to the dormant state*.

Targets assigned to the dormant state are not interrogated since they cannot become collision threats for some considerable time. This assignment is changed to the monitor state when there is any possibility that the aircraft has become a near-term threat, as indicated in Fig. 4-2.

4.1.3 Design Constraints

Given the system structure shown in Fig. 4-2, it remains to specify:

1. The algorithms that are used to determine when the aircraft state should be changed, and,
2. The operations performed for aircraft in each of the four states.

Both of these specifications are strongly influenced by the information that the system is allowed to use concerning the position, motion and capabilities of aircraft and the system parameters that can be varied dynamically.

To draw upon the experience obtained from flight tests of BCAS it was decided to constrain this study of minimum TCAS II design in a number of respects. These constraints are listed in Table 4-1 and are discussed below.

A major impact of the first group of constraints is to exclude TCAS designs that 1) measure received power levels to estimate aircraft range, 2) utilize on-board information concerning the TCAS aircraft that is not available either from the TCAS equipment itself or from the associated Mode S transponder and 3) measure aircraft bearing.

*As discussed in Sections 4.4 and 4.5, the monitor state is sometimes assigned.

TABLE 4-1
DESIGN CONSTRAINTS

I. Collision Information Used by TCAS: Limited to:

1. That obtained from on-board transponder
2. That obtained from data in transmissions from other transponders
3. Range

II. Design Features Adopted from BCAS

1. Filtering on confidence bits and consistency checks
2. Division of time between interrogation/replies and listening for unsolicited transmissions
3. One-second scans
4. Tracking algorithms
5. Antenna diversity switching
6. Omni-directional operation

The first exclusion was adopted to permit use of the BCAS reply processor design in TCAS. The second exclusion was imposed because of the difficulty and expense of providing interfaces to obtain other information. It is a significant exclusion, for if the airspeed of the TCAS-equipped aircraft were available, use of the relative bearing of aircraft would improve the performance of the system markedly. The third exclusion was adopted for two reasons. First, without information concerning the airspeed of the TCAS-equipped aircraft, bearing information is of limited use. Second, preliminary analysis indicated that the operating requirements could be met without its use. Thus in the interests of system simplicity it was excluded.

The result of the above constraints is that the information inputs to the surveillance processor are: the detected bit pattern of solicited and unsolicited transmissions from Mode S transponders, the measured ranges of aircraft that have been successfully interrogated, and the altitude and maximum capable airspeed of the TCAS aircraft itself. In all of these regards the TCAS design is similar to the BCAS design. Similarities also exist at a more detailed design level as is indicated in Table 4-1.

In particular, the same filtering of detected bit patterns is employed to remove those that are clearly erroneous. Also, the system listens for unsolicited transmissions whenever it is not engaged in an interrogation/reply cycle and during such cycles the listening window is that used in BCAS. These time allocations are organized within one second time-frames called scans. The BCAS tracking algorithms are also assumed to be employed, although they have no direct impact on the work reported here. Finally, diversity antennas are used with the BCAS diversity switching algorithm. Although capable of directional operation, the antennas are assumed to be used in a non-directional mode. This last constraint is imposed more for system simplicity than to capitalize upon the BCAS design.

The TCAS design differs from the BCAS design in the areas enumerated in Table 4-2. The first difference pertains not to the TCAS equipment itself, but to the "squitters" transmitted by Mode S transponders. The reasons for this change are discussed in Section 4.2. Items 2, 3, and 4 in the table all reflect the design changes that were made to ensure satisfactory operation at the high aircraft densities in which TCAS is intended to operate.

Roughly stated, the sensitivity of the system is controlled by the minimum triggering level (MTL) that is used for the reception of unsolicited transmissions from Mode S transponders. It is kept at the most sensitive setting for which the interference limiting constraints of Section 5.1 are satisfied.

Since the interrogation and reply links are of roughly equal quality (at least in the absence of heavy Mode C traffic), the power level used to interrogate an aircraft is related to the MTL at which it was detected. If the MTL was 5 dB above the most sensitive (nominal) setting, the interrogation power used will be 5 dB below the maximum (nominal) value. On the other hand, the maximum receiver sensitivity is always used in listening for the reply to

TABLE 4-2

AREA IN WHICH TCAS DIFFERS FROM BCAS

1. Squitter Format
2. Control of MTL for Unsolicited Transmissions
3. Programming of Interrogation Power
4. Information Processing Algorithms
5. Error Correction

an interrogation. Maximum sensitivity is also used in listening for replies from aircraft that are in the track state, but the interrogation powers to these aircraft are related to their ranges in order to control interference.

The development of the algorithms that determine the state assigned to an aircraft is described in Sections 4.3 through 4.6. That development was the major task in the design of the surveillance processor.


The final listed change, error correction, was made for two reasons. First, at the high densities of interest here, Mode C fruit can cause the reliability of the reply link to be substantially less than that of the interrogation link. The use of error correction reduces the chance that this imbalance will compromise the collision protection provided by the system. Second, it is prudent to choose a robust design whenever it does not involve undue complexity. The use of error correction appears to be such a choice.

4.2 False Address Problem

TCAS equipment only addresses interrogations to aircraft whose ID's have been received. However, "false addresses" will sometimes be generated by fruit, multipath, and receiver noise, which corrupt the squitter signal received from a transponder. In fact, in the high density environments for which TCAS is intended, the squitters used by TCAS might generate and duplicate false addresses at a rate that would overburden the system memory and cause a significant number of interrogations to be addressed to non-existent aircraft.

To ensure that this does not occur, it was necessary to reduce the probability that a false address would be received repeatedly. This was done by changing the squitter to the Mode S All-Call format so that error detection could be used. As a consequence, altitude information is no longer contained within the squitters. Altitude information is now extracted from the Mode S surveillance replies that an already identified transponder transmits in response to interrogations from other equipment, when such replies are available. When such replies are not available, for example when over the ocean, altitude remains unknown until the aircraft is interrogated.

The decision to change the squitter format was based on flight test data which suggested that false addresses were far more frequently created by single bit errors than they would be if the bit errors were statistically independent and identically distributed. An illustration of this is given in Fig. 4-3 which shows the number of times each bit of a Mode S reply was received erroneously. For a total of 18,500 receptions in low-density airspace, 5.6% of the replies had errors and about 40% of those errors involved just one out of the 56 reply bits. It is believed that most of these errors were due to multipath, as the fruit rate was low. The increased counts near the end of the reply also support this conclusion.



TIME DURATION = 651 SEC.

TOTAL NO. OF REPLIES DETECTED = 18500

NO. OF THESE IN ERROR = 1039 = 5.6%

NO. IN ERROR AND HAVING A REPEATING ADDRESS = 451 = 2.4%

NO. HAVING A 1-BIT ERROR = 403 = 2.2%

Fig. 4-3. Histogram of single-bit errors.

If these errors are not detected, the consequences are two-fold. First, they increase the computational load and memory requirements of the TCAS equipment. For example, if the error rate is 10% and 20 seconds elapse before a false address is purged from the system, there will be roughly twice as many false addresses in memory as real addresses. More significantly, because 40% of the errors involve only one bit, the rate of repeating false addresses will lead to many wasted interrogations in the high-density airspace for which the system is intended.

The wasted interrogation rate can be reduced somewhat by purging addresses from the processor sooner, but the detection rate then also suffers. Detection studies showed that addresses should not be purged less than 16 seconds following their first receipt. The curve of Fig. 4-4 shows that there will be as many interrogations transmitted to non-existent targets as to real targets when the average single-bit error rate reaches 10%. This is significant since TCAS will not achieve the desired high-density performance if the wasted interrogation rate approaches the valid interrogation rate.

In higher traffic densities the Mode S reply rate is higher and there are more interference replies to corrupt each Mode S reply. Realizing this, the squitter error rate was examined in a denser traffic environment. Figure 4-5 shows results from an encounter flown over New York City in September 1982. The top part of the plot shows the range of the Mode S aircraft as a function of time. The bottom half shows the rate of 1-bit errors detected by the top and bottom antennas on the TCAS aircraft. The rates fluctuated considerably and exceeded 10% a significant fraction of time.

The false address problem can be eliminated by using the Mode S All-Call format for Mode S acquisition. In the all-call format, address errors can be detected and corrected with high probability because the address is transmitted as part of the data field of the reply format and it is protected by an independent parity field, as shown in Fig. 4-6.

Using the Mode S All-Call format results in a slight increase in the Mode S fruit rate because, unlike surveillance replies, the all-call format is not transmitted routinely for other purposes by Mode S transponders. The periodic transmission of an all-call squitter thus adds to the existing Mode S fruit background. However, this additional fruit causes no significant degradation of ground surveillance (Ref. 10).

Another disadvantage of using the All-Call format for squitters is that it does not provide altitude information. However, altitude is not necessary in squitters since (in dense traffic, where altitude information is needed most) a Mode S surveillance fruit with altitude will usually be received shortly following the receipt of an all-call squitter. If a surveillance reply with altitude is not received soon after the squitter, TCAS can interrogate the target to determine its altitude and range.

False addresses interrogated per real aircraft




Fig. 4-4. Interrogations to repeating false addresses.

4-14




Fig. 4-5. Single-bit error rate observed in high density airspace.



Fig. 4- 6. Mode S all-call format.

4.3 Monitor State

4.3.1 Design Considerations

The information available for identifying possible collision threats is 1) the pattern of squitters and fruit received from the aircraft, 2) the altitude information conveyed by the fruit and 3) the number of unsuccessful attempts that have been made to acquire the aircraft.

If the altitude separation is sufficiently large and its rate of decrease is sufficiently small, no interrogation is needed. On the other hand, unless the available altitude information clearly indicates that the aircraft is a non-threat, the other available information must be examined if an interrogation is to be avoided. The most that can be inferred from this information is the degree to which the received power level does, or does not, exceed the detection threshold (MTL). Thus, loosely stated, the processor must decide whether or not the received power level is large enough so that the aircraft may be a threat and whether the link reliability is good enough so that acquisition should be attempted.

One approach to such decision problems that has been found to be effective in many instances is the Sequential Probability Ratio Test provided by Statistical Decision Theory (Ref. 8). An application of that test to the problem at hand suggests the following algorithm*.

Decision Algorithm.— Upon the first receipt of an aircraft ID the aircraft is assigned to the monitor state and a sum initialized at a value C is associated with it. Upon each succeeding receipt of the same ID, the sum is incremented by an amount z ; for each scan during which the ID is not received the sum is decremented by one. The process continues as long as the value of the sum exceeds 0 and is less than a constant Z . When the sum decreases to 0, the ID is purged from the system and any further receipt of that ID causes a newly initialized sum to be formed. When the sum equals or exceeds the constant Z , a test is performed to determine if the aircraft should be assigned to the acquisition state, unless the available altitude information now indicates that this is not necessary. The squitter processing used in BCAS is, in fact, a special case of this algorithm.

The operation of this algorithm is illustrated in Fig. 4-7 for three different example sequences of address detection.

The action of the algorithm on squitters and fruit differs in two ways. First, as discussed in Section 4.2, only squitters are used to enter an aircraft ID into the system. Detected fruit is processed only if its address is already contained in the system. Second, the assessment of the collision threat in altitude involves only the fruit (surveillance replies), since no altitude information is contained in the squitters.

*The detailed specification and performance of the algorithm is presented in Sections 4.3.2 and 4.3.3.

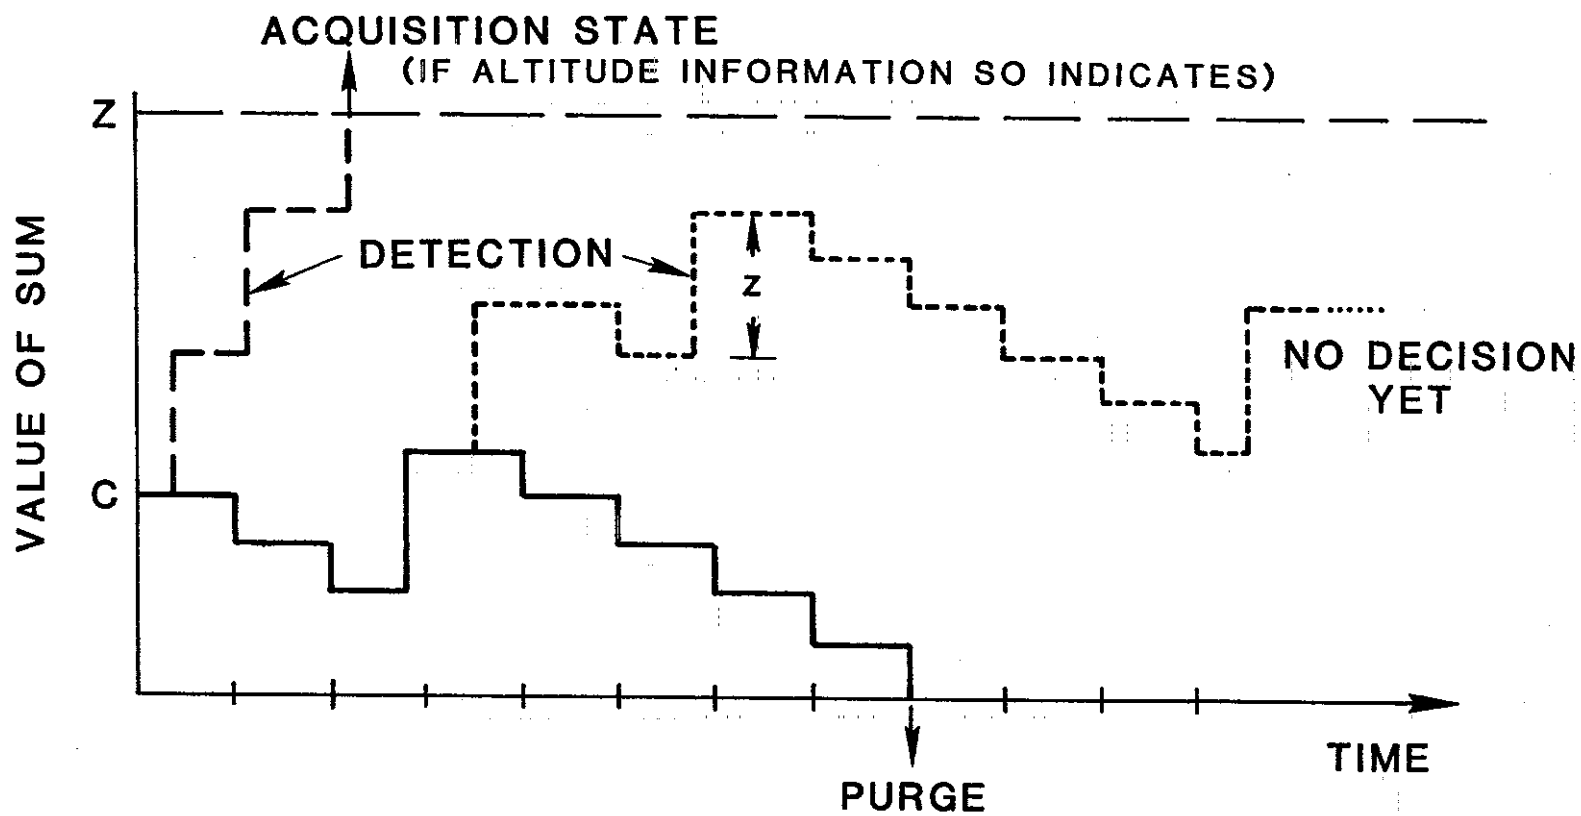


Fig. 4-7 Threat assessment for monitor state processing.

The average value of the sum n scans after its initialization will be $n[zr - Q] + C$, where r is the average number of times the address is detected per scan and Q is the probability that no address is detected during a scan. From this it is apparent that if zr exceeds Q the sum will tend, in time, to exceed Z . If zr is less than Q , the sum will tend, in time, to fall below zero. Clearly, z must be chosen so that the first situation prevails for all detection rates r that can be associated with threatening aircraft.

If the values of r and Q were uniquely and monotonically related to the range of the aircraft, the choice of z would be straightforward. Specifically, the minimum range at which an aircraft could not pose an immediate collision threat would be determined and z would be set equal to Q_m/r_m where r_m and Q_m are the values of r and Q for aircraft at that minimum range. (A somewhat larger value of z would actually be required since the time the algorithm requires to reach a decision tends to infinity as Q/r approaches z .)

Unfortunately, the substantial variations that can occur in transponder power outputs and link losses keep r from being uniquely related to the aircraft range. Thus z must be made large enough to ensure that no threatening aircraft will go uninterrogated. This means that a number of aircraft will be interrogated whose range is so large that they need not have been interrogated. These interrogations cannot be avoided when an aircraft is first detected; for there is then no way of knowing if the detection is the result of an unusually large power from a distant transponder.

On the other hand, once an aircraft has been interrogated, a more discriminating decision can be made concerning it even if a reply is not received, for the absence of that reply indicates that the reliability of the interrogation and/or reply link is not as good as had been thought and no further interrogations should be made until the reliability improves. Since that improvement can be sensed only by a change in the detection rate of the aircraft's squitters and fruit, a higher detection rate should be required for any subsequent interrogations. Thus the parameter z in this processing algorithm should not be a constant but should vary from aircraft to aircraft according to the number of times they have previously been interrogated unsuccessfully.

Minimum Triggering Level.— Another important system parameter is the Minimum Triggering Level (MTL) used to detect squitters and fruit. Setting the MTL to about the minimum received power expected from any threatening aircraft will both facilitate the rapid interrogation of threatening aircraft and reduce the number of interrogations to non-threat aircraft. It is the value of MTL that should be used if only one fixed value is to be employed. However, the value of the MTL cannot be fixed but must instead be adjusted continuously to the most sensitive value that satisfies the interference limiting standard. In this way the collision protection provided is always maximized subject to the constraints imposed. Whether or not the protection is adequate is determined by whether or not the resulting MTL is more sensitive than the minimum value determined above.

Threat Assessment.— A final question to be addressed is the relationship between the assessment of an aircraft's threat potential from altitude information and from the running sum associated with it. For example, should a sum be associated with an aircraft whose altitude separation from the TCAS aircraft is known to be quite large and, if so, what action should be taken when the sum reaches Z?

Part of the answer to this question is clear. Since any altitude information obtained from an aircraft is less ambiguous and more precise than that obtained from a running sum, an aircraft that is determined not to be an immediate threat from the available altitude information should not be interrogated. One might infer from this that the sum need not be initialized for an aircraft until the available altitude information indicates that it may be a threat. However, this would delay the interrogation of aircraft closing in altitude by the time required for the sum to build up to the value Z.

This delay could conceivably compromise the collision protection provided against aircraft with marginal transponder power. Therefore, the sum is associated with an aircraft when the monitor state is first assigned to it and the sum is allowed to evolve independently of the altitude information until the threshold Z is reached. The state will be changed to the acquisition state at that time unless the available altitude information indicates that the aircraft is not a threat. If it is not a possible threat, the evaluation of the running sum continues, but the sum is not permitted to exceed z.

The processing sequence that results from the above decisions is shown in Fig. 4-8 for a single aircraft ID.

To complete the functional description of the processing for aircraft assigned to the monitor state it is necessary to specify:

1. The values of the running sum parameters C, Z and z,
2. The processing of the altitude information.

These tasks will now be addressed in turn.

4.3.2 Parameters of the Running Sum Algorithm

A number of important factors influence the choice of C, Z and z. These factors are discussed below.

1. First, Z should be made large enough, or z made small enough, so that several aircraft will be detected before the running sum reaches Z. Otherwise, many aircraft that are separated in altitude will be interrogated when the sum reaches Z even though no altitude information has been received. Since it is a squitter that causes the sum to be initialized, it seems reasonable to require at least two more detections after initialization before Z can be reached. This will occur if z is less than Z-C. The probability that some

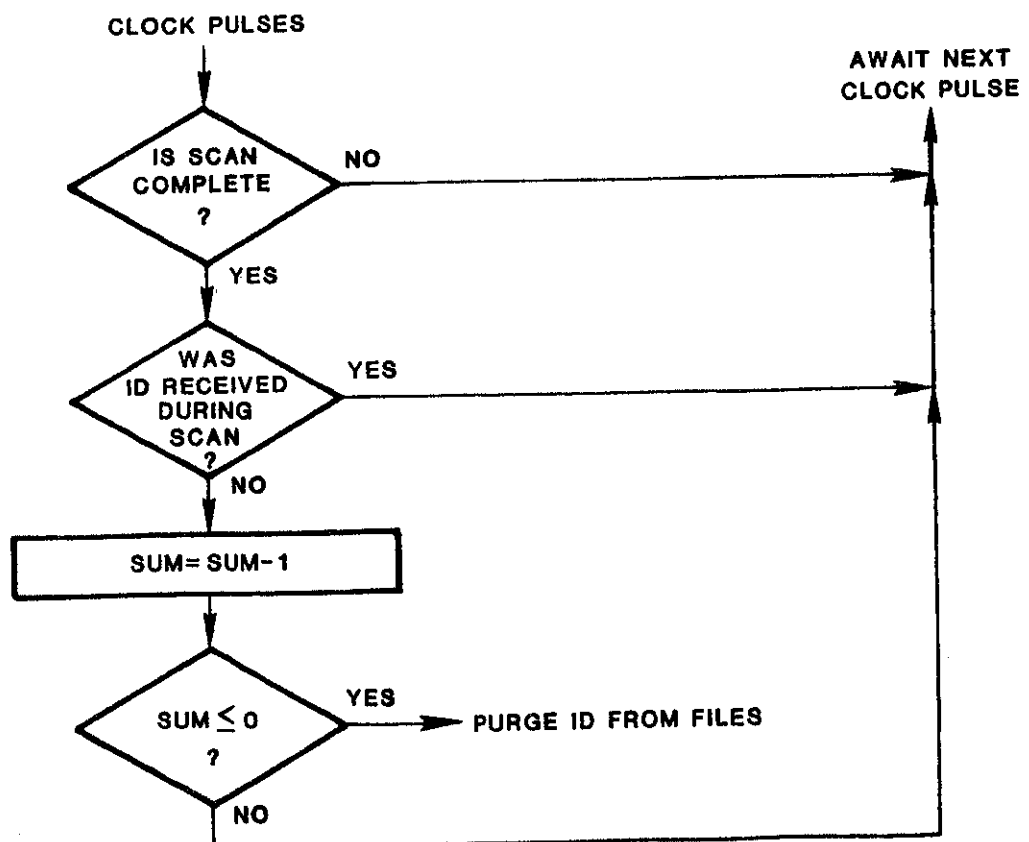
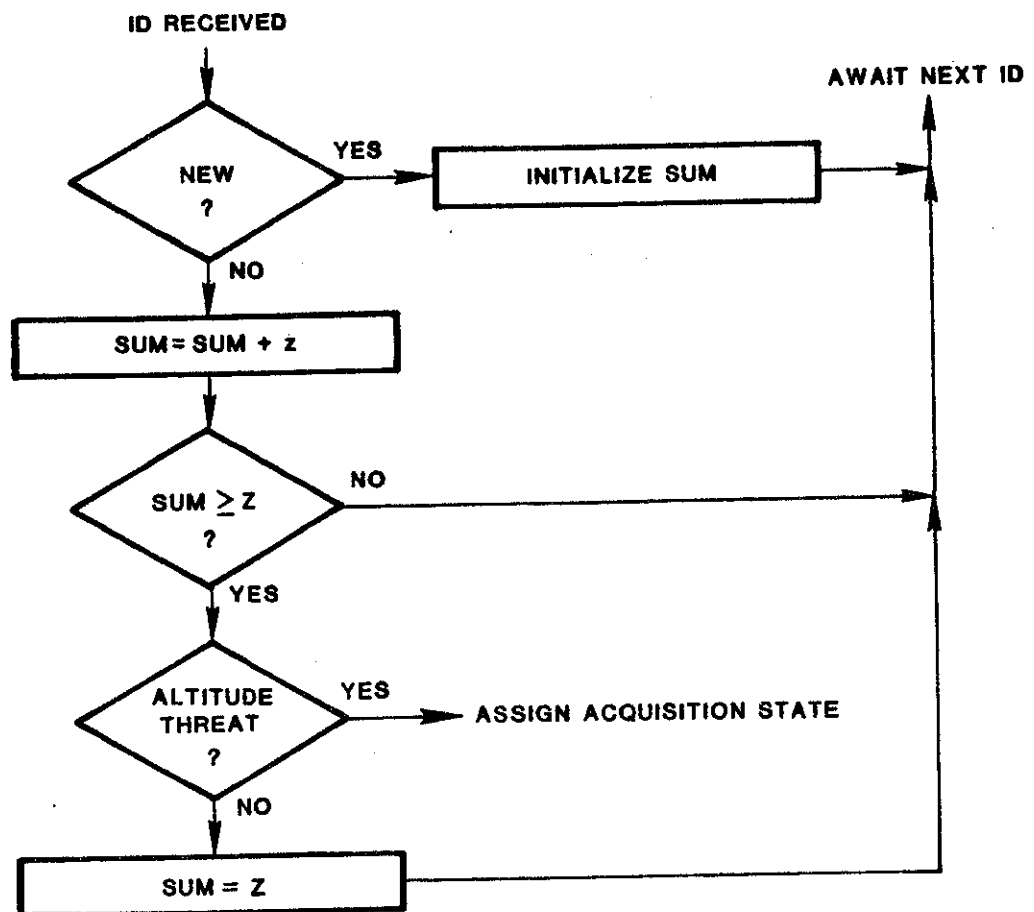


Fig.4-8. Monitor state processing algorithms.

altitude information will be received before the sum reaches Z is then at least 0.5 if the detection probability is the same for squitters and fruit. If z is much less than $Z-C$, the time that elapses before a threatening target is interrogated may become excessive. Therefore z should be on the order of $Z-C$.

2. A second consideration is that threatening aircraft should be assigned to the acquisition state in a timely manner even when the detection probability is varying widely, as it will during deep fades. This is particularly important before the first acquisition attempt. Clearly, the performance cannot be acceptable in all situations, but it seems reasonable to require that the acquisition state be assigned whenever several detections occur in a short period of time, even if the value of the running sum is near zero. This implies that Z be at most a few times z when there is no past history of interrogation failures, that is, let Z/z be at most three.

3. As an aircraft accumulates a history of unsuccessful interrogations, the value of z should be reduced as discussed in Section 4.3.1. These values need not be limited as described in the previous paragraph, since the lengthening history of no replies reduces the probability that a short deep fade is in progress. However, with one exception, z should always be large enough that the threshold Z will be reached in a relatively short time if squitters or fruit begin to be detected on each successive scan. That time is taken to be 10 seconds and therefore Z/z should not exceed 10.

4. An exception arises if repeated interrogations of the aircraft fail to elicit a reply and yet the aircraft continues to be reassigned to the Acquisition state even after z has been reduced to the minimum value specified above. Then it is highly probable that the Mode S transponder being interrogated is not working properly, e.g., it is abnormally insensitive or its power is abnormally high, and z should be reduced even further to avoid wasting interrogations. Indeed, one might argue that no further interrogations should be addressed to it; however a more conservative approach is to only relax the constraints on Z/z by a factor of two, from 10 to 20, and this only in the extreme situation in which the aircraft has been returned to the monitor state from the acquisition state three or more times.

5. At the other extreme, in the absence of altitude information, an aircraft should be assigned to the acquisition state if the probability, P , of detecting its squitters and fruit is sufficiently large, no matter what the past history of interrogations has been. It is obvious that this should be done when the detection probability, P , is one; for there is then no way of estimating just how close the aircraft may be. That assignment should probably also be made when P is as small as $1/4$ or $1/8$ since the antenna switching on the two aircraft could cause three out of four transmissions to occur on an antenna pair for which the path loss is high. The conservative value of $1/8$ was chosen. However, even if P exceeds $1/8$, there is no certainty that the assignment to the acquisition state will always be made; all that can be specified is the probability of its being made. The parameters were selected so that the transition from the monitor state to the acquisition state will be

made with a probability of at least 90% whenever P exceeds $1/8$. Thus, if the transition is not made on the first attempt and the aircraft's transmissions continue to be received, the entire process will be repeated and the probability that it is assigned the acquisition state on one of the first two iterations will be 99%.

6. Finally, an aircraft should not be purged from the system while there is any significant chance that it soon will be reassigned to the monitor state; for if that occurs the history of past interrogations will be lost. On the other hand, to reduce the memory load, an aircraft's ID should be purged when there is little chance of receiving further transmissions from it. A requirement was imposed that aircraft for which P is less than $1/50$ be purged from the system with a probability of 90%.

These factors lead to the set of constraints listed in Table 4.3. These constraints can be translated into numerical limits by drawing upon the performance expressions for Sequential Probability Ratio Tests. The expressions involve

P_1 , the value of P above which it is desired that the acquisition state be assigned,

β , the probability that this assignment is in fact not made,

P_0 , the value of P below which it is desired that the address be purged from the system and α , the probability that this is not done.

These values are $P_1 = 1/8$, $P_0 = 1/50$ and $\alpha = \beta = 0.1$.

Approximate expressions for C , Z and z in terms of P_1 , P_0 , α and β are available in Ref. 8 for the situation in which at most one squitter, or fruit, is received from an aircraft per scan. This situation will arise when the ground interrogation rate of Mode S transponders is small. It is a "worst case" situation for the issues of concern here. The expressions are

$$C = \frac{-\ln \beta}{\ln \frac{(1-P_0)}{(1-P_1)}}$$

$$Z = \frac{-\ln (\alpha\beta)}{\ln \frac{(1-P_0)}{(1-P_1)}}$$

$$z = \frac{\ln (P_1/P_0)}{\ln \frac{(1-P_0)}{(1-P_1)}}$$

TABLE 4-3

CONSTRAINTS ON THE VALUES OF C, Z AND z.

$z \approx Z-C$

$Z/z < 3$ Prior to first interrogation

$Z/z < 10$ After first interrogations but,

$Z/z < 20$ If it appears transponder is malfunctioning

For $P > 1/8$ Acquisition state assigned with probability of at least 90%

For $P < .02$ Aircraft ID purged with probability of at least 90%

Introducing the values of P_1 , P_0 , α and β into these expressions yields $C=20.3$, $Z=40.6$, and $z=16.2$, which are rounded off to

$C=20$
 $Z=40$
 $z=16$

Note that the theory also states that the mean time required to reach a decision when $P = P_1$ is approximately $C/(1+P_1z)$, or about 7 seconds.

These values satisfy the constraints that apply before the first interrogation has been made. Thus they may also be used for that situation. The remaining issue is how to reduce z on subsequent returns to the monitor state from the acquisition state. Table 4-3 implies that z should be no less than four for the first and second return and no less than two for any subsequent return. This suggests that the sequence of values for z be 16,x,4,2 where x is a value to be determined.

Simulation studies of the kind discussed in Section 4.7 have indicated that the performance of the system is not very sensitive to the choice of x . Thus it is appropriate to continue the geometric pattern and take it to be 8.

4.3.3 Altitude Processing

For a target aircraft that may possibly be a threat in range, the function of monitor state altitude processing is to determine whether available altitude information indicates that the aircraft is not a threat in altitude. The processing divides naturally into two parts. In one, estimates of the relative altitude rate are derived from the sequence of fruit replies received from an aircraft that has been assigned to the monitor state. In the other, the threat that an aircraft represents in altitude is evaluated whenever the value of the sum described in Section 4.3.2 becomes at least as large as the threshold Z . These two aspects of the processing are discussed in more detail below.

The information available is the sequence of altitude reports contained within the fruit that have been received from the aircraft. However, only a few of the most recent values are significant. Because of that (and to reduce the storage requirements), the threat assessment is based upon the most recently received altitude and the most recently calculated estimate of the altitude rate. The two primary design questions are then: How should the altitude rate be estimated and how should the threat be assessed?

Rate Estimation.—Rate estimation involves a compromise. An up-to-date estimate of the rate is desired, which implies that the two most recently received values of the altitude should be used in the estimate. On the other hand, the values used must be separated by enough time to ensure that the estimate is not corrupted by the quantization of the altitude reports.

Finally, the time separation should be small enough to ensure that the true altitude rate is being measured. The compromise may be struck in a number of

ways, the approach used in the simulation described in Section 4.7 is as follows:

For each aircraft assigned to the monitor state, an altitude, an altitude rate, and the time at which they apply is retained in a file. Initially, the first altitude report received from the aircraft is stored in the file. Each subsequent altitude report replaces the one that is stored unless the time between the two reports is less than 20 seconds, in which case the newly received report is discarded. When a new altitude is to be stored, it and the altitude it is to replace are used to re-estimate the altitude rate. The new rate then replaces the previously stored rate unless the time separation of the two reports exceeds 120 seconds, in which case no rate estimate is retained.

The above procedure does not always cause the most recently received altitude to be saved in the file. If the most recent altitude were saved, and no other altitudes were recorded, the elapsed time between the stored and newly received altitudes could at times be so small that a useful estimate of the altitude rate could not be obtained. Of course, this could be remedied by retaining additional altitude information in the file, but the approach described here provides satisfactory performance. With this approach, the stored altitude and the altitude rate were valid less than 20 seconds ago unless an altitude report has not been received for 20 seconds in which case they were valid less than 20 seconds before the last received report.

Having chosen the means by which an aircraft's altitude and altitude rate are determined, it remains to specify the means by which the aircraft threat is assessed.

Altitude Threat Assessment.— An aircraft should not be considered a threat if the altitude separation from it is large and will continue to be so for the immediate future. Stated more precisely, an aircraft should be considered a threat and it should be assigned to the acquisition state if either 1) no altitude information is available, or 2) the altitude separation is, or has recently been, less than some critical separation, or 3) the separation could become zero within some critical time. In the simulations reported in Section 4.7 a critical separation of 3,000 feet and a critical time of 60 seconds were used.

In particular, when the sum associated with an aircraft becomes as large as the threshold Z it is assigned to the acquisition state unless the following conditions are satisfied:

1. An altitude has been received from it and
2. When the altitude was stored the vertical separation exceeded 3,000 feet and either
- 3a. The altitude rate was estimated within the last 60 seconds and at that rate the vertical separation of the aircraft could not become zero for at least 60 more seconds or

- 3b. An altitude rate was not estimated within the last 60 seconds but, assuming that the aircraft has been closing in altitude since the last altitude was stored and that the closure rate does not exceed the sum of 6,000 feet per minute plus the magnitude of the rate for the TCAS aircraft, the present vertical separation either exceeds 9,000 feet or the additional time required for it to reach zero exceeds 60 seconds.

The last condition, 3b, pertains to situations in which a recent estimate of the rate is not available, and should not arise often since the parameters of the monitor state processing have been chosen so that several altitude reports will usually be received before the threshold Z is reached. Moreover, it affects the performance of the processor significantly only when there are many aircraft for whom the vertical separation from the TCAS aircraft is rather large but from whom few altitude reports are received.

Condition 3b can occur when the fruit rate is low either because few surveillance replies are requested by other interrogators or because the link geometry is such that they do not reach the TCAS aircraft. Since the former situation can be encountered on oceanic flights and the latter can be encountered in overflights of high density terminal areas, the condition has been retained in the design and is included in the simulations reported in Section 4.7.

4.4 Acquisition State

4.4.1 Functions

The processing of aircraft in the acquisition state is similar to that used in BCAS (Ref. 7) and need be described only in broad outline and in contrast to the BCAS processing.

The functions of the processing are to determine the range of aircraft and to assess the threat they represent. If that threat is significant, the aircraft is assigned to the track state. Otherwise, the aircraft is assigned to the dormant state or the monitor state. In making these assignments it is necessary to limit the number of interrogations to aircraft from whom replies are not received. This limit must balance the goals of ensuring that all threatening aircraft are assigned to the track state and of avoiding unnecessary interrogations that could cause the interrogation limiting algorithm to compromise the collision protection provided by the system. The means of achieving these goals are discussed in turn below.

4.4.2 Threat Assessment

In the acquisition state the threat represented by an aircraft is determined from its altitude separation from the TCAS-equipped aircraft and its slant range.

Altitude Separation.- The altitude information is used in much the same way as it is for the monitor state. The two processes differ only in that altitude information will be obtained from replies to interrogations, rather than from fruit. Thus, except for some minor changes, the altitude processing for the acquisition state is as described in Section 4.3.3. In particular, an aircraft is removed from the acquisition state and reassigned to the monitor state whenever the altitude information indicates that it cannot become an immediate threat.

It may be noted that a transition from the acquisition state to the monitor state is not allowed in Fig. 4-2. It was omitted from the figure and the accompanying text to simplify the initial description. A more complete description which distinguishes between the use of altitude and range information is shown in Fig. 4-9.

Slant Range and Time-to-Endanger.- The range information is used to determine the length of time during which a collision cannot occur when there is no vertical separation between the two aircraft. This time is called the time-to-endanger and is denoted by TE. The available information upon which the calculation of TE is based is the range, the maximum capable airspeeds of the two aircraft and the knowledge that a 250-Kt speed limit exists at altitudes below 10,000 feet. Because this speed limit is sometimes waived, it is assumed that the interrogated aircraft does not obey it. It is assumed that the TCAS aircraft does conform to the speed limit. Thus, above 10,000 feet, TE is the range divided by the sum of the maximum capable airspeeds, and below 10,000 feet, it is the range divided by the sum of the speed limit and maximum capable speed of the interrogated aircraft. A conservative speed limit of 300 knots is used in the system simulation discussed in Section 4.7.

The magnitude of the threat represented by an aircraft is inversely related to TE. The question is: what is the value of TE for which an aircraft should be assigned to the track state? The value must be large enough to ensure that the track state is assigned before the aircraft reaches the threat boundary used by the CAS logic. For 1200-Kt and 500-Kt head-on encounters this boundary is reached when TE equals 33 and 27 seconds, respectively.

Since the two times given above are comparable and since some additional time is required to establish a track that can be used by the CAS logic, there is little advantage in letting the threshold value of TE depend upon altitude. Instead, a single threshold value of 41 seconds was used in the simulations described in Section 4.7. In the absence of the interrogation limiting constraints, the use of a larger threshold would provide added collision protection by causing aircraft to be tracked at greater ranges. However, action of the interrogation limiting algorithm could in fact reduce the range at which aircraft are detected if this threshold were made larger.

4.4.3 Interrogation Parameters

The above discussion assumes that the interrogations made by TCAS elicit replies. It remains to discuss the selection of the rate and the power of interrogations addressed to an aircraft.

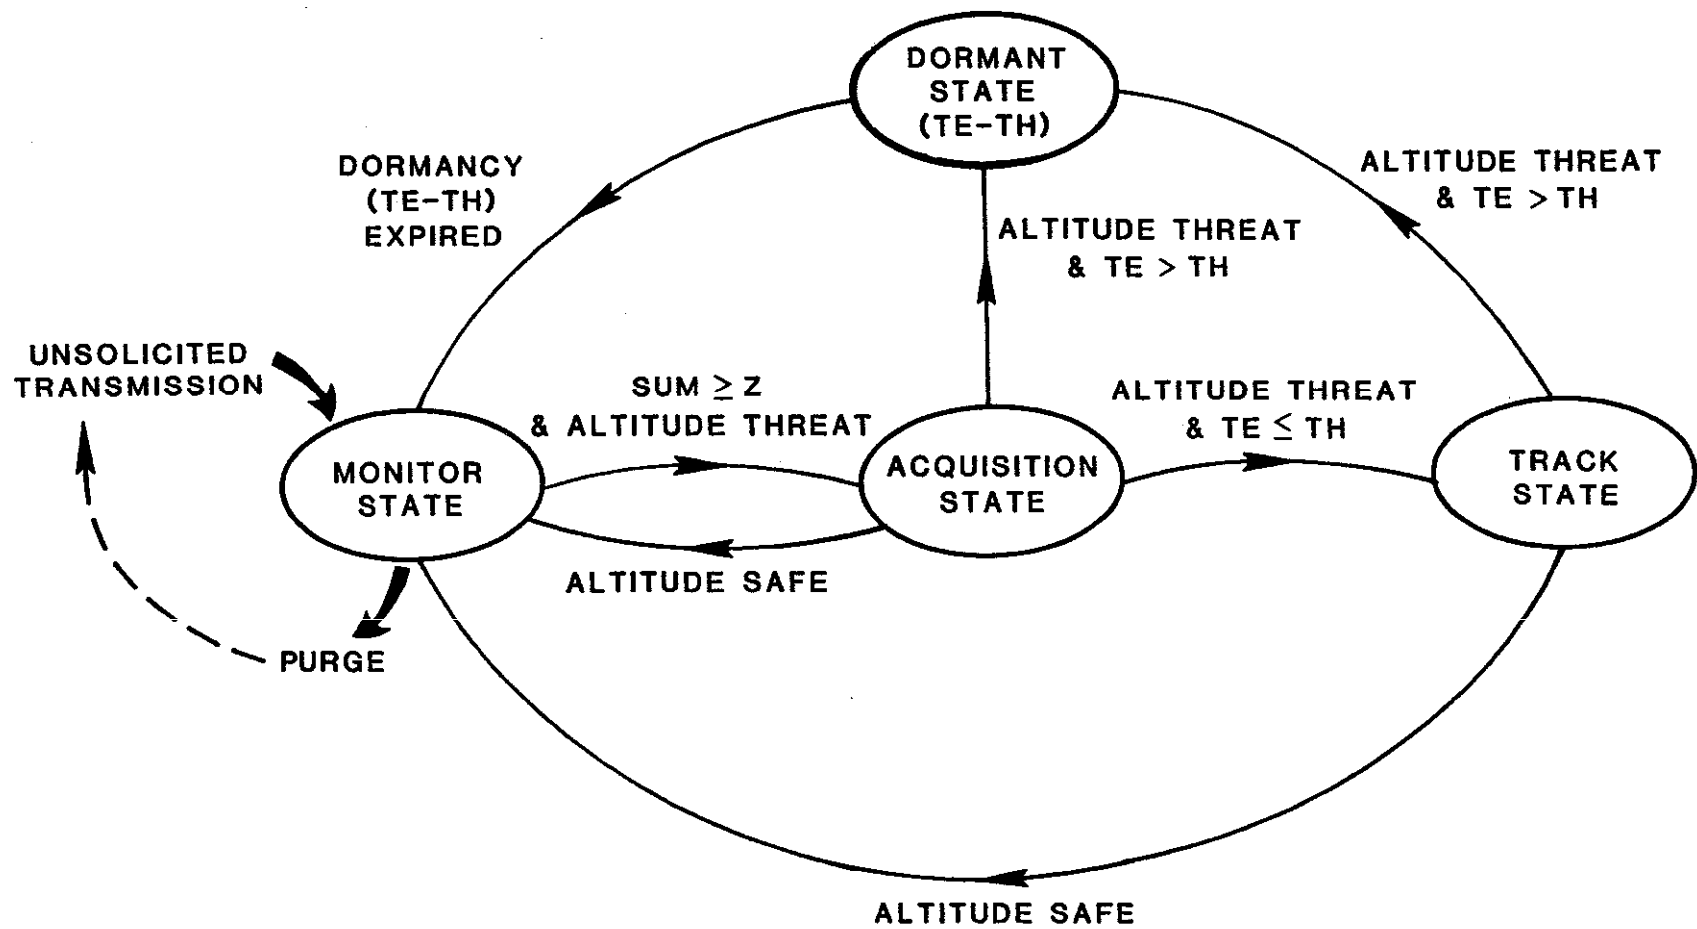


Fig.4-9. Operation of the Mode S surveillance processor.

Parameter Selection Considerations.- Several factors enter into the choice of these parameters. The interrogation rate and power should be sufficient to ensure an adequate reply probability. However, neither the rate nor the power should be excessive, for the resulting action of the interrogation limiting algorithm may then compromise the collision protection that is provided.

It is not necessary for replies to be received from all aircraft assigned to the acquisition state. Some assignments may have resulted from unusually high power emitted by aircraft whose ranges are so great that they cannot be threats. Failure to elicit a reply from such an aircraft will cause the increment z to be changed when the monitor state is reassigned to the aircraft. This, in turn, will prevent its being returned to the acquisition state and reinterrogated until its level of squitter/fruit activity has increased.

The design choices to be made are then: what power level should be used for interrogations and how many interrogations should be made before the acquisition attempt is declared a failure?

Power Level.- The choice of the power reflects the fact that the interrogation link is nominally as reliable as the link for the reception of squitters and replies to interrogations. Thus if the presence of an aircraft were detected with an MTL 6 dB above the minimum value, an interrogation power 6 dB below the maximum value should suffice to elicit a reply.

The balance between the two links is not exact; in any specific situation a substantial imbalance may exist. The only consistent rationale for less interrogation power is that Mode C fruit is not present on the interrogation link. But this does not justify a general reduction in the interrogation power since the proposed power might be inadequate to elicit a reply from a threatening aircraft whose transponder sensitivity is low relative to its power output. Such an aircraft would seem to benefit from an increased interrogation power. However, simulation studies of the type discussed in Section 4.7 indicate that such an increase is not needed to obtain satisfactory performance. Moreover, to adopt an increase in the interests of conservatism could be ill-advised since the constraints imposed by the interrogation limiting algorithm would then be tightened. Hence the decision to match the interrogation power to the MTL.

Interrogation Rate.- The following factors were considered in choosing the maximum number of interrogations allowed during an acquisition attempt. The number must be large enough to ensure that a threatening aircraft is acquired in time for evasive maneuvers to be taken. On the other hand, the number should be small enough to prevent unnecessary restriction of the collision protection by the action of the interrogation limiting algorithm. The choice between the extremes is not critical since the maximum number of interrogations will rarely be employed.

Some guidance in making the choice is provided by the conditions under which the acquisition state is assigned by the monitor state processing. Examination of that processing shows that the acquisition state is assigned n scans after the sum has been initialized only if the number, n_r , of squitters and fruit received is approximately equal to $(20 + B)/z$ where B is the number of scans during which there were no receptions. If only squitters were received, B would equal $n - n_r$ and, for the state change to occur in n scans, n_r would be approximately equal to $(20 + n)/(z + 1)$. Then one could conclude that the reliability of the reply link was $(20 + n)/(z + 1)n$.

As noted in Section 4.3.2, the mean value of n is about 7 for the situations in which it is desired to assign the acquisition state with z equal to 16; thus the link reliability is on the order of $1/4$ when the acquisition state is first assigned. Consequently, if the interrogation and reply links are balanced, an average of about 4 interrogations should be needed to elicit a reply. If fruit are also received, the link reliability will be less than this estimate and more interrogations may be needed. Conversely, the antenna switching for acquisition is not random, as it is during monitor processing, but is determined by the history of successful receptions. Thus, fewer interrogations than four might suffice.

Faced with these uncertainties, and the knowledge that there is little penalty in erring on the high side, it was decided to allow a maximum of 6 interrogations during an acquisition attempt after one or two previous attempts have failed. A larger value, 9, is allowed for the first attempt to reduce the chance of failing to acquire a truly threatening aircraft with a substandard transponder. At the other extreme, after three previous failures, each accompanied by a decrease in z , it is assumed that there is little chance a reply will ever be received. This would suggest that the aircraft not be interrogated further, but conservatism indicates one interrogation on each acquisition attempt after the third.

Simulation studies of the kind discussed in Section 4.7 were used to explore the change in system performance that would result from small variations of the numbers of interrogations presented above. Little change was observed so the choices were adopted.

4.5 Track State

An aircraft that has been assigned to the track state is interrogated regularly and tracked. These operations differ from those used in BCAS in only two regards. First, the interrogation power is varied according to the aircraft range and, second, the altitude processing has been modified to incorporate the improvements introduced in the processing of other states.

The decision to vary the interrogation power with aircraft range stemmed from two factors. One was that there is no reason to use the maximum possible power to interrogate aircraft in the track state when a lower power sufficed to obtain a reply in the acquisition state. The other was that the power used for acquisition interrogations is as large as allowed by the interrogation limiting algorithm. If that power provided a detection range of 20 nmi.,

there is little point in using it to interrogate a target at a range of 2 nmi. If excess power is used to track a close-in aircraft, the range at which other aircraft are acquired will be reduced by the interrogation limiting algorithm, thereby reducing the overall collision protection provided by the system.

The manner in which the interrogation power should be varied with range is not immediately clear. In the absence of any channel fading, a reasonable procedure would be to vary it as $(R/R_0)^2$ where R is the range of the aircraft to be interrogated and R_0 is the surveillance range for maximum power (30 nmi). That is, the power used for a range R should be reduced by $20 \log(R_0/R)$ dB from the power used at the maximum range R_0 . Since link fades due to multipath and aspect angles occur frequently, this manner of varying the interrogation power is not acceptable, but it becomes much more promising when an adequate fade margin is included in it.

Examination of link propagation data indicated that a margin of about 10 dB was more than adequate. Thus the interrogation power to an aircraft at a range R might reasonably be taken to be $10 + 20 \log(30/R)$ dB below the maximum possible power. This power might still exceed that used to (successfully) acquire the aircraft, so we limit the interrogation power to the lesser of the above expression and the power used for acquisition.

The resulting variation with range is shown in Fig. 4-10 and is summarized in the statement: the interrogation power used for tracking is the maximum power for ranges greater than 10 nmi and decreases as the square of the range for ranges of less than 10 nmi: however it never exceeds the interrogation power used for acquisition.

For the issues of interest here, the altitude processing in the Track state is identical to that used for the acquisition state. Thus the monitor state is assigned to an aircraft under the same conditions as it would be if the processing were occurring in the acquisition state. These properties were summarized in Fig. 4-9. Detailed descriptions of the system implementation are given in (Ref. 7).

4.6 Dormant State

This state is assigned when the reply to an interrogation indicates that the target cannot be a threat in range for a time that exceeds the threshold TH . In such situations the aircraft should not be interrogated further until a time $TE-TH$ has elapsed. It is for this time that the aircraft is assigned to the dormant state.

At the end of the interval $TE-TH$ the aircraft may possibly become a threat again so its activity must then be monitored as is that of other aircraft. That is, it must be assigned to the monitor state or purged from the system. Somewhat better performance is obtained by assigning it to the monitor state. This is particularly true if the file on the aircraft's altitude and altitude rate is updated during the time it is assigned to the dormant state and that information is retained when it is assigned to the

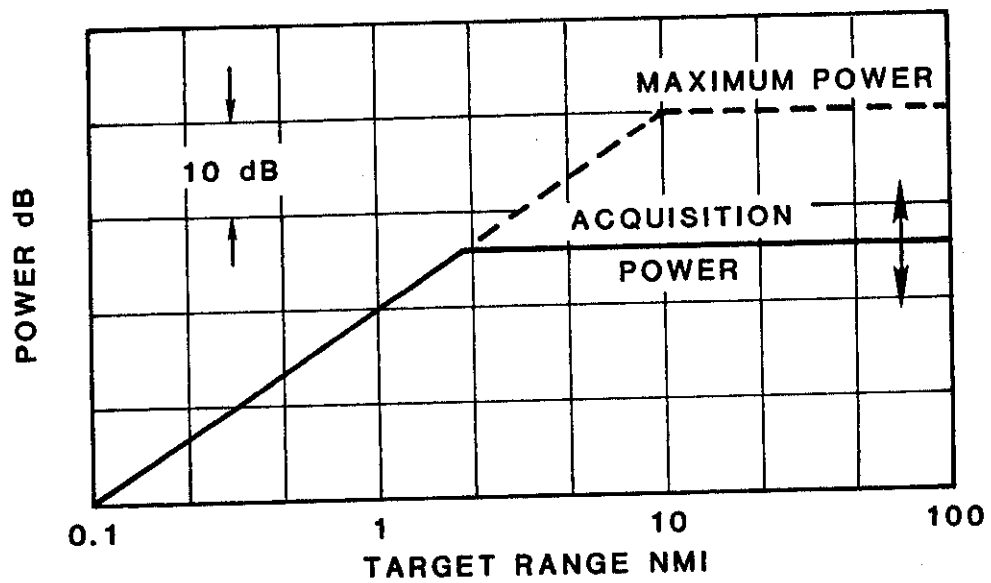


Fig. 4-10. Track state interrogation power.

monitor state. This approach will result in a larger number of aircraft being assigned to the monitor state than would be if it were purged when the interval TE-TH has elapsed. However, after assignment to the monitor state from the dormant state, most aircraft are soon purged from the system in any case. Thus assigning the monitor state to them does not stress the storage or processing capabilities of the system.

4.7 Performance Evaluation

The performance of the surveillance processor is indirectly coupled to the operational environment through the interrogation limiting algorithm. That coupling manifests itself through the value that is assigned to the MTL used for the detection of squitters and Mode S fruit. Thus the performance of the processor can be evaluated by first determining the MTL values for which satisfactory collision protection is provided and then determining the value that the MTL will assume in the operational environments of interest.

The results of the first step of that process are discussed here. The conclusion is that satisfactory protection is provided when the MTL is raised as high as 6 dB above the nominal value of -74 dBm. As discussed in Section 5.5, in the intended operational environments the MTL will not be raised by more than 6 dB at low altitudes or 3 dB at high altitudes. Thus the system can provide the desired collision protection in the intended operating environments.

4.7.1 Performance Goals

TCAS II is intended to provide collision protection in several different operational environments. Here the extremes represented by the low-altitude high-density environment and the high-altitude, low-density environment will be used to measure the acceptability of the design described in Sections 4.1 through 4.6. The transition from low to high altitude occurs at an altitude of 10,000 feet.

Below 10,000 Feet.- At altitudes below 10,000 feet TCAS II is intended to provide collision protection from aircraft on head-on collision courses at relative airspeeds of 500 kts. In such encounters the "Threat Boundary" used in planning evasive maneuvers is crossed 27 seconds before collision. It is mandatory that the aircraft be assigned to the track state before that boundary is crossed. To allow some time for the planning of evasive maneuvers it is desired that, with a 90% probability, it be assigned at least five seconds earlier.

The above goal should be met when the TCAS II is in an environment of transponder-equipped aircraft that are uniformly distributed in an area out to a range of 5 nmi with a density of 0.3 per nmi², and are uniformly distributed in range beyond 5 nmi. That is, the number, $N(R)$, of aircraft within a range R is given by:

$$N(R) = 0.3 R^2$$

for $R < 5$ nmi and by

$$N(R) = 7.5 (R/5)$$

for $R \geq 5$ nmi.

Above 10,000 Feet.— At altitudes above 10,000 feet TCAS II is to provide protection against head-on collisions at closing speeds of 1200 kts, but the peak density of aircraft is only 0.06 per nmi². At these speeds the threat boundary is crossed 33 seconds before collision. Again to allow some time for the planning of evasive maneuvers, it is desired that the aircraft be assigned to the Track state at least five seconds before the Threat Boundary is crossed with a probability exceeding 90%.

The density of aircraft in which this requirement must be met is uniform in area for ranges less than 10 nmi, and is uniform in range for larger ranges. That is, the number of aircraft, $N(R)$, within a range R of the TCAS is given by:

$$N(R) = R^2$$

for $R < 10$ and by

$$N(R) = 100 (R/10)$$

for $R \geq 10$ nmi.

Other Considerations.— Several other factors influence the system's performance. These include: the number of other TCAS units operating in the area, the fraction of the transponder-equipped aircraft that carry Mode S transponders, the distribution of altitude and airspeed for those aircraft and, finally, the number of aircraft that are generating Mode C fruit. All but the last factor influence only the value of the MTL used for the detection of squitters and fruit. Since the MTL is treated as a free parameter in this section, only the fruit level needs to be specified. A worst case assumption is made that no Mode S ground sensors are operating near the TCAS-equipped aircraft so the fruit environment is that associated with the given spatial distribution of aircraft when all of them carry ATCRBS transponders.

4.7.2 Models

Simulation models were combined with non-real time processing of flight test data to evaluate the system performance. Those evaluations involve 1) the probability that a received signal of a given power level will be detected in a given ATCRBS fruit environment, 2) the rate at which squitters and Mode S fruit are generated by a transponder and 3) the distribution of the power levels received by the TCAS and by the Mode S transponders it interrogates.

The model for the distribution of power levels was essentially that used in earlier studies of BCAS (Ref. 4). The exception was that the random

scan-to-scan variation in the TCAS power and sensitivity was eliminated so that the dynamic performance of a single TCAS unit was described rather than the static performance of an ensemble of such units. Squitters are generated at the rate of one per second, by design, but the generation rate for Mode S fruit depends upon the operational environment. This rate is conservatively estimated at one per second.

The expression for the detection probability was derived from the results available for an environment in which the ATCRBS transponders are uniformly distributed in area. That expression is (Ref. 7, p. 9).

$$P_D = P_0(P) P_f [P - 10 \log(\rho / 0.06)]$$

where P is the received power level, $P_0(P)$ is the detection probability in the absence of fruit, ρ is the (uniform) density of ATCRBS transponders and $P_f[\cdot]$ is a function that accounts for the effects of ATCRBS fruit. The above expression is for the situation in which error correcting decoding is not employed. The approximate effect of error correction decoding is to replace ρ by $\rho/2$, i.e., to reduce the fruit density by a factor of two.

The function $P_f[\cdot]$ has been determined by careful simulation for a uniform ATCRBS environment but not for the environment of interest here. However, a simple analysis suggests that in general $P_f[\cdot]$ is given approximately by the expression

$$P_f[y] = \exp - N(y/2)$$

where $N(y/2)$ is the average number of ATCRBS fruit that overlap a Mode S signal and that are received at a power level exceeding $y/2$.

For a uniform density of ATCRBS transponders the above approximation to $P_f[\cdot]$ agrees reasonably well with the result obtained by simulation (Ref. 7, p. 9). Therefore it was used for the non-uniform distributions specified in Section 4.7.1.

4.7.3 Results

The performance of TCAS in the head-on encounters described in Section 4.7.1 was evaluated by simulating the operation of the Mode S surveillance processor and driving that simulator either with an RF link simulator that generated the models described in Section 4.7.2 or with flight test data recorded by the Airborne Measurement Facility (AMF) (Ref. 7).

The RF link simulator was used as the driver during much of the TCAS development because it could be used to model a wide variety of situations. Since those models did not include a number of possibly important effects such as multipath, the available flight test data recorded on AMF tapes during the BCAS development was used to validate the overall performance of the system. In particular, for the collision encounters specified in Section 4.7.1, the probability that the aircraft would be assigned to the track state at least t seconds before the projected collision time was determined from both the flight test data and the link simulator.

The desired probability was obtained by configuring the simulator and driver for a head-on collision at the desired airspeed, altitude, and fruit environment. When the RF link simulator was used as a driver, this merely entailed setting the parameter values to the desired level. When the AMF tapes were used, the rate at which the recorded encounters were sampled was adjusted to scale the apparent relative speeds to the desired value, and some of the samples were corrupted to simulate the desired ATCRBS fruit environment. A series of encounters were then run and analyzed to determine the probability of interest. The results are discussed below.

4.7.3.1 Low Altitude Encounters

Figures 4-11 through 4-14 show the probability that an aircraft whose maximum capable airspeed is 300 kts will be assigned to the track state at least t seconds before collision when it is on a head-on collision course with the TCAS aircraft at an altitude of less than 10,000 feet with a relative airspeed of 500 kts. A larger maximum capable airspeed would cause the aircraft to be assigned the track state even sooner.

The results are for the situation in which the peak aircraft density is 0.3 per nmi² and both power programming and error correcting decoding are employed. As will be discussed subsequently they are also valid when the peak density is 0.15 and neither power programming nor error correction is used. In each figure the projected collision time is taken to be zero and the time at which the threat boundary is crossed is indicated by a vertical line.

Performance with the RF Link Simulator.- Figures 4-11 and 4-12 were obtained by running 300 encounters with the RF link simulator and plotting the fraction of the runs for which aircraft were assigned the track state at least t seconds before the projected collision. Thus for the encounters described by the rightmost curve in Fig. 4-11 all of the aircraft were assigned to the track state about 20 seconds before the threat boundary was crossed.

Figure 4-11 applies to normal operation of the surveillance processor with MTL's raised 6, 9, and 12 dB above nominal for the detection of Mode S squitters and fruit. For MTL's raised by 6 and 9 dB, 90% of the aircraft are assigned the track state about 20 seconds before the threat boundary is crossed. The performance differs little for these values because nearly all of the aircraft are assigned to the dormant state well before the threat boundary is crossed and are reassigned to the monitor state only when they are close to the threat boundary and the link reliability is even higher. Thus the performance for these MTL's is determined by the time required to assign an aircraft to the track state when the link reliability is high. In such situations one can cause the aircraft to be assigned to the track state T seconds earlier by merely increasing the threshold TH from its nominal value of 41 seconds to $41 + T$ Seconds. This can be done so long as the aircraft are still detected and assigned the dormant state well before the new threshold is crossed.




Fig. 4-11. Low altitude collision protection for normal operation with MTL's of 6, 9, and 12dB.




Fig. 4-12. Low altitude collision protection with the bottom antenna disabled for MTL's of 6, 9, and 12dB.




Fig. 4-13. Low altitude collision protection inferred from 7 flight tests conducted over land.




Fig. 4-14. Low altitude collision protection inferred from 6 flight tests conducted over water.

The situation changes when the MTL is increased by 12 dB. Then a significant number of aircraft are not detected until the time-to-endanger is less than TH and the performance curve is determined by the time at which the aircraft are first detected. Even then the performance of the surveillance processor is satisfactory in that 90% of the aircraft are assigned to the track state 15 seconds before the threat boundary is reached. To provide a scale of reference, it will be seen in Section 5.5 that the MTL increase does not exceed 6 dB in the environments for which TCAS is designed to operate.

As indicated earlier, the RF link simulator used to obtain the above results does not realistically model the effects of multipath and fades due to shadowing. Some measure of the magnitude of these effects can be gained by introducing a 20 dB fixed loss in the simulated bottom-mounted antenna. The probabilities that were obtained when the encounters described above were repeated with this loss inserted are shown in Fig. 4-12.

It is apparent from Fig. 4-12 that the loss of the bottom antenna has very little effect upon the track probability when the MTL is raised 6 dB. Essentially all of the aircraft are still assigned to the track state about 20 seconds before the threat boundary is reached. The effect of the added loss is more pronounced when the MTL is raised 9 or 12 dB, but even then, at least 90% of the aircraft are assigned to the track state before the threat boundary is crossed. However, for an MTL increase of 12 dB, small changes in the model for the system noises may cause significant changes in the time at which 90% of the aircraft are in track. That is, the performance will be much more robust when the MTL is raised 6 dB than when it is raised 12 dB.

Performance with AMF Data.— Further evidence that the Mode S surveillance processor will provide satisfactory collision protection for the head-on encounters under discussion was obtained by driving the simulated processor with AMF tapes of thirteen head-on encounters. The characteristics of the encounters are described in Table 4-4. As discussed above the relative airspeed and fruit environment were scaled to the values of interest here. In particular, the encounters were speeded up to a closing speed of 500 kts rather than the actual airspeeds of the aircraft listed in the table.

The six encounters flown over water exhibited substantially inferior performance compared to the flights that occurred over land at the same altitude. This was probably due to multipath interference, but other causes such as equipment failures cannot be ruled out. Because of the disparity in performance between the two kinds of flights, the track probability was determined for each set separately. Figures 4-13 and 4-14 show the results obtained from the over-land and over-water flights, respectively, for MTL increases of 6, 9, and 12 dB. Each figure also contains the curve from Fig. 4-11 for an MTL increase of 12 dB.

The performance obtained with the over-land AMF tapes is very similar to that obtained with the RF link simulator. This implies that the link reliabilities in the over-land flights were large enough that the Dormant

TABLE 4-4

CHARACTERISTICS OF RECORDED FLIGHT ENCOUNTERS.

AIRCRAFT		
TCAS	OTHER	SURFACE
B727	BONANZA	LAND
C580	C421	..
C580	C172	..
C421	BONANZA	..
C421	C172	..
C421	CHEROKEE	..
C421	CHEROKEE	..
B727	BONANZA	WATER
C580	C421	..
C580	C172	..
C421	BONANZA	..
C421	BONANZA	..
C421	CHEROKEE	..

Note: All encounters were head-on at about 5,000 feet MSL.

state was assigned well before the threat boundary was reached. In contrast, the performance obtained with the over-water flights differed markedly from the over-land performance at MTL increases of 9 and 12 dB, but was comparable at an MTL of 6 dB. Detailed examinations of the records indicated that the difference was caused by signal fades for which the link was unreliable at MTL increases of 9 and 12 dB, but for which it was still reliable at an MTL increase of 6 dB.

Conclusions.— Satisfactory collision protection against the stipulated head-on encounter is provided in all situations when the MTL is no more than 6 dB above nominal, and is not provided when the MTL is raised by 9 dB. If the over-water AMF tapes were not included in the analysis, an MTL increase of 12 dB might be acceptable but the protection would then be sensitive to the details of the link disturbances.

The above conclusions are based upon simulations in which the peak aircraft density was 0.3 per nmi² and both power programming and error correcting decoding were used in the surveillance processor. However, they are also valid for a situation in which the peak aircraft density is 0.15 per nmi² and neither power programming nor error correcting decoding is used. There are two reasons for this.

First, the parameters of the power programming were chosen so that they did not compromise the collision protection provided by the system when the MTL is fixed. Thus, the removal of power programming does not effect the results presented in Figs. 4-11 through 4-14. Second, the aircraft density influences the collision protection afforded at a given MTL setting only through the ATCRBS fruit associated with it. Thus changing the peak density from 0.3 to 0.15 will improve the performance by reducing the interference from such fruit. As discussed in Section 4.7.2 that improvement has been estimated to be equivalent to a factor-of-two increase in the argument of the function $P_f[\cdot]$. On the other hand, the elimination of error correcting decoding has been estimated to be equivalent to a factor-of-two decrease in the argument of $P_f[\cdot]$. Thus the two factors cancel and the link simulator parameters remain unchanged.

4.7.3.2 High Altitude Encounters

The collision protection provided at altitudes above 10,000 feet was determined in much the same way as it was for lower altitudes.

Performance with the RF Link Simulator.— Figures 4-15 and 4-16 give the probability that an aircraft will be assigned to the track state at least t seconds before collision when it is on a head-on collision course with the TCAS aircraft at an altitude of more than 10,000 feet with a relative airspeed of 1200 kts and both aircraft have a maximum capable airspeed of 600 kts.




Fig. 4-15. High altitude collision protection for normal operation with MTL's of 0, 3, and 6 dB.




Fig. 4-16. High altitude protection with the bottom antenna disabled for MTL's of 0,3, and 6dB.

The results were obtained with the RF link simulator and are for the situation in which the peak density is 0.06 per nmi² and both power programming and error correction decoding are employed. For the reasons given at the end of Section 4.7.3.1 they also apply to the situation in which the peak density is 0.03 and neither power programming nor error correction is used.

The interpretation of the figures is similar to that of Fig. 4-11 and requires little elaboration. It is clear from Fig. 4-15 that when both antennas are operating normally, the performance is satisfactory for MTL increases of 0 and 3 dB, but not for 6 dB. For purposes of comparison in the high-altitude environments for which TCAS is intended, the MTL will not exceed 3 dB.

A measure of the robustness of the above result is provided by Fig. 4-16 in which it is assumed that a fixed loss of 20 dB is inserted in one of the antennas. Even with this loss, at least 90% of the aircraft are assigned to the track state before the threat boundary is crossed when the MTL is raised by 0 or 3 dB.

Performance with AMF Data.-Only two of the encounters listed in Table 4-4 were started at large enough ranges to be useful in evaluating the performance of the surveillance processor against high-speed aircraft. For one of these (C421, Bonanza, land) the track state was assigned 40 seconds before collision for all three values of the MTL. For the other (B727, Bonanza, water) it was assigned 40 seconds before collision when the MTL was raised by 0 and 3 dB and 33 seconds before collision when it was raised by 6 dB. These times are consistent with those obtained with the RF link simulator.

5. INTERFERENCE LIMITING

Interference limiting is carried out by each TCAS II unit to keep interference effects to other systems at an acceptably low level. As described in Chapter 2 (Sec. 2.1 and 2.5) the interference limiting standards previously developed for BCAS had to be modified in the TCAS development program for several reasons: (1) to provide for directional interrogation, (2) to control self suppression, and (3) to control the fruit generated by TCAS.

Interference limiting standards have been developed in a form suitable for adoption as a National Standard. These standards, described in the next section, are inequalities that specify maximum values of interrogation power and interrogation rate. A given TCAS II unit conforms to these standards by means of interference limiting algorithms (Sec. 5.2), which are not standardized in detail. For example, a directional unit and an omnidirectional unit may employ different interference limiting algorithms, as long as the standards are satisfied.

5.1 Interference Limiting Standard

The interference limiting standards consist of three inequalities to be satisfied by each airborne interrogator. They are summarized in Fig. 5-1. The three inequalities correspond, respectively, to three interference phenomena: (1) air-to-air effects on transponder reply ratio, (2) suppression of the on-board transponder, and (3) generation of Mode C fruit.

These inequalities were originally derived analytically. Subsequently they have been tested through a comprehensive and detailed simulation study at the Electromagnetic Compatibility Analysis Center (ECAC) in Annapolis.

5.1.1 Derivations

The following derivations illustrate the nature of the issues involved.

5.1.1.1 Reply Ratio

A limit of 2% has been placed on the reduction in transponder reply ratio caused by TCAS II. This is a conservative basis for interference limiting since a drop in reply ratio of 2% would not significantly affect the reliability of tracking aircraft from a ground-based interrogator.

An initial question is how to allocate the 2% total into its two parts, (1) effects on transponders in other aircraft, and (2) effects on own transponder. The total could be divided into two fixed equal parts (1% each), or into two fixed unequal parts, or into variable parts at the discretion of each manufacturer.

A variable allocation would be undesirable since it could result in the following situation. Imagine two populations of TCAS II interrogators, type A in which 1.9% of the 2% drop in reply ratio is allocated to suppression of

$$(1) \quad \sum_{i=1}^I \frac{P(i)}{250} < \text{the smaller of } \left[\frac{280}{NT + 1}, 18 \right]$$

$$(2) \quad \sum_{i=1}^I M(i) < 0.01 \text{ sec.}$$

$$(3) \quad \frac{1}{B} \sum_{k=1}^K \frac{PA(k)}{250} < \text{the smaller of } \left[\frac{80}{NT + 1}, 5 \right]$$

where the variables in these inequalities are defined as follows:

I = total number of interrogations transmitted by own TCAS II in a 1-second period.

i = index number for all interrogations; $i = 1, 2, \dots, I$.

$P(i)$ = total radiated power (in watts) from the antenna for the i^{th} interrogation.

NT = onboard estimate of TCAS II interrogators within 30 nmi, obtained by counting TCAS Broadcast Interrogations, detected with a transponder receiver threshold of -74 dBm.

B = beam sharpening factor (ratio of 3-dB beamwidth to beamwidth resulting from interrogation sidelobe suppression).

$M(i)$ = duration of the self suppression (or "mutual suppression") interval for own transponder associated with the i^{th} interrogation.

K = total number of Mode C interrogations transmitted by own TCAS II in a 1-second period.

k = index number for Mode C interrogations; $k = 1, 2, \dots, K$.

$PA(k)$ = total radiated power (in watts) from the antenna for the k^{th} Mode C interrogation.

Fig. 5-1. Interference limiting standard.

own transponder, and type B in which 1.9% of the 2% drop in reply ratio is allocated to air-to-air effects. It follows that the transponders on all of the type A aircraft would experience a total degradation considerably more than 2%. Such conditions are avoided if the allocation associated with each effect is fixed and standardized.

The next question is whether the division should be into equal or unequal parts. No reason has become evident to prefer allocating more than half of the total to either of the two mechanisms, so the allocation adopted as the standard is simple equality: 1% for each.

Derivation of the inequality to limit air-to-air effects to 1% begins with an idealized situation, and then in a series of steps, removes the idealizations one-by-one.

Step 1. Idealized Model. Imagine a population of airborne TCAS interrogators, uniformly distributed with a density D (interrogators/nmi²), all transmitting omnidirectionally, all transmitting at a power of 250 watts (the total amount radiated from the antenna), and all interrogating at a common rate I (interrogations/sec). The question is: what is the maximum value of I such that the rate of interrogations received at a victim transponder of MTL = -74 dBm, referred to the antenna (which is the nominal MTL), satisfies:

$$(\text{average reception rate}) \quad (35 \text{ s}) \quad 0.01$$

To answer this, it is necessary to know how many interrogators are within range. Under the stated conditions, the interrogation range is 30 nmi*. Thus letting $T(30)$ be the number of interrogators within a 30 nmi radius:

$$\text{average reception rate} = T(30) \quad I$$

where,

$$T(30) = (30 \text{ nmi})^2 \quad D$$

Thus the maximum value of I is:

$$I = \frac{280}{T(30)}$$

Step 2. Other Power Levels. Generalize the situation by allowing the interrogation power P to be any value, but the same for all interrogations. The interrogation range becomes:

$$R = 30 \text{ nmi} \quad (P/250)^{1/2}$$

*Interrogation range refers to the range at which the power margin is 0 dB. Its value can be calculated (and confirmed to be 30 nmi) using the method given in Ref. 4, page 2.

and the maximum interrogation rate becomes:

$$I = \frac{280}{T(R)}$$

But since

$$T(R) = T(30) \left(\frac{R}{30} \right)^2 = T(30) \left(\frac{P}{250} \right)$$

the relationship limiting interrogation rate can be written:

$$\left(\frac{P}{250} \right) I = \frac{280}{T(30)}$$

Step 3. A Mix of Powers and Rates. Generalize further by allowing different rates and powers for different interrogators. Each interrogator is constrained to operate at some rate and power whose product satisfies the result in Step 2. The issue is to show that the reception rate is still the same as in Step 2, namely 280/sec.

Let f_1, f_2, f_3, \dots denote fractions of the interrogator population corresponding to different rate-power values.

The interrogators constituting the fraction f_1 transmit at a rate = I_1 and power = P_1 , where:

$$\left(\frac{P_1}{250} \right) I_1 = \frac{280}{T(30)}$$

and where:

$$f_1 + f_2 + f_3 + \dots = 1$$

Since the density of type i interrogators is $D f_i$, it follows that the reception rate from all of the type i interrogations is $280f_i$. Thus the total reception rate is just:

$$280f_1 + 280f_2 + 280f_3 + \dots = 280$$

Step 4. Different Powers From Each Interrogator. Generalize further by allowing a mix of powers to be transmitted by any one interrogator providing they satisfy:

$$\frac{P}{250} = \frac{280}{T(30)}$$

where the summation includes all the interrogations in 1 second. The issue is to show that the reception rate is still 280/sec.

Let f_1, f_2, f_3, \dots be defined as in Step 3, and let P_{ij} and I_{ij} denote the power and rate of the interrogations transmitted by an interrogator of type i and power level j . Since the result in Step 2 can be stated:

$$\begin{array}{l} \text{(reception rate} \\ \text{at victim)} \end{array} = \frac{P}{250} \times I \times T(30)$$

it follows that the receptions due to the i - j interrogations occur at a rate:

$$\left(\frac{P_{ij}}{250} \right) \times I_{ij} \times T(30) \times f_i$$

The total reception rate is the sum of such contributions:

$$\begin{array}{l} \text{(total reception} \\ \text{rate)} \end{array} = \sum_i \sum_j \times \frac{P_{ij}}{250} \times I_{ij} \times T(30) \times f_i$$

Since the constraint on each interrogator causes the j -summation to equal $280/T(30)$, the total reception rate is just:

$$\begin{array}{l} \text{(total reception} \\ \text{rate)} \end{array} = \sum_i \frac{280}{T(30)} \times T(30) \times f_i = 280/\text{sec.}$$

Step 5. Elevation Patterns. The results so far apply to idealized omnidirectional antenna patterns. Now consider realistic elevation patterns for aircraft antennas mounted on the top and bottom of the fuselage (still omnidirectional in azimuth).

Elevation effects depend on which antennas are involved: whether interrogations are transmitted from top and bottom, and whether reception is via the top and bottom antenna. The bottom-to-bottom case approximates the ideal omnidirectional characteristics, since as elevation is increased above 0 degrees, the gain of the transmitting antenna decreases (due to increasing obstruction by the fuselage) while the gain of the receiving antenna increases (due to an improvement in the geometry relative to the ground plane). These two effects tend to counteract each other, and the same is true as elevation is decreased. The resulting coverage pattern is similar to omnidirectional, except for being less at very high and very low elevation angles. Thus the limiting formula developed above may reasonably be applied to bottom-to-bottom interrogation, and may be expected to be conservative in the sense that the total received rate will be somewhat reduced by the departures from the ideal at very high and very low elevation angles.

In the case of transmission from a top antenna to a bottom antenna, the coverage pattern is considerably different. It agrees with the bottom-to-bottom coverage at 0 degrees but has more margin above and less margin below. These two departures from omnidirectional behavior may be said to counteract each other: for a given receiving transponder, those interrogators at lower altitudes contribute more (relative to omnidirectional behavior), and those interrogators at higher altitudes contribute less. Here again, the limiting formula developed above appears to be a serviceable control on the total reception rate.

In regard to coverage, top-to-top links behave like bottom-to-bottom links, while bottom-to-top links behave like top-to-bottom links. Thus it seems reasonable to use the formula developed in Step 4, applying the formula independently of whether the interrogations are transmitted from top or bottom antenna.

Step 6. Azimuth Patterns. The results developed up to this point apply to interrogations transmitted omnidirectionally in azimuth. Now the situation is generalized to include directional interrogation. Given that the interrogators all satisfy the formula given in Step 4, the issue is to show that the average reception rate is still 280/sec.

Decompose the total population of interrogations into:

- types of interrogators, $i = 1, 2, 3, \dots$
- classes of interrogations from type i interrogators:
 $j = 1, 2, 3, \dots$, each class having a power P_{ij} and rate I_{ij}
- subdivisions of the i - j interrogations into azimuth sectors small enough to have approximately constant antenna gain, G_{ijk} ; let A_{ijk} be the azimuth width of this sector.

Since the result from Step 2 can be stated:

$$\left(\begin{array}{l} \text{reception rate} \\ \text{at victim} \end{array} \right) = \frac{P}{250} \times I \times T(30)$$

it follows that the receptions due to the interrogations associated with G_{ijk} are at a rate

$$(G_{ijk} \text{ reception rate}) = \frac{P_{ij}}{250} \times G_{ijk} I_{ij} \times \frac{A_{ijk}}{360^\circ} \times T(30) \times f_i$$

$$\left(\begin{array}{l} \text{total reception} \\ \text{rate} \end{array} \right) = T(30) \sum_i f_i \sum_j \frac{P_{ij}}{250} \times I_{ij} \times \left[\sum_k G_{ijk} \frac{A_{ijk}}{360^\circ} \right]$$

The k summation is just the average antenna gain in azimuth, which is unity. Thus this expression reduces to the form treated in Step 4, and simplifies to:

$$\text{total reception rate} = 280/\text{sec.}$$

Step 7. Imperfect Knowledge of Density. Up to this point, the constraint on rate and power:

$$\sum \frac{P}{250} < \frac{280}{T(30)}$$

has been expressed in terms of the density D of TCAS interrogators, through the factor

$$T(30) = \pi (30 \text{ nmi})^2 D$$

which is the average number of interrogating aircraft within 30 nmi. The next question addressed is how to implement this constraint, or an approximation to it, on board each interrogating aircraft where an exact knowledge of D is not available.

One obvious approach is to have each aircraft count the actual number of interrogators within 30 nmi and use this count N as an estimate of T(30). This would probably work well when T(30) is large, since some aircraft would by chance obtain a higher than average value of N and others would obtain a lower than average value. When T(30) is large, these chance deviations would be small fractions of the mean value, so that the penalty resulting from a larger-than-average value of N would not be severe, and furthermore the total reception rate would be nearly the same as if each interrogator had used the exact value of density. There is, however, a bias, due to the fact that:

$$\text{average } \frac{1}{N} > \frac{1}{\text{average}(N)}$$

the bias is in the direction which would increase interference if this simple rule were used. The bias is small when T(30) is large, but can become exceedingly large when T(30) is small. Consider the case in which some interrogating aircraft obtains a count N=0. Then using the constraint:

$$\frac{P}{250} < \frac{280}{N}$$

this aircraft would be able to interrogate at arbitrarily high rates and powers, and so a reception rate of < 280/sec. could not be assured. This form of limiting standard would be unsatisfactory.

Consider the simple change of adding 1 to N.

$$\sum \frac{P}{250} < \frac{280}{N+1}$$

This change effectively biases the total interferences back in the other direction (reducing interference). It also eliminates the problem associated with occurrences of $N=0$. Furthermore, this change has a negligible effect when $T(30)$ becomes large, under which conditions there was no need for such a change. This formulation, therefore, seems to be a satisfactory way of dealing with the imperfect knowledge of density.

Step 8. Non-Uniform Aircraft Density. In reality, of course, aircraft density is not uniform as has been the idealization throughout the above. Higher densities around metropolitan areas are to be expected and have been observed through measurements (Ref. 5).

Even where density is not constant, it seems reasonable to use the same interference limiting standard as derived in Step 7. This limiting inequality has a built-in adaptability to density; rather than being based on any prespecified density, the inequality causes each interrogator to adjust to the local density around that interrogator. For example, in any region where there is a uniform rate of change of density, each interrogator would be controlled by the average density in a region centered at that aircraft. A victim transponder would receive interrogations from a higher density side and a lower density side. The higher density side would have more numerous interrogators, but with each transmitting at a proportionately reduced rate; and vice versa for the low density side. Thus the total effect at the victim transponder would be approximately the same as if the density were uniform.

5.1.1.2 Suppression of Own Transponder

Suppression of own transponder can be limited to 1% or less by constraining

$$\frac{\sum M(i)}{1 \text{ sec.}} < 0.01$$

where the summation is over 1 second, and where the extent of the on-board transponder suppression period accompanying the i th interrogation, $M(i)$, may vary as a function of i . This is rewritten to appear in the limiting standard in the form

$$\sum M(i) < 0.01 \text{ sec.}$$

5.1.1.3 Fruit

The basis for the fruit-limiting inequality is that the Mode C fruit generated by TCAS should not be greater than 20% of the present peak transponder reply rate. Such an increase will not significantly affect the

performance of the ground-based surveillance system. Furthermore, the peak reply rate has steadily declined over the last decade as a result of programs to reduce overinterrogation. It is expected that this trend will continue and that the rates will decline even more when existing sensors are replaced with Mode S sensors.

Currently, the peak Mode C reply rate in areas of intense ground interrogation activity is approximately 200 replies in a one-second period, provided that all interrogators are operating normally (Ref. 8). (Omnidirectional sensors interrogating at high rates, or sensors operating without sidelobe suppression can result in reply rates considerably higher than 200 per second; but these are not normal operating conditions).

Thus, for any transponder, the Mode C reply rate due to TCAS interrogations, RRT, must be less than 0.2 times 200 per second. That is,

$$RRT < 40 \text{ per sec.}$$

RRT is proportional to the number of detectable whisper-shout sequences received by the transponder each second (reduced by a transponder beam-sharpening factor) and it is proportional to the average number of replies transmitted by the transponder in response to each whisper-shout sequence.

That is,


$$RRT = \frac{1}{B} \times (SW) \times (RPW) < 40 \text{ per sec.,}$$

where B is the beam sharpening factor, SW is the total number of whisper-shout sequences detected by the transponder each second, and RPW is the average number of replies transmitted by the transponder in response to a whisper-shout sequence.

The significance of the beam sharpening factor is illustrated in Fig. 5-2 for a four-beam directional antenna. The area in which transponder replies are generated is a subset of the area in which the whisper-shout interrogations can be detected, because the P2 beam-sharpening control pattern suppresses transponders outside. (For example, measurements of the Dalmo-Victor four-beam antenna indicate that the detection area is approximately 20% larger than the reply area. So, for that antenna, B = 1.2).

Using reasoning identical to that of the derivation of the first inequality presented above, the sum of the whisper-shout sequences detected each second is

$$SW = (NT + 1) \sum \frac{P_{max}}{250},$$



$$\frac{1}{B} = \frac{\text{reply area}}{\text{detection area for W-S interrogations}}$$

Fig. 5-2. Beam sharpening factor, B.

where NT is the number of other TCAS units within a nominal 30-nmi detection range, and Pmax is the power (in watts) of the highest-power interrogation transmitted in each whisper-shout sequence.

In the specific whisper-shout sequence used in the December 1982 Los Angeles testing, Pmax is one-fifth of the sum of the total radiated powers for the individual Mode C interrogations, PAi. That is

$$\frac{P_{\max}}{250} = \frac{1}{5} \sum_i \frac{PA(i)}{250}$$

This whisper-shout sequence has been experimentally determined to generate approximately 2.5 replies per transponder on average; thus RPW = 2.5.

Substituting these factors into the above equations and rearranging terms gives the third inequality in the standard form:

$$\frac{1}{B} \sum_i \frac{PA(i)}{250} < \frac{80}{NTA+1}$$

To this inequality, a fixed upper limit is added to control interrogations in cases when NT is small. This limit is based on the power sum values (left hand side of the above inequality) for the particular designs developed in this program, the designs tested in December 1982 in Los Angeles (Fig. 3-15). These power sum values are

$$\frac{1}{B} \sum_i \frac{PA(i)}{250} = \begin{matrix} 5 & \text{omnidirectional design} \\ 2.5 & \text{directional design} \end{matrix}$$

The third limiting inequality becomes

$$\frac{1}{B} \sum_i \frac{PA(i)}{250} < \text{the smaller of } \left[\frac{80}{NT+1}, 5 \right]$$

Thus the limit on the right hand side remains constant as NT increases up to 15.

A similar fixed upper limit is added to the first limiting inequality. Here again the value of the limit for NT = 15 is taken as a fixed upper limit even for lower values of NT.

$$\sum_i \frac{P(i)}{250} < \text{the smaller of } \left[\frac{280}{NT+1}, 18 \right]$$

5.1.2 Interference Simulation

Following the analytical derivation of interference limiting standards, these standards are being tested through an interference simulation conducted by ECAC. This is a large scale simulation encompassing an extensive RF environment of many transmitters and receivers, while also including a detailed representation of events at the microsecond level. A number of scenarios in the Los Angeles Basin were simulated. The simulation includes specific ground-based SSR's whose locations and transmitting characteristics (such as transmitter power, antenna scan rate, and interrogation repetition frequency) are taken from a Master File of existing interrogators. Aircraft traffic is represented as a set of specific aircraft locations and types, taken from the traffic model in Ref. 9. A very large amount of computer time is required to run the simulation for each scenario. A detailed description of the simulation is given in Ref. 10.

The simulated scenarios are in pairs: with and without TCAS activity. The subject whose performance is being examined is the SSR at Long Beach. The simulation determines for each scenario:

- % in track, the percentage of aircraft in track at a given time
- % updated, for the aircraft in track, the percentage whose tracks are updated with a new measurement of range and altitude in a given scan

The main simulation results are in this form, relating to performance attributes that may be evident to users (that is, to air traffic controllers using the SSR displays). Simulation results were also generated for more detailed performance attributes, such as reply ratio and fruit rate, which would not be directly evident to users.

1030 MHz Broadcast. At an early stage in the simulation study it was observed that there was a potential problem with interference limiting in regard to the estimation of NT. NT is the means by which a TCAS II unit estimates the local density of TCAS II interrogators (Fig. 5-1). At that time in the study, the concept for estimating NT was to count aircraft according to receptions of their squitters (which are transmitted at 1090 MHz). It was soon realized that this counting was made quite inaccurate by the effects of fruit. As a result it was decided to change the concept for estimating NT to a technique based on 1030 MHz broadcasts. The interference conditions in the 1030 MHz band are much less severe. In this concept, each TCAS II unit spontaneously transmits self-identifying broadcasts at a rate of one in 10 seconds. The simulation study showed that the NT inaccuracy problem was overcome using this concept.

Main results. The simulation study is not yet complete. Interim results for the main performance attributes are given in Table 5-1. Results are given for three traffic models, the highest density case having 743 aircraft within 60 nmi of LA International Airport. The middle case, 474 aircraft, corresponds approximately to the high density condition for which TCAS II is being designed.

TABLE 5-1

INTERIM RESULTS FROM INTERFERENCE SIMULATION

Scenario				Main results	
Total number of aircraft	Avionics mix			% in track	% updated
	Mode A,C	Mode S	TCAS II		
328	75%	25%	0	79	95
328	75%	14%	11%	79	95
474	75%	25%	0	80	92
474	75%	14%	11%	80	92
743	75%	25%	0	73	91
743	75%	14%	11%	73	91

These results show that the TCAS activity has no effect on SSR performance. The bases from which the interference limiting standards were derived are in fact low enough so that the presence of TCAS II aircraft in large numbers would not be evident to users of ground surveillance equipment.

5.2 Interrogation Limiting Algorithm

For the specific TCAS II designs developed for testing in Los Angeles, the full power Mode C interrogations are at:

250 watts, omnidirectional design (Fig. 3-5)
80 watts, directional design (Fig. 3-15)*

The corresponding values of the Mode C power sum as limited by the standard are:

$$\frac{1}{B} \sum \frac{P_A}{250} = 4.9, \text{ omnidirectional design} \\ 4.8, \text{ directional design.}^*$$

These are within the maximum limit of 5, but are not far below. Thus these peak power levels are nearly the maximum values permitted by the standard.

The purpose of the interrogation limiting algorithm is to ensure that the TCAS equipment conforms to the interference limiting standard of Section 5.1. This is accomplished by controlling the nominal range at which the presence of an aircraft is first detected.

To control the Mode C detection range the number of transmitted whisper-shout levels is varied. If the range is to be reduced, the highest power interrogation last used is omitted, thereby causing some distant aircraft not to receive an interrogation.

The detection range for Mode S-equipped aircraft is controlled by varying the MTL used to detect squitters and fruit. As discussed in Section 4.4.3, this variation is matched by a change in the power used to interrogate aircraft assigned the acquisition state. The two controls are coordinated to keep the detection range in the forward direction comparable in Mode S and Mode C.

5.2.1 Structure

The algorithm exercises control through the application of the four steps that are discussed below and which are embodied in the flow diagram shown in Fig. 5-3. The steps involve interference limiting inequalities (1), (2), (3) given in Fig. 5-1. In evaluating these inequalities, 16-second averages of the Mode S parameters are used, and current or anticipated values of the Mode C parameters are used.

* Obtained using $P_{\max} = 320 \text{ watts} \times (90^\circ/360^\circ) = 80 \text{ watts}$, and $\beta = 90^\circ/(360^\circ/2.5)$, from Sec. 3.2.5.

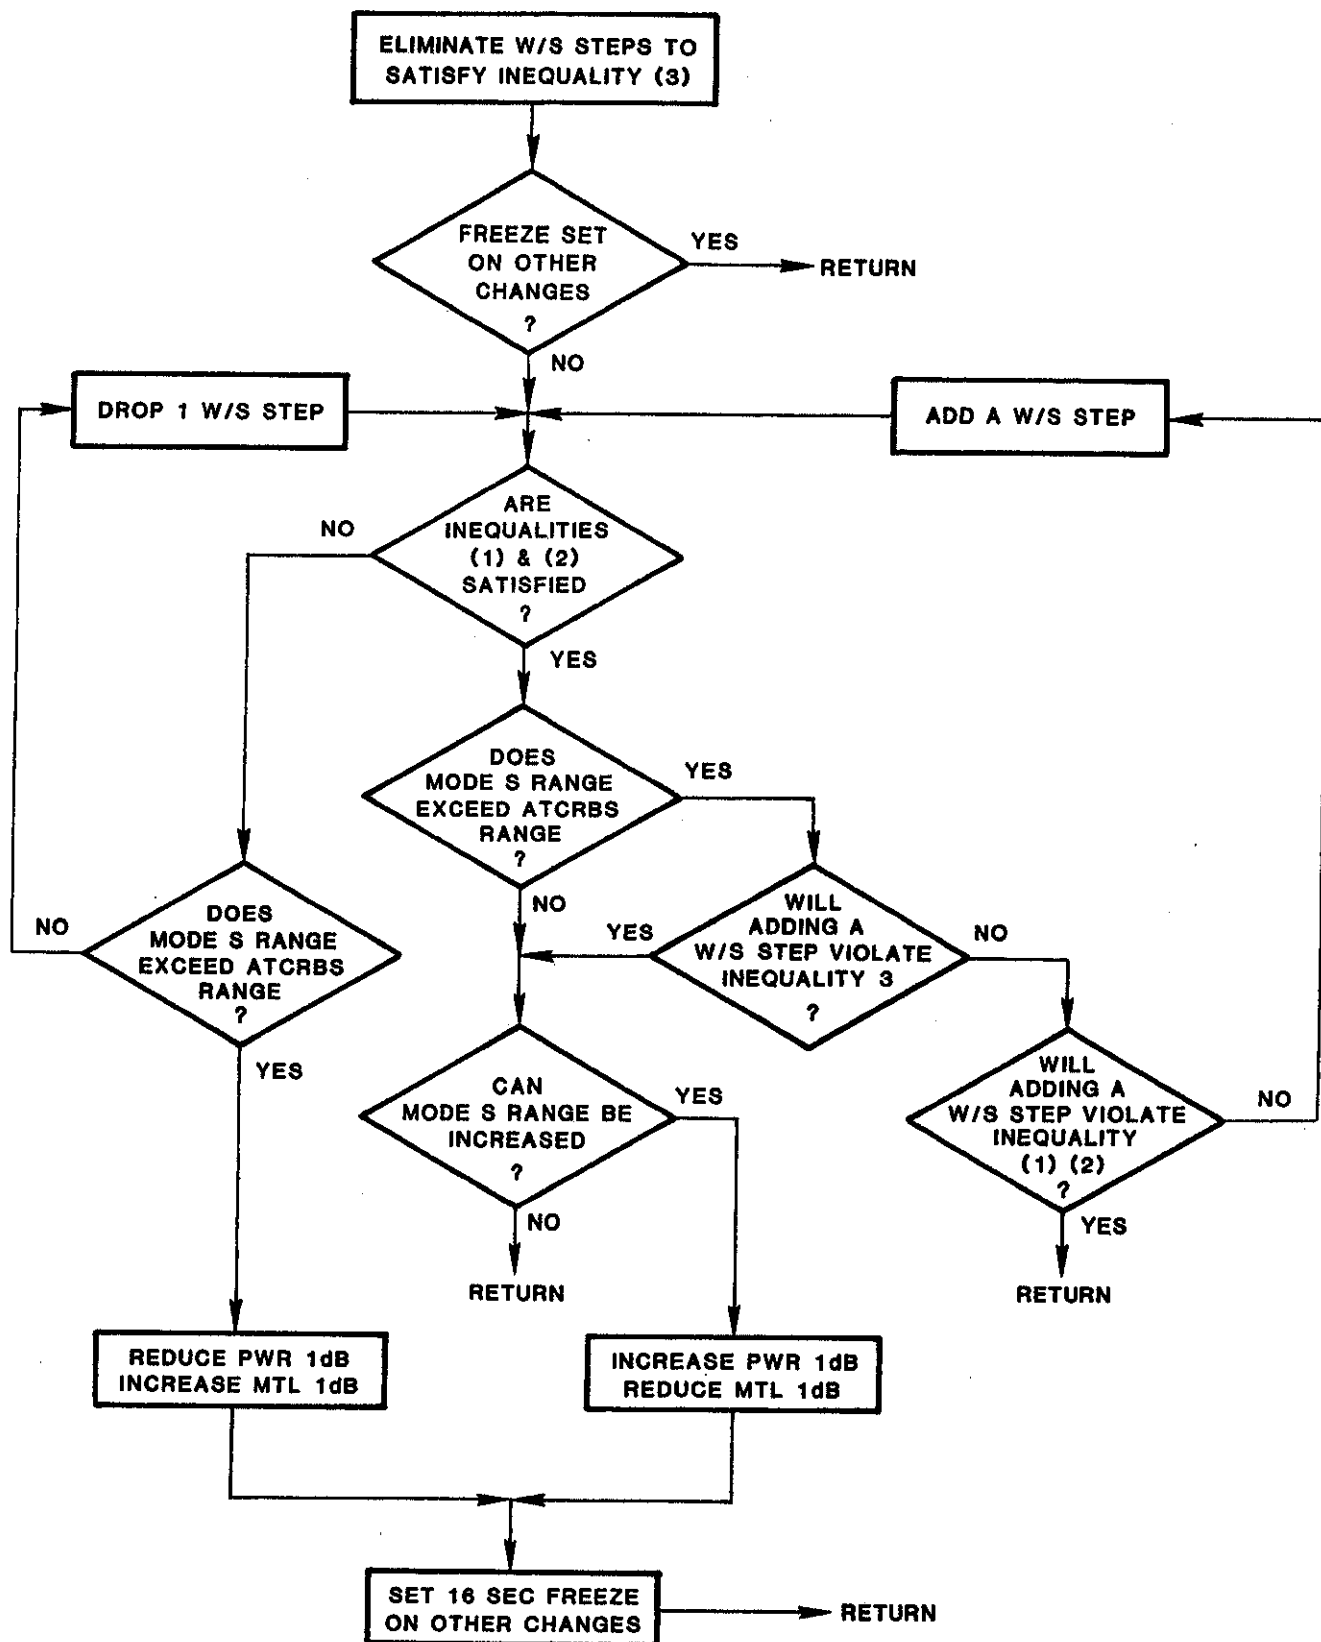


Fig. 5-3, Interrogation limiting flow diagram.

given in Fig. 5-1. In evaluating these inequalities, 16-second averages of the Mode S parameters are used, and current or anticipated values of the Mode C parameters are used.

The first step in the control process is to reduce the number of whisper-shout levels tentatively scheduled for use during the present scan if either

- a) Inequality (3) is violated, or
- b) Inequality (1) or (2) is violated and the Mode S surveillance range of the last scan does not exceed the Mode C surveillance range that would result from use of the scheduled whisper-shout sequence.

Whisper-shout levels are eliminated in the order dictated by the design of the Mode C processor and the number of levels eliminated is just large enough to ensure that neither of the above conditions is satisfied. The whisper-shout level tentatively scheduled for use is initialized at that used on the last scan.

The relative ranges for Mode S and Mode C surveillance are determined from the estimated maximum power densities seen by head-on collision targets with Mode S and Mode A,C transponders respectively. If the transponder sensitivities were identical, the Mode S range would be more or less than the Mode C range according to whether the Mode S power density was more or less than the Mode C density. Since Mode A,C transponders may have somewhat lower sensitivities than Mode S transponders, the Mode C range is assumed to be greater than the Mode S range if, and only if, the Mode C power density exceeds the Mode S power density by 3 dB. The power density is determined by the power input to the antenna and the antenna radiation pattern.

The second step in the controlling process is to reduce the Mode S interrogation power last used for acquisition by 1 dB and to increase the MTL used to detect Mode S squitters and fruit by 1 dB if

- c) Inequality (1) or (2) is violated and the Mode S surveillance range of the last scan exceeds the Mode C surveillance range that would result from use of the scheduled whisper-shout sequence.

Once such a change has been made the only other change allowed during the ensuing 16 seconds is a reduction in the number of whisper-shout levels if such is needed to satisfy Inequality (3). This 16-second freeze allows the effect of the Mode S changes to become apparent since the 16-second averages used in Inequalities (1) and (2) then will be determined by the behavior of the system since the change.

The third step is to add a whisper-shout level to those tentatively scheduled when it is not prevented by a 16-second freeze and the following conditions are satisfied:

- d) Inequalities (1), (2), and (3) are satisfied and will continue to be after the level is added, and,
- e) The Mode S surveillance range of the last scan exceeds the Mode C surveillance range that would result from use of the scheduled sequence.

As many levels are added as possible without violating d) or e) above.

Finally, if condition d) above is satisfied, but condition e) is not, an estimate (see 5.2.2) is made of the effects of increasing the Mode S interrogation power for acquisition by 1 dB and reducing the MTL for detecting Mode S squitters and fruit by 1 dB. If the estimate indicates that Inequalities (1) and (2) will not both continue to be satisfied, the 1 dB change is not made. If the estimate indicates that they will both continue to be satisfied, the 1 dB change is made and no further changes in either the Mode C or Mode S parameters are made for the ensuing 16 seconds except as described in connection with condition c).

5.2.2 Parameter Estimates

The estimate of the consequences of increasing the Mode S interrogation power, and decreasing the MTL for detecting squitters and fruit, by 1 dB is based upon the last available 16 second averages of the following Mode S parameters.

- PI_A : the contribution to Inequality (1) of acquisition state interrogations
- PI_T : the contribution to Inequality (1) of the track state interrogations
- IA : the contribution to Inequality (2) of acquisition state interrogations
- IT : the contribution to Inequality (2) of track state interrogations
- f : the fraction of aircraft in the track state that were interrogated with the maximum allowable interrogation power on the last scan.

The contribution of the different interrogations to the inequalities are separated because they are affected differently by the power change under consideration. For example, the acquisition state contribution will always increase, partly because the increased surveillance range causes more targets to be acquired per unit time and partly because a larger interrogation power is used for all acquisition interrogations. On the other hand, the track-state contribution will change only if the interrogation power to some track-state aircraft equals the interrogation power used for acquisition. The question is: what changes in these quantities are expected to result from the 1 dB changes in the MTL and the interrogation power used for acquisition?

The change in I_A . A 1-dB increase in power should increase the detection range about 12%, from R_0 to $1.12 R_0$. To first order, the rate at which aircraft are assigned to the Acquisition state should be proportional to range so the percentage increase in the rate should be roughly equal to the increase in detection range, i.e., 12%. Thus the estimated value of I_A after the change is $1.12 (I_A)^\circ$ where $(I_A)^\circ$ is the measured value of I_A before the change.

The change in PI_A . Since the interrogation power has increased 1 dB, or 25% the estimated value of PI_A after the change is $(1.25)(1.12)(PI_A)^\circ$ or $1.4 (PI_A)^\circ$ where $(PI_A)^\circ$ is the value of PI_A before the change. To provide some margin against the oscillation that might result from under-estimating these increases, the following estimates were adopted and are used in the simulations described in Section 5.2.3.

$$I_A = 1.25 (I_A)^\circ \quad (4)$$

$$PI_A = 1.5 (PI_A)^\circ \quad (5)$$

The changes in I_T and PI_T . If the interrogation power for all of the aircraft assigned the track state is less than that used for acquisition, neither I_T nor PI_T should change appreciably when the interrogation power for acquisition is increased 1 dB and the MTL for squitters/fruit is decreased 1 dB.* The change should still be negligible when a small fraction of the track state aircraft are interrogated at the acquisition power. Therefore, for values of f no larger than 0.1 it will be assumed that the values of I_T and PI_T are not changed by the 1 dB change in the MTL and acquisition interrogation power. That is, for f no larger than 0.1,

$$I_T = (I_T)^\circ \quad (6)$$

$$PI_T = (PI_T)^\circ \quad (7)$$

where $(I_T)^\circ$ and $(PI_T)^\circ$ are the values measured before the change in Mode S parameters.

For values of f exceeding 0.1 the effect of the change upon I_T and PI_T depends upon the distribution of aircraft and the conditions under which they are assigned to the track state. For a uniform distribution of aircraft and for the surveillance algorithms discussed in Chapter 4, the number of aircraft assigned to the track state will increase by about 25%. That is, for f greater than 0.1,

$$I_T = 1.25(I_T)^\circ \quad (8)$$

*Some changes will occur if additional aircraft are assigned to the Track state.

The change in the value of PI_T depends in detail upon the distribution of the track-state aircraft that are interrogated at maximum power. For simplicity, the change will be over-estimated by assuming that all Track state aircraft are interrogated at maximum power**. Then, for f greater than 0.1,

$$PI_T = (1.25)^2 (PI_T)^{\circ} \approx 1.5 (PI_T)^{\circ} \quad (9)$$

where $(PI_T)^{\circ}$ is the average value of PI_T before the change in Mode S parameters is made.

Equations 4 through 9 provide the needed estimates of the effects of changing the Mode S parameters upon Inequalities (1) and (2). To determine whether or not the change is feasible, the average values of $(I_A)^{\circ}$, $(I_T)^{\circ}$, $(PI_A)^{\circ}$ and $(PI_T)^{\circ}$ last used in evaluating the inequalities are replaced by the values given in the above equations. If the inequalities are still satisfied the change is made. Otherwise, it is not.

5.2.3 Performance Evaluation

Two questions arise concerning the operational performance of the interrogation limiting algorithms: do they cause the interference limiting standard to be met and do they result in a value of the MTL for which the collision protection is satisfactory? These questions were answered by simulation for the environments of interest. The conclusion is that the interference limiting standard is met and that the MTL for squitters and fruit will be small enough to achieve the desired collision protection.

The aircraft environments in which the protection is to be provided are discussed below. Then the essential features of the simulation are described. This is followed by a discussion of the results obtained.

5.2.3.1 Operating Environments

The environments in which protection is to be provided were discussed in Section 4.7.1. Two of them are low-altitude low-speed environments for which it was found that the MTL could be raised by 6 dB without sacrificing the desired protection. They differ in that one has a peak aircraft density of $0.3/\text{nmi}^2$ and pertains to the situation in which both error correction decoding and power programming are employed while the other has a peak density of $0.15/\text{nmi}^2$ and is used when neither error correction nor power programming is employed.

The other two environments involve high-altitude, high-speed encounters for which it was found that an MTL increase of no more than 3 dB results in satisfactory collision protection. One of these environments has a peak aircraft density of $0.06/\text{nmi}^2$ and is used when both error correction and power programming are employed. The other applies when neither of these techniques is employed and has a peak density of $0.03/\text{nmi}^2$.

**For a uniform-in-area distribution, the error in the estimate is not large.

In Section 4.7.1 several parameters of the operational environment were left unspecified since they did not influence the performance quantities of interest there. Those parameters, which will now be specified, are: the number of other TCAS units operating within 30 nmi of the TCAS unit under consideration, the fraction of the transponder-equipped aircraft that carry Mode S transponders, the altitude distribution of those aircraft, and the distribution of their airspeeds.

The number of other TCAS operating within 30 nmi is specified to be 30. The basis on which the other parameters were chosen is as follows. First, it will be assumed that all of the transponders are Mode S. This is a worst-case assumption since, as the fraction of Mode S transponders increases, the MTL for squitters and fruit increases, thereby reducing collision protection.

Two altitude distributions will be employed. In one the aircraft are uniformly distributed in altitude between two limits that can be specified arbitrarily. In the other, their density is that shown in Fig. 5-4 which is derived from measurements made at Long Beach, California and will be called the Long Beach altitude density. In that measurement the altitudes of aircraft above 14,500 feet were not recorded. The 15% of the aircraft that were found to be above that altitude are uniformly distributed from 14,500 feet to 40,000 feet.

The speeds of the aircraft are taken to be random variables whose probability density varies with altitude. The density used at altitudes of less than 10,000 feet reflects the large fraction of low speed aircraft that are encountered there. It is a truncated decaying exponential that begins at an airspeed of 70 kts and is of the form $\exp(\text{speed}/30 \text{ knots})$. For altitudes above 10,000 feet a uniform density is employed with the range of speeds being 200 to 400 kts below 15,000 feet and 300 to 600 kts above 15,000 feet.

5.2.3.2 Simulation of the Environment

To evaluate the performance of the system, the motion of a TCAS-equipped aircraft moving through the environments described above was simulated. The environments were simulated by assigning random altitudes and airspeeds to the aircraft. They were also assigned headings that were uniformly distributed around the compass and positions that were uniformly distributed within a square whose size could be specified. The density was controlled by varying the number of aircraft in the square.

The TCAS aircraft flew at an assigned airspeed and followed an arbitrarily specified altitude profile. The square moved along with the TCAS aircraft which was always at its center. Each of the other aircraft maintained a constant airspeed, altitude, and heading except when they reached the edge of the square. Then they were removed from the simulation and reintroduced at a point on the opposite side of the square with the same airspeed, altitude, and heading.




Fig. 5-4. Altitude density function.

The spatial density of aircraft that results from this simulation varies slightly with time. Examples of that variation are shown in Fig. 5-5 for the situation in which 200 aircraft with the Long Beach altitude density are initially distributed over a square that is 25.75 nmi on a side. The speed of the TCAS aircraft for this figure was 250 kts. For purposes of comparison, the ensemble average density that would result from a uniform distribution of aircraft is also shown. It is apparent from the figure that the simulation provides a relatively constant and uniform density of aircraft.

The simulated environment described above was combined with the link simulation described in Section 4.7.2 to create the signal environment in which TCAS is intended to operate. Those signals were then used as inputs to the simulation of the Mode S surveillance processor to determine the variation in the MTL caused by the action of the interrogation limiting algorithm. The results of those simulations are summarized below first for low-altitude operations and then for high-altitude operations.

5.2.3.3 Low-Altitude Results

Representative simulation results for the low-altitude environment are shown in Fig. 5-6. There the variation with time of the MTL used for the detection of Mode S squitters and fruit is labeled MTL and the three curves labeled "energy", "number", and "fruit" are the normalized values of the interference limiting inequalities given, respectively, by (1), (2), and (3) of Fig. 5-1. The normalizations are such that an inequality is satisfied if the value is no larger than one and is violated if it exceeds one. The MTL value in the figure is the deviation from nominal.

The figure is for a TCAS at 5,000 feet with an airspeed of 250 kts and a maximum capable airspeed of 300 kts in an environment of 200 aircraft that were initially distributed uniformly within a 25.75-nmi square. That corresponds to an aircraft density of $0.3/\text{nmi}^2$ within the square. Altitudes were assigned to the aircraft in accordance with the Long Beach altitude distribution of Fig. 5-4. Finally, 30 TCAS aircraft were operating within 30 nmi of the TCAS unit being simulated.

The salient features of the results are as follows. First, the interference limiting inequalities are satisfied throughout the simulation. Second, the largest of the three normalized limits is always nearly equal to one, so the largest possible surveillance range is being maintained. Third, the MTL for the detection of squitters and fruit varies from its nominal value by a maximum of 3 dB and is usually either 1 or 2 dB higher. Thus it is at least 3 dB less than the maximum increase of 6 dB for which satisfactory collision protection at low altitudes is assured.

Throughout the simulation the number of whisper-shout levels used remained constant, as can be inferred from the invariance of the fruit limit. Its value was 81. The MTL, rather than the number of whisper-shout levels, changed because the estimated surveillance range for Mode S targets continued to exceed that for Mode C targets.

AIRCRAFT DENSITY (AIRCRAFT/1-NMI RANGE BAND)




Fig. 5-5 Density of aircraft at three points in time.

Conditions: Aircraft density = $0.3/\text{NM}^2$

TCAS airspeed 250 kt.

Thirty other TCAS operating within 30 NM.

Long Beach altitude distribution:

TCAS altitude = 5000 ft.




Fig. 5-6. Variation of normalized interference limits and MTL in simulated high density environment.

Table 5-2 gives the interval over which the MTL varied in a series of simulations that differ in varying degrees from the one just described. In all of them the interference limiting inequalities were satisfied throughout a 200- to 300-second simulation.

The first row of the table corresponds to the simulation described by Fig. 5-6. The second differs in that the aircraft were uniformly distributed in altitude from 0 to 10,000 feet. Although the MTL change is affected by the change in altitude distribution it remains small enough to provide satisfactory collision protection for the encounters discussed in Section 4.7.

When the TCAS equipped aircraft is either climbing or descending, the MTL can increase beyond the values just discussed because aircraft are assigned to the Acquisition state as the altitude band about the TCAS sweeps over them. Row three of the table shows this effect when the TCAS-equipped aircraft descends from 11,000 feet to 5,000 feet at a rate of 3,000 feet per minute in the Long Beach altitude environment. The descent causes the MTL change to peak at 5 dB. An examination of the simulation record showed that this peak persisted for about 70 seconds.

Altitude changes have a more significant effect when they are more rapid or involve a descent from a low density airspace into a high density airspace. Then the number of aircraft assigned to the Acquisition and Track states increases, in part, because they enter the altitude band at a greater rate and, in part, because the rate at which aircraft enter the band exceeds the rate at which they exit from it. This is illustrated by the fourth row of the table which describes a descent from 15,000 feet to 5,000 feet at a rate of 5,000 feet per minute for the Long Beach altitude distribution. The corresponding variations of the MTL and the three normalized interference limiting inequalities are given in Fig. 5-7. Note the peak transient value of 6 dB for the MTL change as the processor attempts to interrogate all of the aircraft that have suddenly become potential collision threats. Even at this peak value the desired collision protection is provided. Moreover if the maximum capable airspeed were larger than 300 kts, the initial value of the MTL would be increased and its peak value would be further decreased.

The next three table entries show the benefits of power programming and error correction decoding for the situation described by Fig. 5-6 and the first row of the table. If error correction decoding is used, but power programming is not, the MTL will vary from 5 to 6 dB above nominal rather than from 0 to 3 dB. This is still acceptable, but little margin is then left to allow for transients during descents. If neither error correction decoding nor power programming is used, the MTL remains at 6 dB above nominal. The use of power programming alone causes the MTL to vary from 2 to 4 dB above nominal.

TABLE 5-2

VARIATION OF MTL FOR A LOW ALTITUDE ENVIRONMENT

Variation of MTL for a TCAS with an airspeed of 250 Kts and a maximum capable airspeed of 300 Kts in an environment of 200 aircraft. Thirty other TCAS are operating within 30 nmi. Except, as noted the aircraft are initially distributed according to the Long Beach altitude density and are uniformly distributed in a square of width 25.75 nmi, to give a density of $0.3/\text{nmi}^2$, and error correction decoding and power programming are used.

<u>TCAS ALTITUDE PROFILES</u>	<u>MTL VARIATION</u>	<u>NOTES</u>
Level: 5,000 ft	0-3 dB	
Level: 5,000 ft	2-4	(1)
Descent: 11,000 to 5,000 ft; 3,000 FPM	1-5	
Descent: 15,000 to 5,000 ft; 5,000 FPM	1-6	
Level: 5,000 ft	5-6	(2)
Level: 5,000 ft	6	(3)
Level: 5,000 ft	2-4	(4)
Level: 5,000 ft	2-4	(5)

(1) Uniform altitude density 0 to 10,000 ft

(2) No Power Programming

(3) Neither Power Programming Nor Error Correction Decoding

(4) No Error Correction Decoding

(5) Neither Power Programming Nor Error Correction Decoding, 38-nmi Square, Aircraft Density $0.14/\text{nmi}^2$

Conditions: Aircraft density = $0.3/\text{NM}^2$
 TCAS airspeed 250 kt.

Long Beach altitude distribution
 Thirty other TCAS operating within 30 NML




Fig. 5-7 Variation of normalized interference limits during a rapid decent.

Conditions: Aircraft density = $0.3/\text{NM}^2$ Long Beach altitude distribution
 TCAS airspeed 250 kt. Thirty other TCAS operating within 30 NML




Fig. 5-7 Variation of normalized interference limits during a rapid decent.

Although a TCAS that employs neither error correction nor power programming is not intended to provide collision protection in the density of aircraft discussed above, it is intended to provide such protection when the peak density drops to $0.15/\text{nmi}^2$. As indicated by the last row of the table, the MTL will not increase by more than 6 dB; hence the desired protection will be provided.

5.2.3.4 High Altitude Results

Figure 5-8 shows the variation of the normalized interference limits and the MTL for the extreme situation in which a TCAS-equipped aircraft with an airspeed of 600 kts descends from an altitude of 29,000 feet to an altitude of 11,000 feet at a rate of 5,000 feet per minute in an environment of 200 other aircraft of which 30 are TCAS-equipped. The aircraft are initially distributed uniformly over a 57.3 by 57.3 nmi square, and are distributed in altitude according to the Long Beach density. The figure corresponds to the situation in which both power programming and error correction are employed. The MTL is nominal for most of the descent but increases by 3 dB as the aircraft descends into the more densely populated airspace below 14,500 feet. Thus, the performance is acceptable even in this extreme situation.

The interval over which the MTL varies in a number of situations is given in Table 5-3. The first row of the table applies to the situation just described and the second differs from it in that the TCAS altitude is constant at 25,000 feet. In the later instance the MTL does not change during the entire simulation.

The third row applies when the TCAS altitude is 11,000 feet instead of 25,000 feet. It reflects an unrealistic situation in that the TCAS airspeed is taken to be 600 kts at this altitude. However, it is a useful example in that it illustrates the inability of any system to satisfy the interference limiting standard and provide collision protection in all situations. In particular, the peak value of the MTL change is 4 dB which exceeds the value for which collision protection can be assured. An examination of the simulation record shows that this peak persisted for one 30-second period out of 300 seconds. Thus even in this unrealistic situation a substantial amount of protection is provided.

In the above simulations only a very small fraction of the 200 aircraft had altitudes near enough to that of the TCAS aircraft to be interrogated by it. A measure of the number of co-altitude aircraft against which satisfactory collision protection can be provided is given by row four of the table. It indicates that the MTL will not exceed 2 dB when 30 other TCAS-equipped aircraft are co-altitude with the TCAS unit in question and are contained within a square of width 57.3 nmi.

Row five of the table provides another measure of the system's robustness. It applies to the situation in which a TCAS-equipped aircraft at 25,000 feet overflew a high density terminal area containing 200 aircraft within a 25.75 nmi square corresponding to a density of $0.3/\text{nmi}^2$. These

Conditions: Aircraft density = $0.06/\text{NM}^2$ Long Beach altitude distribution.
 TCAS airspeed 600 kt. Thirty other TCAS operating within 30 NM.




Fig. 5-8. Variation of normalized interference limits and MTL during a high altitude descent.

aircraft are distributed in altitude according to the Long Beach density and 30 of them carry operating TCAS units.

The above results are for systems that employ error correction decoding and power programming. As discussed in Section 4.7.1, systems that employ neither are intended to provide satisfactory collision protection at aircraft densities of at most $0.03/\text{nmi}^2$ with no more than 15 TCAS units operating within 30 nmi. That is, the MTL change should not exceed 3 dB under these conditions. The entries in rows six and seven of Table 5-3 show that protection is provided under these conditions even when the TCAS aircraft descends from 29,000 feet to 11,000 feet at 5,000 fpm. Indeed the MTL remains nominal throughout the descent.

5.3 Transponder Suppression

Airborne measurements of interrogation backscatter have been carried out to assess the required duration of self suppression from the TCAS interrogator to the on-board Mode S transponder.

5.3.1 Need for Re-examination of Mutual Suppression

To avoid interference between the various L-band transmitters on an aircraft (for example, a DME interrogator and an SSR transponder), it is common practice for them to interact through an arrangement of "mutual suppression". When such a unit transmits at L-band, it supplies a suppression pulse to a mutual suppression bus. Each system receiving the suppression pulse can make use of this information to disregard any receptions during this brief period, often simply by gating off the receiver for the duration of the suppression pulse.

In TCAS II it is appropriate for the TCAS II interrogator to suppress the onboard Mode S transponder, both of which operate at the same radio frequency (1030 MHz). During the BCAS development program it was realized that the transponder should be kept in suppression for considerably longer than the duration of the transmitted interrogation because backscattered echos from the terrain beneath the aircraft would often cause the transponder to reply. Such replies interfere with TCAS surveillance, both because of the addition to the fruit environment they constitute and because they occur in the active range window of the BCAS or TCAS receiver.

The duration of transponder suppression in the BCAS design as conservatively set at 200 μs , and extensive airborne testing showed that this period was long enough to prevent self interrogation. As BCAS evolved into TCAS, this suppression time needed to be reexamined because of the increase in the number of interrogations per second.

Measurements. Direct measurements of interrogation backscatter were made using the Airborne Measurements Facility (AMF). Mode C and Mode S interrogations were transmitted alternating between top and bottom antenna, and all pulses detected at 1030 MHz were recorded.

TABLE 5-3

VARIATION OF MTL FOR A HIGH ALTITUDE ENVIRONMENT

Variation of the MTL for a TCAS with an actual and maximum capable airspeed of 600 Kts. Except as noted the environment contains 200 aircraft that are initially distributed according to the Long Beach altitude density and are uniformly distributed within a 57.3 nmi square to give a density of 0.06/nmi², error correcting decoding and power programming are used, and 30 other TCAS are operating within 30 nmi.

<u>TCAS ALTITUDE PROFILES</u>	<u>MTL VARIATION</u>	<u>NOTES</u>
Descent: 29,000 to 11,000 ft; 3,000 FPM	0-3 dB	
Level: 25,000 ft	0	
Level: 11,000 ft	0-4	
Level: 25,000 ft	0-2	(1)
Level: 25,000 ft	2-3	(2)
Level: 25,000 ft	0	(3)
Descent: 29,000 to 11,000 ft; 5,000 FPM	0	(3)

- (1) 30 co-altitude aircraft in the environment (not Long Beach); density 0.01/nmi²
- (2) 25.3 nmi square giving a density of 0.3/nmi²
- (3) 100 aircraft, density 0.03, 15 TCAS; neither power programming nor error correction decoding employed.

These measurements were carried out using a Cessna 421 aircraft in the Boston area. Two flights were conducted, one for Mode C interrogations, the other for Mode S interrogations. The Mode C interrogations consisted of two whisper-shout suppression pulses followed by two interrogation pulses (S1, S2, P1, P3) transmitted omnidirectionally at 250 watts total radiated power. In each flight, the measurements began at takeoff, after which the aircraft climbed to 12,000 ft. altitude, then proceeded toward the ocean, continued for a period over the ocean, while descending, and then returned and landed. At the time of the flights (2 March 1982) there was snow cover over a portion of the route.

Results. The results of these airborne measurements are shown in Figs. 5-9 and 5-10. Shaded regions in these figures indicate the time periods during which significant receptions were evident.

Certain patterns in the data are recognizable. For example, at the time of crossing from land to ocean in the Mode C flight, there appears to be an echo of the transmitted P3 pulse, received at a time 25 μ s after the P3

MODE C INTERROGATIONS




Fig. 5-9. Interrogation backscatter measurements (Mode C interrogation.)

MODE S INTERROGATIONS




Fig. 5-10. Interrogation backscatter measurements (Mode S interrogation.)

transmissions. This agrees with the calculated delay time for an echo reflecting from the ocean surface directly beneath the aircraft. The fact that this echo was seen for a bottom antenna interrogation and not for a top antenna interrogation is not surprising, and the fact that the P3 echo is evident over water but not over land also is reasonable.

The region of significant reception did not extend beyond about 100 μ s following the start of interrogation. This was true throughout the flights: at all altitudes and over ocean as well as land. Considering top and bottom antennas separately, and considering Mode C and Mode S separately, the resulting limits of backscatter duration were:

- 50 μ s -- top antenna, Mode C
- 60 μ s -- bottom antenna, Mode C
- 70 μ s -- top antenna, Mode S
- 90 μ s -- bottom antenna, Mode S.

In view of the wide range of altitudes and surface reflection conditions experienced in these flights, it seems unlikely that the extent of backscatter will exceed these values in operational use of TCAS II. Thus, these values of transponder suppression duration were adopted in the TCAS II baseline design...

These time periods are much less than the 200 μ s time period used in BCAS. They are small enough so that they easily satisfy the self-suppression limiting constraint. Thus it is not necessary to pursue the possibility of modifying the transponder's interrogation decoder (Sec. 2.5).

REFERENCES

1. J.D. Welch, V.A. Orlando, "Traffic Alert and Collision Avoidance System (TCAS): A Functional Overview of Minimum TCAS II," Project Report ATC-119, Lincoln Laboratory, M.I.T (8 April 1983), FAA-RD-DOT-FAAPM-83/10.
2. J.D. Welch, V.A. Orlando, "Active Beacon Collision Avoidance System (BCAS) Functional Overview," Project Report ATC-102, Lincoln Laboratory, M.I.T. (17 December 1980), FAA-RD-80-127.
3. W.H. Harman, et al, "Active BCAS: Design and Validation of the Surveillance Subsystem," Project Report ATC-103, Lincoln Laboratory, M.I.T. (17 December 1980), FAA-RD-80-134.
4. W.H. Harman, "Effects of RF Power Deviations on BCAS Link Reliability," Project Report ATC-76, Lincoln Laboratory, M.I.T. (7 June 1977), FAA-RD-77-78, Lincoln Laboratory, M.I.T. (7 June 1977), FAA-RD-77-78.
5. W.H. Harman, "Air Traffic Density and Distribution Measurements," Project Report ATC-80, Lincoln Laboratory, M.I.T. (3 May 1979), FAA-RD-78-45.
6. N. Broste, et al, "Preliminary Evaluation of Active Beacon Collision Avoidance System Performance (Simulated): Protection and Alarms," Technical Report MTR-79-W00135, MITRE Corp. (April 1979).
7. P.H. Mann "Simulation of Surveillance Processing Algorithm Proposed for the DABS Mode of BCAS," Project Report ATC-82, Lincoln Laboratory, M.I.T. (16 February 1978), FAA-RD-77-138.
8. A. Wald, Sequential Analysis (Wiley, New York, 1947).
9. "User's Manual for the Los Angeles Basin Standard Traffic Model," MITRE Corp., FAA-RD-73-88.
10. T. Keech, G. Patrick, "Impact of Minimum TCAS II on ATCRBS and Mode S in the Los Angeles Basin," Project Report, ECAC (to be published).

APPENDIX A


MEASURED RELIABILITY OF REPLY DECODING

Bench tests were undertaken to determine the reliability of Mode C reply detection and decoding when overlapping replies are received. These tests, performed on the Lincoln Laboratory TEU, were intended to provide a basis for standards against which other reply processor equipment can be compared.

The TEU was supplied with an input of three replies overlapping in time by various amounts. The replies were input at RF, and were non-coherent. In each of 22 tests, the amounts of reply overlap were varied systematically in the manner shown in Fig. A-1. In different tests, different combinations of reply code, reply carrier frequency, and received reply power level were used, as listed in Table A-1. Note that in tests 1 through 6, the reply codes (6020, 4030, and 4420) contain three information pulses each. In the remaining sixteen tests, the reply codes (6520, 4760, and 6730) contains 5, 6 and 7 information pulses respectively, which may be expected to cause more severe reply garbling.


Each test consisted of a large number of trials. The data from each test was analyzed to determine the percentage of trials in which reply A was detected at the correct range and also the percentage of trials in which reply A was detected at the correct range and correctly decoded. These same percentages were also determined for reply B and reply C. The results are given in Table A-2.

OVERLAP TIMING, TYPE X



NOTES: Time step (1.45/24) μs. 26 trials at each position.
Total number of trials = 19474.

OVERLAP TIMING, TYPE Y



NOTES: Time step = (1.45/24) μs. 26 trials at each position.
Total number of trials = 18746.

OVERLAP TIMING, TYPE Z

Same as Y except $T_1 = 21.025 \mu s$, $T_2 = 22.475 \mu s$.

Fig. A-1. Reply overlap scenarios tested.

TABLE A-1
REPLY PROCESSOR TEST CONDITIONS

Test No.	Overlap Timing (Fig. A-1)	Reply A			Reply B			Reply C		
		Code (ABCD)	RF (MHz)	Power (dBm)	Code (ABCD)	RF (MHz)	Power (dBm)	Code (ABCD)	RF (MHz)	Power (dBm)
1	X	6020	1090	-60	4030	1093	-63	4420	1090	-57
2	X	6020	1090	-60	4030	1087	-63	4420	1090	-60
3	X	6020	1090	-60	4030	1093	-63	4420	1087	-57
4	X	6020	1090	-60	4030	1093	-63	4420	1087	-60
5	X	6020	1090	-60	4030	1087	-63	4420	1093	-57
6	X	6020	1090	-60	4030	1087	-63	4420	1093	-60
7	Y	6520	1090	-60	4760	1090	-60	6730	1090	-60
8	Z	6520	1090	-60	4760	1090	-60	6730	1090	-60
9	Y	6520	1090	-60	4760	1093	-60	6730	1090	-63
10	Z	6520	1090	-60	4760	1093	-60	6730	1090	-63
11	Y	6520	1090	-60	4760	1087	-60	6730	1090	-57
12	Z	6520	1090	-60	4760	1087	-60	6730	1090	-57
13	Y	6520	1090	-60	4760	1093	-57	6730	1087	-63
14	Z	6520	1090	-60	4760	1093	-57	6730	1087	-63
15	Y	6520	1090	-60	4760	1087	-57	6730	1093	-63
16	Z	6520	1090	-60	4760	1087	-57	6730	1093	-63
17	Y	6520	1090	-60	4760	1090	-63	6730	1087	-60
18	Z	6520	1090	-60	4760	1090	-63	6730	1087	-60
19	Y	6520	1090	-60	4760	1090	-57	6730	1093	-60
20	Z	6520	1090	-60	4760	1090	-57	6730	1093	-60
21	Y	6520	1090	-60	4760	1093	-63	6730	1087	-57
22	Z	6520	1090	-60	4760	1093	-63	6730	1087	-57

TABLE A-2
TEST RESULTS

Test No.	Reply A		Reply B		Reply C	
	Detection	Decoding	Detection	Decoding	Detection	Decoding
1	97	65	76	41	99	67
2	98	71	94	62	98	63
3	97	65	80	42	98	65
4	98	65	87	48	98	65
5	98	71	89	53	99	63
6	99	72	94	61	99	63
7	89	62	64	49	86	60
8	96	60	93	48	93	60
9	90	62	66	49	86	59
10	97	61	95	50	93	59
11	89	61	64	48	88	63
12	96	60	94	47	96	64
13	89	62	65	50	86	58
14	97	62	95	51	93	59
15	89	62	65	50	86	59
16	97	62	95	50	93	60
17	90	62	63	48	89	64
18	97	61	92	44	96	64
19	89	61	64	49	87	61
20	96	59	94	47	95	63
21	90	63	64	48	89	63
22	97	61	94	45	96	64