

© 1983 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

A TRAFFIC ALERT AND COLLISION AVOIDANCE SYSTEM FOR GENERAL AVIATION

Dr. Vincent A. Orlando, Leader
Dr. Jerry D. Welch, Assistant Leader
Dr. William H. Harman, Staff Member

Systems Design and Evaluation Group
M.I.T. Lincoln Laboratory
Lexington, Massachusetts 02173

Abstract

One component of the Federal Aviation Administration approach to independent aircraft separation assurance is known as the Traffic Alert and Collision Avoidance System I (TCAS I), which employs passive or active techniques for the detection of nearby transponder-equipped aircraft. This paper gives the results of a study conducted by Lincoln Laboratory of simple techniques for the passive and active detection of transponders. Filter criteria that may be used to restrict passive detections to potentially threatening aircraft are described and evaluated. These techniques and criteria were used in a candidate passive detector whose performance was evaluated in flight against targets of opportunity. A candidate low-power active interrogator was also evaluated through link calculations and airborne measurements. The results indicate that a low-power active interrogator can provide more reliable detection of nearby aircraft and a lower false alert rate than any of the simple passive techniques considered. The active technique generates insignificant levels of interference and, unlike a passive system, also provides protection in regions where there are no ground interrogators.

Introduction

TCAS Concept

In recent years the development of airborne collision avoidance systems has focused on concepts that make use of the transponders carried for ground air traffic service purposes and hence do not impose the need for special avionics on board the detected aircraft. Such systems have the advantage that they can provide immediate protection against collisions involving a significant and growing fraction of the aircraft population.

A system based on this technique is known as the Traffic Alert and Collision Avoidance System (TCAS-I). TCAS, like its predecessor BCAS (Beacon Collision Avoidance System [1]), is designed to provide protection against aircraft equipped with both the existing SSR and future SSR Mode S transponders. The fundamental purpose of TCAS is to provide a separation assurance capability that is able to operate in all airspace without reliance on ground equipment. The TCAS concept encompasses a range of capabilities that includes (a) TCAS I, a low cost, limited performance version, and (b)

TCAS II, which is intended to provide a comprehensive level of separation assurance in all current and predicted airspace environments through the end of this century.

TCAS I

TCAS I has the capability of detecting transmissions from nearby transponders and advising the pilot when the characteristics of any transmission indicate that it might be a threat. The replies detected may have been elicited by ground station interrogations or by spontaneous transmissions of Mode S transponders (passive TCAS I) or may have resulted from low power interrogations transmitted by the TCAS I equipment itself (active TCAS I).

This paper focuses on suitable techniques for detecting nearby transponders while generating such low levels of radio frequency interference that unrestricted implementation could be permitted with no undesirable interference effects to the current or future SSR. It provides a comparison of simple passive and active detection techniques.

Passive Detection

The principal problem with passive detection is control of false alarms. As TCAS I operates in higher traffic densities its effectiveness will be reduced if it alarms frequently. Thus, some means is needed to filter or restrict the triggering of pilot advisories so they occur only on transmissions received from potentially threatening aircraft, that is, aircraft that are close in both range and altitude. There are only a limited number of characteristics of a passively received reply that can be used as simple filter criteria. The most useful appeared to be:

1. Received power: Received power can be used two ways: a) the received power can be compared to a fixed threshold to reject transmissions from aircraft at long range, b) power may be tracked to determine how range is changing as a function of time.
2. Aircraft altitude: Transmissions from off-altitude aircraft may be rejected two ways: a) the inherent off-altitude rejection provided by the aircraft antenna patterns may be used, b) the altitude code may be detected and compared with own altitude.

3. Time-after-interrogation: If an aircraft is close and in the same ground interrogator beam as the TCAS I aircraft, range information may be inferred by comparing the time-of-arrival (at the TCAS I aircraft) of its transponder reply with the TCAS I transponder reply time.

Thus there are five distinct simple techniques for detection filtering based on these three characteristics. Each of these techniques is described in this section and an indication of expected performance is given.

Received Power Thresholding

The purpose of power thresholding is to distinguish between aircraft that are within a given volume of local airspace and those that are outside of this volume. Unlike the active mode of aircraft detection, in which replies from distant aircraft can be eliminated on the basis of time delays (i.e., range), passive mode detection does not have a direct measure of detection range.

The use of a power threshold criteria filter is complicated by the large variance in transponder reply power and transponder antenna gains observed in actual aircraft installations. The variation of the detected power from a population of general aviation aircraft, all at the same range, has been found to be more than 20 dB [2]. One consequence of the large variation in received power from transponder to transponder is that when the threshold is set to detect most aircraft at a nominal close range, some aircraft will still be detected at long ranges. Table 1 summarizes this effect, showing calculated detection performance for a nominal sensitivity setting of -57 dBm based on data from Ref. 2. It tabulates the range for a given detection reliability for the two types of targets. The detection range is greater for air carrier targets because their transponders are, on average, more powerful.

TABLE 1
RANGE PERFORMANCE (CALCULATED) FOR
RECEIVED POWER LEVEL THRESHOLDING TECHNIQUE

Range For A Given Detection Reliability			
Target Type	Detection Reliability		
	90%	50%	10%
Gen Aviation	1.5 nmi	2.8 nmi	5.5 nmi
Air Carrier	3.3 nmi	5.9 nmi	10.5 nmi

Table 2 shows the maximum closing speeds that could be handled while providing a 30-sec warning. The resulting closing speeds at the 50%-reliability range are 340 kt for general aviation and 710 kt for air carrier targets. These are about the

highest closing speeds a GA aircraft would expect to encounter. However, the general aviation closing speed handled at the 90%-reliability-range of 1.5 mile is only 180 kt. Thus, some of the targets will not be detected early enough to provide a 30-second warning.

TABLE 2
ENCOUNTER SPEED PERFORMANCE (CALCULATED) FOR
RECEIVED POWER LEVEL THRESHOLDING TECHNIQUE

Maximum Encounter Speed for 30-Second Warning			
Target Type	Detection Reliability		
	90%	50%	10%
Gen Aviation	180 Kt	340 Kt	660 Kt
Air Carrier	400 Kt	710 Kt	1260 Kt

Received Power Level Tracking

While the large variance in the received power level of a population of transponders makes it difficult to determine range based on absolute power measurements, one can also measure the power variation observed versus time from a single transponder to attempt to identify transmissions received from approaching aircraft and to reject those received from departing aircraft. If all other link factors are constant, an increase in received power of 6 dB over a time T means that the range to the detected aircraft has decreased to one half its original value and that the range will become zero in the next interval of T seconds if the radial speed remains constant. This indicates that Tau (the time to closest approach) can be expressed as a function of differential received power and measurement time. An equation for Tau as a function of differential received power (ΔP) observed over a time (Δt) is shown in Table 3. Values of Tau in seconds for several values of ΔP and Δt are also shown.

In order to evaluate the accuracy of Tau estimation based on power tracking, an analysis was performed on air-to-air surveillance data for seven planned encounters. During these encounters the threat aircraft was actively interrogated at a rate that permitted range to be measured as a function of time so that true Tau could be calculated.

While there was significant scattering of the power tracking Tau estimate compared to the true Tau, two conclusions were noted:

1. The estimate of covering/diverging status was correct most of the time.
2. A small value of estimated Tau was usually an indication of a true threat condition, i.e., true Tau less than 30 seconds.

TABLE 3.

TAU (τ) DERIVED FROM POWER TRACKING

$$\tau = \frac{\Delta t}{1 - 10^{(\Delta P/20)}}$$

Power Diff, ΔP (dB)	Time Difference, Δt (Sec)			
	1	4	6	10
6	1	4	6	10
4	2	7	10	17
2	4	15	23	39
0	"	"	"	"
-2	-5	-19	-29	-49
-4	-3	-11	-16	-27
-6	-2	-8	-12	-20

These observations are illustrated in Table 4 where the sample measurements are categorized by average power change over a six-second time interval. For example, the first row indicates that of 12 cases where a +6 dB increase in average power was measured in one six-second interval, 10 of the cases occurred where the true Tau was < 30 sec and 2 occurred when the threat aircraft was diverging.

It should be noted that power tracking requires reply-to-reply correlation. This correlation is easy for Mode S replies because of the unique address code. It is somewhat difficult for Mode C or discrete code Mode A replies and very unreliable for other ATCRBS cases when there are enough aircraft present to result in a finite probability that two or more targets have the same code.

Altitude Discrimination by Antenna Pattern Filtering

Measurements of typical transponder antenna patterns [3, 4, and 5] indicate that if both the transponder antenna and the passive transponder detector antenna are bottom mounted, the vertical coverage of the power thresholding technique described above is restricted by antenna patterns and air frame blockage to roughly ± 5000 ft if the target is a GA aircraft, and ± 12000 ft if it is an air carrier aircraft as shown in Fig. 1. Air-to-air measurements demonstrating this effect are shown in Fig. 2.

Altitude Code Filtering

The second way to filter off-altitude targets is to determine the code contained in the detected reply. This detection is more reliable if a top-mounted antenna is used since this improves the protection from code errors due to multipath. Mode A replies

are of no value since they contain no altitude data. Unfortunately, ATCRBS replies are not uniquely labelled as Mode A or C. Some Mode A replies can be rejected by checking the code bits and rejecting those that contain illegal altitude codes. Rejecting those combinations will not eliminate all Mode A replies. Fortunately, all 1200 code replies are illegal. Further, the probability of a discrete Mode A code causing an altitude alert appears small and has not been observed in the data analyzed.

The overlay of a ± 1000 foot altitude code filter on the antenna patterns of Fig. 2 is shown in Fig. 3. It is seen that the technique appears most useful when detecting the higher-power air carrier aircraft, which are all equipped for altitude reporting.

TABLE 4.
POWER TRACKING MEASUREMENT PERFORMANCE
(Based on $\Delta t = 6$ Seconds)

ΔP dB	True Tau			Probability That True Tau < 30 sec	Prob. of Correct Conv/Div. Category
	< 30 sec	> 30 sec	Opposite Sense		
> 6	10	0	2	0.83	0.83
5	1	1	2	0.25	0.50
4	10	3	1	0.71	0.93
3	5	5	0	0.50	1.00
2	4	19	1	0.17	0.96
1	6	24	6	0.12	0.82
0	1	22	0	0.04	-
-1	1	7	7	0.07	0.53
-2	0	11	1	0.0	0.92
-3	0	6	0	0.0	1.00
-4	1	4	1	0.17	0.83
> -5	0	9	0	0.0	1.00
Total Data Points	37	111	21		




Fig. 1 Antenna pattern filtering.




Fig. 2 Antenna pattern filtering results, bottom antenna.




Fig. 3 Altitude code filtering.

Time-After-Interrogation

The principle used in time-after-interrogation filtering is shown in Fig. 4. A TCAS I aircraft (A) and a threat aircraft (B) are both illuminated by the beam of a ground interrogator. The ATCRBS interrogation arrives first at the TCAS I aircraft

and a short time later at the threat aircraft. The reply generated by the threat aircraft is seen to arrive at the TCAS I aircraft in the interval following the TCAS I reply. The fact that the TCAS I aircraft replies to the same interrogation as the threat aircraft limits the closest range from which replies can be received due to what can be called the "ATCRBS blind spot effect". For example, if the aircraft are close together, the reply from the threat aircraft can be received at the TCAS I aircraft while the TCAS I transponder is itself transmitting a reply and blanking the TCAS I receiver.

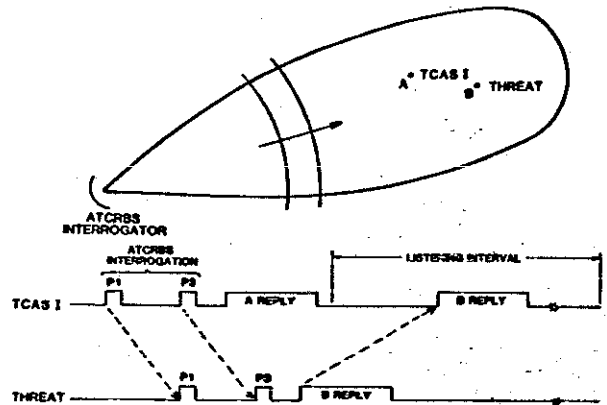


Fig. 4 Time-after-interrogation filtering.

The general geometry for the ATCRBS blind spot effect is shown in Fig. 5. Two envelopes are shown. The outer is the blind spot envelope for which the threat reply would overlap some portion of the TCAS I reply, the inner envelope is for the clear detection of only the F2 pulse of the threat reply. It is obvious that a pulse detection approach must be used if the blind spot envelope is




Fig. 5 Time-after-interrogation blind spot geometry.

to be kept small enough to allow detection of aircraft within 2 miles. The outer envelope also gives an indication of how the acceptance volume changes relative to the location of the ATCRBS interrogator. If the listening window is set to accept pulses from aircraft up to 2 nmi farther away from the interrogator than the TCAS I aircraft, the acceptance volume increases as the threat range decreases with respect to the ATCRBS interrogator.

With one ATCRBS interrogator, passive detection provides a useful reduction in acceptance volume compared to the power thresholding technique. When a second interrogator is considered, the effectiveness of passive detection in reducing acceptance volume decreases because the resultant acceptance volume is the union of the acceptance volumes for each interrogator. With more than 3 or 4 interrogators, the time-after-interrogation filter appears to provide very little additional filtering compared to the power thresholding technique. An example of this multi-interrogator effect is shown in Fig. 6, which compares the alert rate measured at 8500 feet in the Boston area using only power thresholding with the rate measured when using power thresholding and time-after-interrogation filtering. For these measurements the time-after-interrogation acceptance window was set at 2 nmi. The figure shows nearly equal alert rates for the two techniques.




Fig. 6 Time-after-interrogation filtering performance - Boston area.

A serious false alarm mechanism for the technique occurs for ATCRBS targets at altitudes up to around 5000 ft. At these altitudes, backscatter multipath from the TCAS I transponder's reply may have sufficient amplitude to be detected in the listening window. An example of the effect is shown in Fig. 7. A substantial number of pulse detections above the -57 dBm threshold are seen to occur up to 4500 feet with some pulses still detected at 5000 feet.

Evaluation of Filter Criteria

Two of the techniques studied appear unsuitable due to difficulties in handling ATCRBS replies:

Power Tracking - Reply correlation is needed to support power tracking. This correlation becomes very unreliable for non-Mode C and non-discrete-replies in higher density airspace where filtering is needed most.

Time-After-Interrogation - At low altitude, ATCRBS detections will be unreliable due to backscatter multipath. At high altitude the performance will be reduced due to the fact that the TCAS I will become visible to more interrogators. The filter will be least effective in an area with a high interrogator density, where filtering is needed most.

The three remaining filter criteria (power level thresholding, antenna pattern and altitude code filtering) formed the basis for measurements of passive detector performance using the following techniques:

1. Continuous listening except during own transponder replies.
2. Received power thresholding with a nominal sensitivity of -57 dBm.
3. Altitude code filtering to: (1) reject replies outside a nominal ± 1000 foot band, (2) reject replies with invalid altitude codes, and (3) accept replies with empty Mode C brackets.

It does not appear feasible to use both antenna pattern and altitude code filtering in a passive detector with a single antenna because the bottom-mounted antenna location required for antenna pattern filtering will lead to frequent multipath-induced bit errors in the detected Mode C code.

Experience with passive detection has indicated a high alarm rate if an alert is triggered on every accepted reply. False brackets are frequently synthesized by pulses of closely-spaced ATCRBS replies. It is therefore necessary to set a minimum threshold on the number of replies per second that must be received to trigger an alert. Since a terminal sensor elicits approximately 12-16 replies per beam dwell and may use an interrogation sequence of Mode A/Mode A/Mode C, the highest fixed threshold that can be used when there is only a single interrogator is 4 Mode C replies in a one-second interval.




Fig. 7 1090 MHz reply backscatter.

Once an alert has been triggered, the alert should stay on for 5 seconds to provide the pilot the opportunity to observe it, and to avoid continuous retriggering in the case where the replies are being elicited by a single terminal interrogator and thus are only received once per scan.

Passive Detection Performance Measurements

Two sets of in-flight performance measurements on targets-of-opportunity were analyzed in order to quantify the performance of the passive detector.

1. Active/Passive - The equipment was configured to interleave active TCAS II interrogations with passive listening once per second. The active data provided the true target information needed to evaluate passive detection acquisition range, warning time and false alarm probability.

2. Passive Only - The equipment was configured to operate as a real-time TCAS I equipment and measure alert rates on two flights from Boston to Washington.

Active/Passive Measurements

Data on targets-of-opportunity were collected at 8500 feet in the Boston area. The following results are based on an analysis of one-hour and twenty minutes of flight data.

Performance Results

The data yielded approximately 2000 aircraft-seconds of data on 35 different aircraft. To increase the sample size, calculations of acquisition range and warning time were performed on the total set of aircraft regardless of the results of altitude filtering. The following performance measurements were calculated from this set of data.

Acquisition Range - The range at which the passive reply count initially exceeded four replies/second was determined for each of the 35 acquired aircraft. The results are plotted as an acquisition range histogram in Fig. 8. Only two of the aircraft in the sample were non-Mode C equipped and are presumably general aviation aircraft. Note that these two aircraft were detected at close range (as predicted by the link calculations).

Warning Time - The time from initial acquisition until the time of closest approach was noted for the 10 aircraft in the sample whose minimum range was 3 nmi or less, since this would be the subset of most immediate interest to the pilot of the TCAS I aircraft. The results are presented in Table 5.

Probability of Surveillance False Alarm - Alerts due to Mode C detections were very reliable. Only 5% of the alert time could not be correlated with active traffic measurements. A much higher false alarm rate was noted for non-Mode C alerts. With a threshold of 4 replies/second, 53% of the alert time caused by non-Mode C detections could not be correlated with traffic detected by active measurement.




Fig. 8 Power thresholding acquisition performance.

TABLE 5.

PASSIVE DETECTION WARNING TIME.

(TARGET WITH MINIMUM RANGE \leq 3 NMI)

ACQUISITION RANGE (NMI)	MINIMUM RANGE (NMI)	WARNING TIME* (SEC)
5.5	1.5	30
2.7	2.5	5
3.0	3.0	0
2.8	2.8	0
3.2	3.0	5
7.5	0.7	125
0.5	0.5	0
2.0	1.6	20
1.8	1.1	20
2.8	2.6	20

*Time of Alert Prior to Time of Closest Approach

Alert Rate - Figure 9 shows the alert rate performance for the 80 minute flight in terms of the percent of time the alert was "on" for each of 10, 8-minute intervals. Results are shown with and without altitude code filtering and demonstrate the effectiveness of code filtering in reducing alerts in environments with high Mode C equipage.

Passive Only Measurements with Bottom Antenna

Passive data on targets-of-opportunity were conducted on flights from Boston to Washington, where: (1) a Mode C acceptance band of ± 1500 feet was used, and (2) the antenna was bottom-mounted. Results for one flight are shown in Fig. 10. Note the high alert rates over New York and on descent into Washington National Airport.




Fig. 9 Altitude code filtering performance - Boston area.




Fig. 10 Passive detector alert rate: Boston to Washington at 8,500 ft.

Low Power Active Detector

An interference analysis was conducted to explore the possibility of using an active proximity detector for TCAS I. The purpose of the analysis was to determine a power level that could be used by TCAS I aircraft in the highest density environments and still cause no significant interference effects on the environment. The calculation (which was based upon the Los Angeles high density model [6]), assumed that one-half of all aircraft were active TCAS I equipped, and allowed the total interference effect of TCAS I operation to be 10% of the interference caused by TCAS II operation. The results of the analysis indicate that an active TCAS I using a time-power product equivalent to one 5-watt Mode C interrogation per second has negligible impact on the SSR system. This is equivalent to a 10-watt Mode C interrogation every two seconds, etc. [7].

Calculated Performance

A link analysis was performed for a low power TCAS I interrogator in order to estimate the possible utility of this technique. This analysis also includes the performance of a 4-watt Mode C interrogation once per second since measured data at that power level were already available.

The calculated performance is shown in Table 6. Performance out to the range of principal interest for visual acquisition (about 2 nmi) is seen to be adequate.

TABLE 6

CALCULATED VALUES OF TRACKING PROBABILITY FOR A LOW POWER TCAS I DETECTOR

Range (nmi)	Interrogator Power (at antenna)			
	4 watts	5 watts	10 watts	20 watts
1	0.90	0.93	0.97	0.99
2	0.67	0.72	0.84	0.93
3	0.47	0.53	0.69	0.83
4	0.33	0.38	0.56	0.72

Measured Performance

Data for a flight from Boston to New York and return were analyzed in order to obtain measurements of performance of a low power active TCAS I in an actual in-flight environment. Attention was focused on the enroute portion of the flight, at 8000 feet southbound and 9000 feet northbound. A total of 70 minutes of flight was examined, which provided data on 16 aircraft targets-of-opportunity. TCAS II surveillance data from both top and bottom antennas were used to establish range/altitude truth. Replies from just the lowest level (4 watt) interrogation from the top antenna were examined to identify the portions of each flight path during which low power interrogations were successful. Figure 11 gives results for one leg of the flight. In Fig. 12, the results for aircraft within the principal threat zone ($\pm 10^\circ$) have been presented in terms of probability of successful detection for each one-nmi range band. Also plotted is the calculated performance. The match between airborne measurements and the calculated performance is good considering the number of tracks observed.

Summary

Passive Detection

Several simple techniques for passive filtering were evaluated. Those that were found to be useful were combined in a candidate passive detector that was evaluated with flight test data. The results show that initial acquisition ranges can vary from





Fig. 11 Active TCAS I performance - Boston to New York.



* PERCENT OF AIRCRAFT FROM WHICH REPLIES ARE ELICITED BY A 4-WATT INTERROGATION, FOR AIRCRAFT WITHIN $\pm 10^\circ$ IN ELEVATION ANGLE

Fig. 12 Active TCAS performance as a function of range.

one to eleven miles and demonstrate the difficulty of the passive detection technique to effectively discriminate aircraft range. Altitude code filtering is an effective technique for reducing the alert rate in environments of high Mode C equipage.

Active Detection

It appears feasible to use a low power interrogator to greatly improve air-to-air surveillance performance. A time-power product equivalent to one 5-watt Mode C interrogation every second is low enough in power and rate that all interference effects resulting from these interrogations are acceptably small.

A limited set of data evaluated for a low-power active interrogator agreed with calculated link performance and showed adequate performance out to about 2 nmi for a 4-watt interrogator. Performance at 2 nmi and beyond could be enhanced by increasing the power and decreasing the interrogation rate. One 20-watt interrogation every four seconds would seem to be a suitable design. The false alarm rate of the active detector should be low, due to the use of range gating and a top-mounted antenna. If Tau calculations are performed, the false alarm performance of the active TCAS I should be acceptable.

References

1. 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.
2. G.V. Colby, E.A. Crocker, "Final Report Transponder Test Program", Project Report ATC-9, Lincoln Laboratory, M.I.T. (12 April 1972) FAA-RD-72-30.
3. G.J. Schlieckert, "An Analysis of Aircraft L-Band Beacon Antenna Patterns", Project Report ATC-37, Lincoln Laboratory, M.I.T. (15 January 1975) FAA-RD-74-144.
4. D.W. Mayweather, "Model Aircraft L-Band Beacon Antenna Pattern Gain Maps", Project Report ATC-44, Lincoln Laboratory, M.I.T. (16 May 1975) FAA-RD-75-75.
5. K.J. Keeping, J.C. Sureau, "Scale Model Pattern Measurements of Aircraft L-Band Beacon Antennas", Project Report ATC-47, Lincoln Laboratory, M.I.T. (4 April 1975) FAA-RD-75-23.
6. "User's Manual for the Los Angeles Basin Standard Traffic Model", MITRE Corp., FAA-RD-73-88, p. 73.
7. V.A. Orlando, et al, "TCAS I Design Evaluation," Project Report ATC-114, Lincoln Laboratory, M.I.T. (24 September 1982) FAA-RD-82-12.

This work was performed under the sponsorship of the Federal Aviation Administration.

The views and conclusions contained in this document are those of the contractor and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the United States Government.