

AD-776 798

FALSE TARGET ELIMINATION AT ALBUQUERQUE  
USING ARTS-III SOFTWARE

Alan G. Cameron

Massachusetts Institute of Technology

Prepared for:

Electronic Systems Division

12 March 1974

DISTRIBUTED BY:

**NTIS**

**National Technical Information Service**  
**U. S. DEPARTMENT OF COMMERCE**  
5285 Port Royal Road, Springfield Va. 22151

DOCUMENT CONTROL DATA - R&D <b>AD-776 798</b>		
<i>(Security classification of title, body of abstract, and indexing annotation must be entered when the overall report is classified)</i>		
1. ORIGINATING ACTIVITY (Corporate author)  Lincoln Laboratory, M. I. T.	2a. REPORT SECURITY CLASSIFICATION Unclassified	2b. GROUP
3. REPORT TITLE  False Target Elimination at Albuquerque Using ARTS-III Software		
4. DESCRIPTIVE NOTES (Type of report and inclusive dates) Technical Note		
5. AUTHOR(S) (Last name, first name, initial)  Cameron, Alan G.		
6. REPORT DATE 12 March 1974	7a. TOTAL NO. OF PAGES 30	7b. NO. OF REFS 1
8a. CONTRACT OR GRANT NO. F19628-73-C-0002  b. PROJECT NO. 649L  c. d.	9a. ORIGINATOR'S REPORT NUMBER(S)  Technical Note 1974-12	9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)  ESD-TR-74-134
10. AVAILABILITY/LIMITATION NOTICES  Approved for public release; distribution unlimited.		
11. SUPPLEMENTARY NOTES  None	12. SPONSORING MILITARY ACTIVITY  Air Force Systems Command, USAF	
13. ABSTRACT  Lincoln Laboratory has been assisting the U.S. Air Force Weapons Laboratory in the study of ATCRBS false targets caused by reflections from buildings at Albuquerque International Airport/Kirtland AFB.  This paper discusses the false target situation there and proposes two algorithms for correcting it in ARTS-III software. The simpler of the two appears applicable to the Albuquerque situation today, and is easily extendable to correct false targets caused by new buildings. Since the process appears directly applicable to many FAA Secondary Radar installations, a more complex algorithm is also presented, which is suitable for use in very high density terminal areas.		
14. KEY WORDS  false targets                      air traffic control                      surveillance systems ARTS-III                              radar reflections		

MASSACHUSETTS INSTITUTE OF TECHNOLOGY  
LINCOLN LABORATORY

FALSE TARGET ELIMINATION AT ALBUQUERQUE  
USING ARTS-III SOFTWARE

A. G. CAMERON

*Group 41*

TECHNICAL NOTE 1974-12


12 MARCH 1974

Approved for public release: distribution unlimited.

LEXINGTON

*ia*

MASSACHUSETTS



The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology, with the support of the Department of the Air Force under Contract F19628-73-C-0002.

This report may be reproduced to satisfy needs of U.S. Government agencies.

## ABSTRACT

Lincoln Laboratory has been assisting the U. S. Air Force Weapons Laboratory in the study of ATRBS false targets caused by reflections from buildings at Albuquerque International Airport/Kirtland AFB.

This paper discusses the false target situation there and proposes two algorithms for correcting it in ARTS-III software.\* The simpler of the two appears applicable to the Albuquerque situation today, and is easily extendable to correct false targets caused by new buildings. Since the process appears directly applicable to many FAA Secondary Radar installations, a more complex algorithm is also presented, which is suitable for use in very high density terminal areas.

Accepted for the Air Force  
Eugene C. Raabe, Lt. Col., USAF  
Chief, ESD Lincoln Laboratory Project Office

---

\*Some of the concepts on which these algorithms are based are also applicable to other ATRBS processors such as that contained in DABS. Many of the details of the false target algorithms designed for the ATRBS mode of DABS are different, however.

## TABLE OF CONTENTS

<u>Section</u>		<u>Page</u>
	ABSTRACT . . . . .	iii
1	INTRODUCTION . . . . .	1
2	BACKGROUND . . . . .	2
3	THE FALSE TARGET MECHANISM . . . . .	4
4	THE BASIC APPROACH . . . . .	11
5	A MORE SOPHISTICATED APPROACH . . . . .	18
6	RECOMMENDATIONS AND CONCLUSIONS . . . . .	24

## LIST OF ILLUSTRATIONS

<u>Figure</u>		<u>Page</u>
1	False target geometries . . . . .	5
2	Typical link parameters . . . . .	6
3	Reflection geometry general case . . . . .	8
4	Simplified geometry $R_A \gg R_R$ . . . . .	10
5	Close-in reflection geometry at Albuquerque . . . . .	12
6	Simple false target identification algorithm . . . . .	13
7	False target elimination algorithm employing tracking . . . . .	19

SECTION 1  
INTRODUCTION

Lincoln Laboratory is presently assisting the Air Force Weapons Laboratory in the development of a software false target elimination process for ARTS-III. The software process will be installed and tested initially on the FAA-operated ARTS-III at Albuquerque; several studies of ARTS-III and other data from many installations suggest that its use throughout the FAA secondary surveillance system could result in widespread performance improvements.

During the course of a study\* by Lincoln Laboratory of the overall performance of the ATC Radar Beacon system, as determined by analysis of ARTS-III extractor data, it became apparent that the highly deterministic nature of the false-target-producing mechanism made it amenable to software correction. The report on that analysis proposed such a fix; since that time, several variations on the basic concept have been developed in connection with analyses performed on data from other sites. This paper reviews the basic concepts behind software elimination of false targets, and applies them specifically to the situation at Albuquerque.

---

\*A. G. Cameron and D. H. Pruslin, "Empirical Assessment of ATCRBS," Project Report ATC-16, Lincoln Laboratory, M. I. T. (31 October 1973).

## SECTION 2

### BACKGROUND

The problem of false targets due to reflections from buildings and obstructions has been recognized as severe since the FAA secondary surveillance radar system was first implemented. At present, it and the problem of weak targets are considered by most controllers to be the most severe limitations on the operational performance of ATCRBS. Although their relationship is not apparent at first, the two problems are closely coupled. Present runlength thresholds are set high in order to discriminate against false targets; in that process they unfortunately also discriminate against weak targets. An effective solution to the false target problem would improve weak target performance, since it would allow runlength thresholds to be substantially reduced.

Due to the recognized severe operational impact of the false target problem, especially in a semi-automated surveillance and control system such as NAS, much attention has been devoted attempts to correct it. The improved interrogation sidelobe suppression technique ( $I^2SLS$ ) was developed and implemented in most FAA sites to reduce false target levels by inhibiting all aircraft not in the mainbeam with a suppression transmission ( $P_1$  and  $P_2$ ). This would prevent replies to reflected interrogations following close behind the omnidirectionally transmitted suppression pulses. The effectiveness of

this technique in the field is not well-known. For example, it does not appear to eliminate many false targets at Boston, Milwaukee, Andrews AFB and other sites; on the other hand, Improved SLS appears to be operating quite effectively at Albuquerque. These differences have yet to be resolved.

Software procedures for identifying and rejecting false targets have been proposed by MITRE, UNIVAC, Lincoln Laboratory, and others. These have been generally simple (e. g., whenever the same discrete-code appears twice in one scan, drop the one at greater range), but have not been implemented since a) they are not generally effective against nondiscrete code targets, and b) they all leave a small but significant likelihood that a legitimate target will be mistakenly identified as false and inadvertently dropped. For a false target elimination process to be successful, the likelihood of such improper identification and subsequent rejection must be extremely low per scan, and far lower still over a sequence of scans. The process proposed here appears to satisfy that requirement.

A tradeoff must be addressed in this area, similar to that considered in radar threshold selection. At one extreme, the system can be designed to assume that doubtful targets are real, and thus not identify all the false targets. At the other extreme, the system might assume them to be false, thus occasionally erroneously tagging actual traffic as false targets. Fortunately, it appears that the already low likelihood of either situation can be further reduced by taking advantage of the scan-to-scan pattern of false targets to increase the confidence with which the decisions are made.

## SECTION 3

### THE FALSE TARGET MECHANISM

Extensive empirical analyses have confirmed that the reflection process can be characterized quite precisely by a few simple laws of geometrical optics. Three types of reflection geometries have been observed (Fig. 1); the most common type (Fig. 1a), in which both interrogation and reply follow the same reflecting path, has been seen throughout the data, usually involving several reflectors at each site. The ability of a surface to sustain a reflection process of this sort depends on its orientation and on whether the reflected signal is of sufficient strength to trigger the transponder on the uplink and be detected on the downlink. Since both links are generally quite highly overpowered to compensate for antenna-null induced fading, even reflected signals which are highly attenuated by the reflection process can cause false targets. Attenuation in the reflection process results both from geometrical considerations (small reflectors at great distances do not reradiate much signal power), and from the typically low reflectivities of the reflector materials. As in the conventional radar equation, when the reflector is small compared to a Fresnel zone (the usual case), the additional free space pathloss encountered in the path from the interrogator to the reflector is quite large, and is not completely compensated for by the reflector aperture and gain. Figure 2 compares the signal levels which might be encountered on typical reflecting paths with





Fig. 1. False target geometries.




(a) REFLECTOR SMALL  
(compared to a Fresnel zone)

UPLINK

ERP, ...	49.5 dBW
pathloss, $r_1$	(104) dB
reflector gain, $r_2$	77 dB
pathloss, $r_2$	(123) dB
	<u>-70.5 dBm</u>

WELL ABOVE MTL




(b) REFLECTOR LARGE  
(compared to a Fresnel zone)

UPLINK

ERP, ...	49.5 dBW
pathloss	(125) dB
reflection losses	(6) dB
	<u>-51.5 dBm</u>

WELL ABOVE MTL



(c) NO REFLECTOR  
(direct path = 22 mi.)

UPLINK

ERP, ...	49.5 dBW
pathloss (125) dB	
	<u>-45.5 dBm</u>

WELL ABOVE MTL

Note: Figures idealized to neglect fading due to small-time-delay multipath

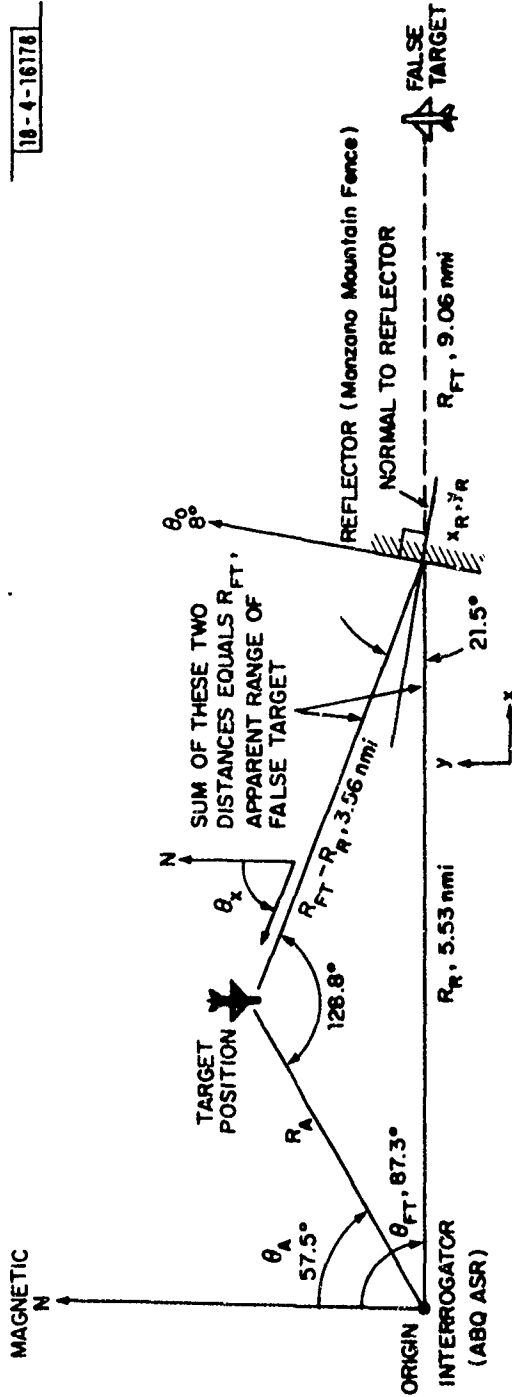
Fig. 2. Typical link parameters.

direct path signal levels. We have noted that in practice the geometrically induced limitations appear to dominate; wooden, glass, and concrete buildings are quite common reflectors. However, small distant buildings rarely support the reflection process regardless of how they are constructed.

In the general case (Fig. 3) in which the positions of both the aircraft and the reflector are arbitrary, locating the reflector involves the general solution of an ellipse. This solution is shown in Fig. 3; the parameters used in the numerical example there were taken from an actual false target declaration in the Albuquerque data caused by the Manzano Mountain fence.

The inverse process, determining the location of a potential false target when given the reflector parameters and actual target position is somewhat complicated. As the expressions in Fig. 3 show, transformation from polar to rectangular coordinates is required, involving solution of several trigonometric functions. Fortunately, this process is rarely necessary; usually reflector/aircraft geometry is such that the solution simplifies considerably. Situations such as the one in Fig. 3 arise only when the reflector is at considerable distance from the interrogator (comparable to that of the aircraft). In order to support the reflection process at that range, the reflector must be quite large; the Manzano Mountain fence is one of very few distant reflectors seen in ARTS-III data that is of sufficient size (for its range) to cause problems. Its orientation, such that it illuminates a frequently traveled flight path, results in a high incidence of false targets.

Since reflection processes such as this one are atypical, and lead to considerable complexity, it seems proper at this initial stage of study to concentrate on the more commonly seen situation in which the geometry allows considerably simpler solutions.



TO DETERMINE REFLECTOR LOCATION

- INTERROGATOR AND ACTUAL AIRCRAFT COMPRISE FOCI OF AN ELLIPSE; REFLECTOR LIES ON PERIPHERY
- $R_A = 2ae$ ,  $R_{FT} = 2a$ ;  $e = R_A / R_{FT}$   
(recall method of constructing an ellipse using two pins and a piece of string)
- HERE,  $a = \text{SEMIMAJOR AXIS} = 4.53 \text{ nmi}$   
 $e = \text{ECCENTRICITY} = 0.286$
- $R_R = \frac{a(1-e^2)}{1-e \cos[\theta_{FT} - \theta_A]}$
- REMAINING SIDES AND ANGLES FOUND BY SUBTRACTION AND:  
 $\frac{a}{\sin A} = \frac{b}{\sin B} = \frac{c}{\sin C}$

TO DETERMINE AIRCRAFT LOCATION

- CONVERT TO RECTANGULAR COORDINATES
- $x_R = R_R \sin \theta_{FT}$      $x_A = R_A \sin \theta_A$
- $y_R = R_R \cos \theta_{FT}$      $y_A = R_A \cos \theta_A$
- $x_A = R_R \sin \theta_{FT} + (R_{FT} - R_R) \sin \theta_x$
- $y_A = R_R \cos \theta_{FT} + (R_{FT} - R_R) \cos \theta_x$   
WHERE  $\theta_x = 2\theta_0 - \theta_{FT}$
- KNOWN QUANTITIES ARE  $R_R$ ,  $\{\theta_{FT}\}$ ,  $R_A$ ,  $\theta_A$ ,  $\theta_0$
- INVERT (or iterate) TO FIND  $R_{FT}$ ,  $\theta_{FT}$

Fig. 3. Reflection geometry (general case).

When the reflector is close to the interrogator, \* such that the paths from both to the aircraft are essentially parallel, the situation shown in Fig. 4 results. Here, false target range and azimuth are related to those of the actual aircraft in a simple and straightforward manner. That this situation is by far the most frequently observed in practice allows software implementation of the false target elimination process to be accomplished with a minimum of computation. The simplified equations of Fig. 4 were used to locate the two strong reflectors nearby the interrogator at Albuquerque; the numerical example in that figure is one of the false target instances resulting from the fence to the north of west.

---

\*The same situation arises when the reflector is in the immediate vicinity of the aircraft; this is rarely seen in practice.

18-4-16179

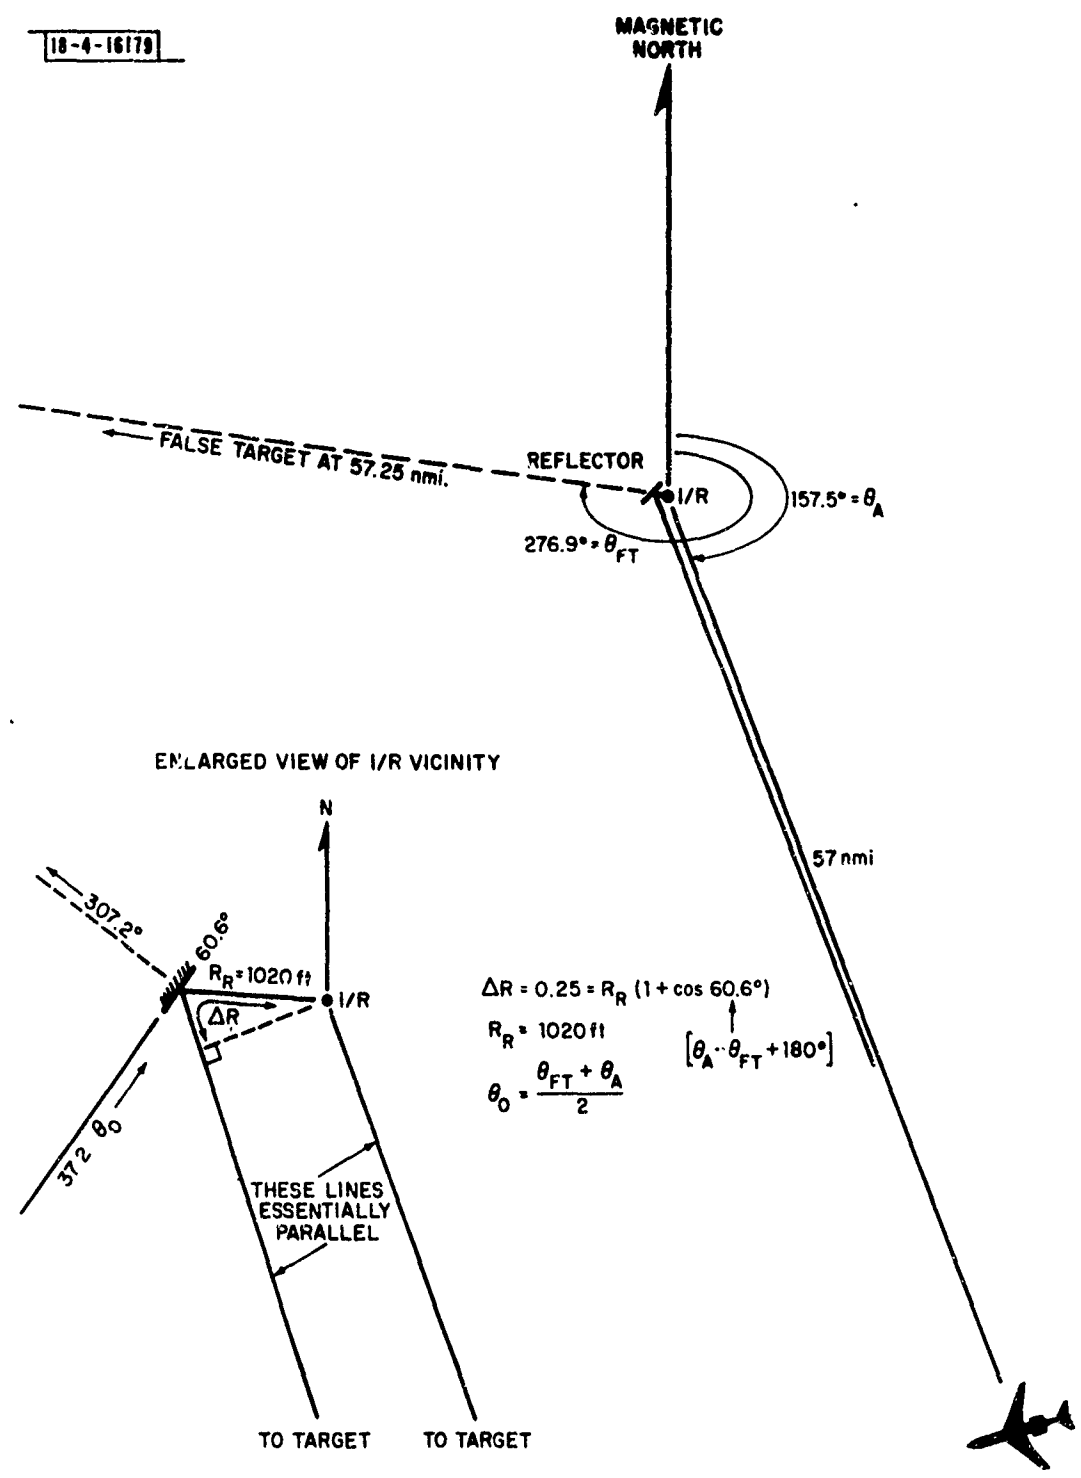


Fig. 4. Simplified geometry ( $R_A \gg R_R$ ).

## SECTION 4

### THE BASIC APPROACH

This section describes the basic false target elimination algorithm proposed for Albuquerque, first in general terms, and then specifically tailored to the Albuquerque situation. The process involves recognition of aircraft in regions where they can cause false targets, calculation of where those targets should be, search of those areas to see if correlated targets are present, and identification of those targets as false.

The first step in the process is to identify all aircraft which are in the regions illuminated by the reflectors, and which could thereby produce false targets.

The illuminated regions are simply defined as azimuthal wedges (Fig. 5). Whenever a target declaration azimuth falls within one of these wedges, its range, azimuth, identification, and altitude are stored for further processing (Fig. 6), along with the parameters of the particular reflector,  $\theta_0$  and  $\Delta R$ .

These parameters are defined in Fig. 4, and allow calculation of the position at which a false target would occur from the position of the actual aircraft causing it. In particular:

10-4-1010

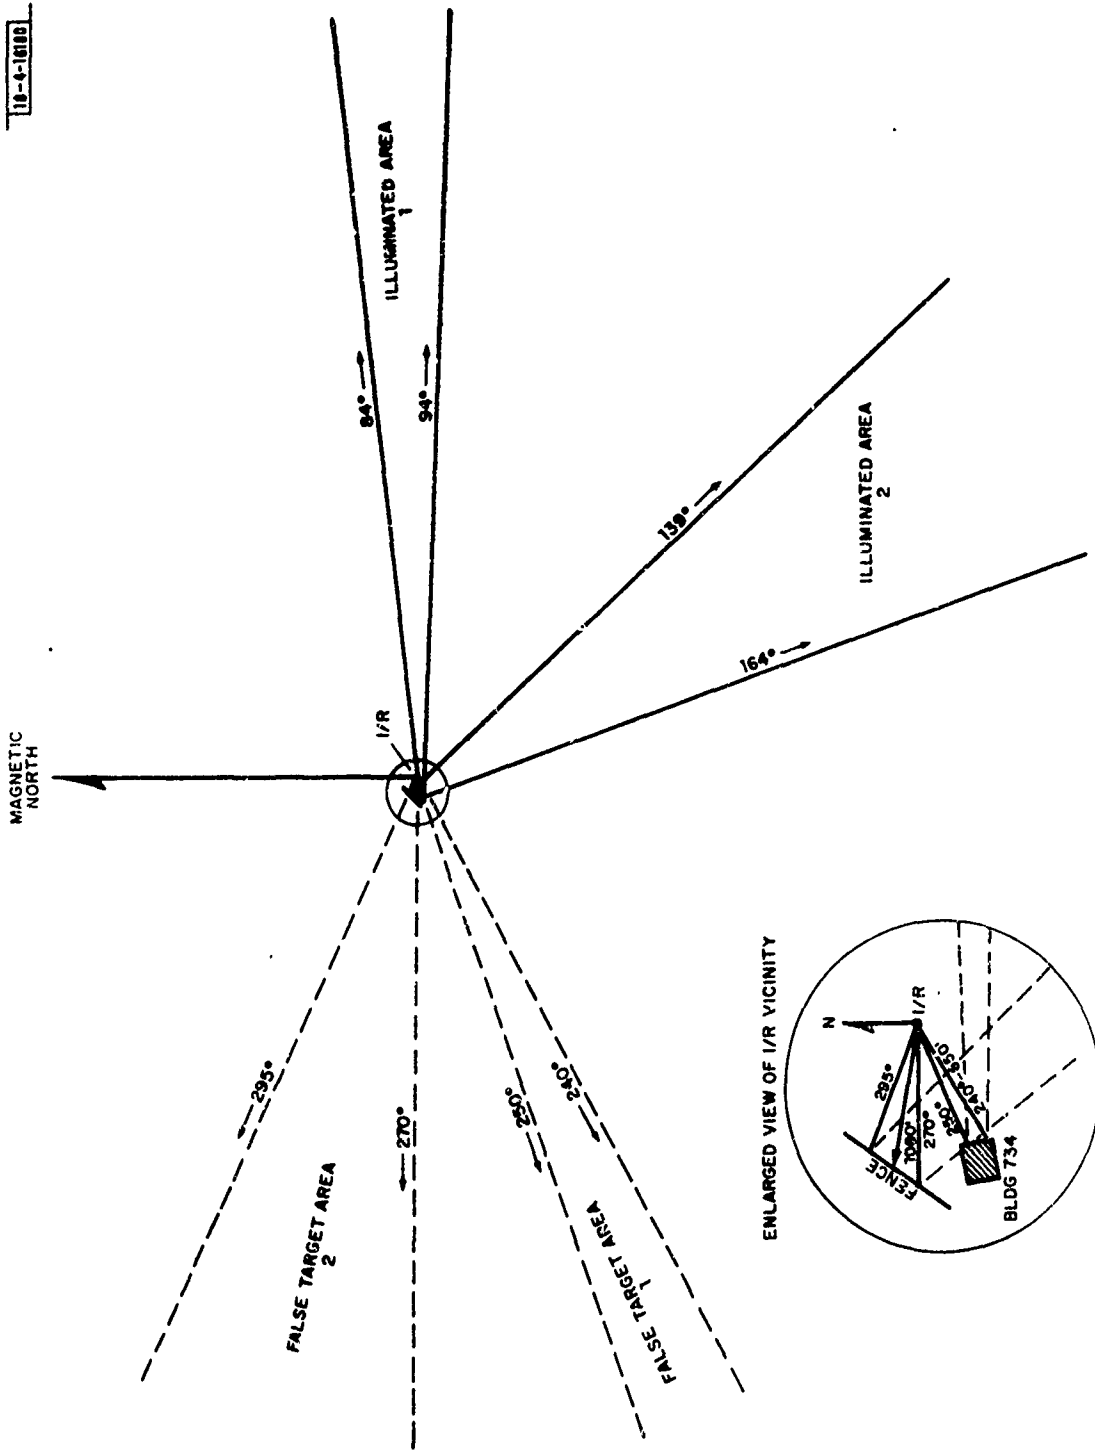


Fig. 5. Close-in reflection geometry at Albuquerque.



$$R_{FT} = R_A + \Delta R$$

$$\theta_{FT} = 2\theta_o - \theta_A .$$

Here, A denotes the actual target and FT denotes the false one.

The process next creates a window of size  $2 d\theta$  by  $2 dR$  around the expected false target location, and examines the subsequent target declarations occurring in the next scan to determine if they, a) fall within the window, b) agree in code and altitude with the actual target, and c) are not updated positions of tracked aircraft which were at one time outside the illuminated area. If all of these conditions are satisfied for a particular target declaration, it is concluded to be false and tagged with a special symbol (e. g., an "F").

Window size is determined by the precision with which the false target position can be calculated, and by the distance over which the target can move between the times when it and the false target occur (typically one to three seconds apart). Manual solutions, in which aircraft position has been interpolated between the two target reports adjoining the false target to the instant at which the false target appears, regularly yield errors less than  $\pm 1/16$  nmi (one range cell) and  $\pm 0.5^\circ$ . Additional error results from the fact that high-speed aircraft could change position by as much as  $1/2$  nmi and  $3^\circ$  during the interval between their legitimate declaration and the time at which they next cause a false target. Thus, window size depends primarily on uncertainty in instantaneous position due to aircraft motion; basing window position solely on aircraft position as of the last declaration leads to a window of moderate

size; basing it on instantaneous (interpolated or extrapolated) position allows the use of an extremely small window.

This technique could conceivably flag a legitimate target as false, if that target was in the right place at the right time, squawking the right code and altitude. It is evident that the probability of that event - albeit very small - is proportional to window size. What is of interest here is whether the window size that results from basing window location solely on previous declared position is small enough to ensure that the probability of declaring a real target as false is maintained at an acceptably low level. Also, do the further reductions in that level that result from using the smaller window based on interpolated data warrant the complexity of the interpolation software? In a low-density environment, it would appear that the likelihood of a legitimate aircraft appearing in a window of moderate size (say, 1 mile by  $6^\circ$ ), and agreeing in code and altitude\* with the aircraft whose presence has caused the window to be generated is exceedingly small. In addition, given that unlikely event, it would be highly unlikely that the relationships between the velocities and headings of the two aircraft would be such that the situation would persist over many scans. In short, it seems appropriate to develop the initial version of false-target-elimination software around the assumption that a window based solely on previous position is sufficiently small; this eliminates the need for interpolation, and the tracking/correlating process that would be necessary in that situation.

---

\*For aircraft not equipped with altimeters, presence or absence of empty brackets could be checked. Since these aircraft are the most likely users of nondiscrete codes (e. g., 1200), perhaps consideration should be given to a more widespread discrete code assignment procedure.

The ultimate output of the process described above and diagrammed in Fig. 6 would thus be a flagging of all targets determined to be false. The determination process would occur independently from scan to scan, and the way in which controllers treated flagged targets would, to some extent, be influenced by the number of scans over which they were flagged as false.

The general procedure described here is now applied specifically to the Albuquerque situation. There, three reflecting surfaces regularly result in declared false targets. One, the Manzano Mountain fence, results in the complicated geometry of Fig. 3, and is understood to cause false targets that do not lead to severe operational problems. The others are the structure (Bldg. 734) located 850 ft away, to the south of west ( $240$  to  $250^\circ$ ), and the fence to the north of it (1000 to 1150 ft from the radar at  $270$  to  $295^\circ$ ). The two are oriented such that they appear from far away to illuminate azimuthal wedges at  $84$  to  $94^\circ$  and  $139$  to  $164^\circ$ , respectively (Fig. 5).

Whenever a target declaration azimuth falls within one of these wedges, its parameters are stored for further processing, and the false target location is determined by:

$$R_{FT} = R_A + 3/16 r_{mi}$$

$$\theta_{FT} = 3798 - \theta_A \text{ ACP}$$

for targets between  $84$  and  $94^\circ$ , and:

$$R_{FT} = R_A + 4/16 \text{ nmi}$$

$$\theta_{FT} = 4942 - \theta_A \text{ ACP's}$$

for targets between 139 and 164°.

Now, the regions defined by  $R_{FT} \pm 1/2$  nmi and  $\theta_{FT} \pm 3^\circ$  are searched for potential false targets. Any declaration falling within one of these regions is examined for code and altitude agreement; if that agreement is noted, the target is labelled false.

Automated processes taking past history into account in determining the certainty with which targets are declared false are possible, perhaps desirable; these all require that tracking logic be employed, and are all, therefore, somewhat more complicated to implement. The degree of added complexity must be weighed against the additional benefits derived in order to determine whether a process involving tracking is more desirable than the simple one described here. That determination is beyond the scope of this report; much detailed information about the operation of the ARTS-III tracker is needed before it can be properly made. However, the following section discusses briefly a possible approach to false target elimination making use of ARTS-III tracking.

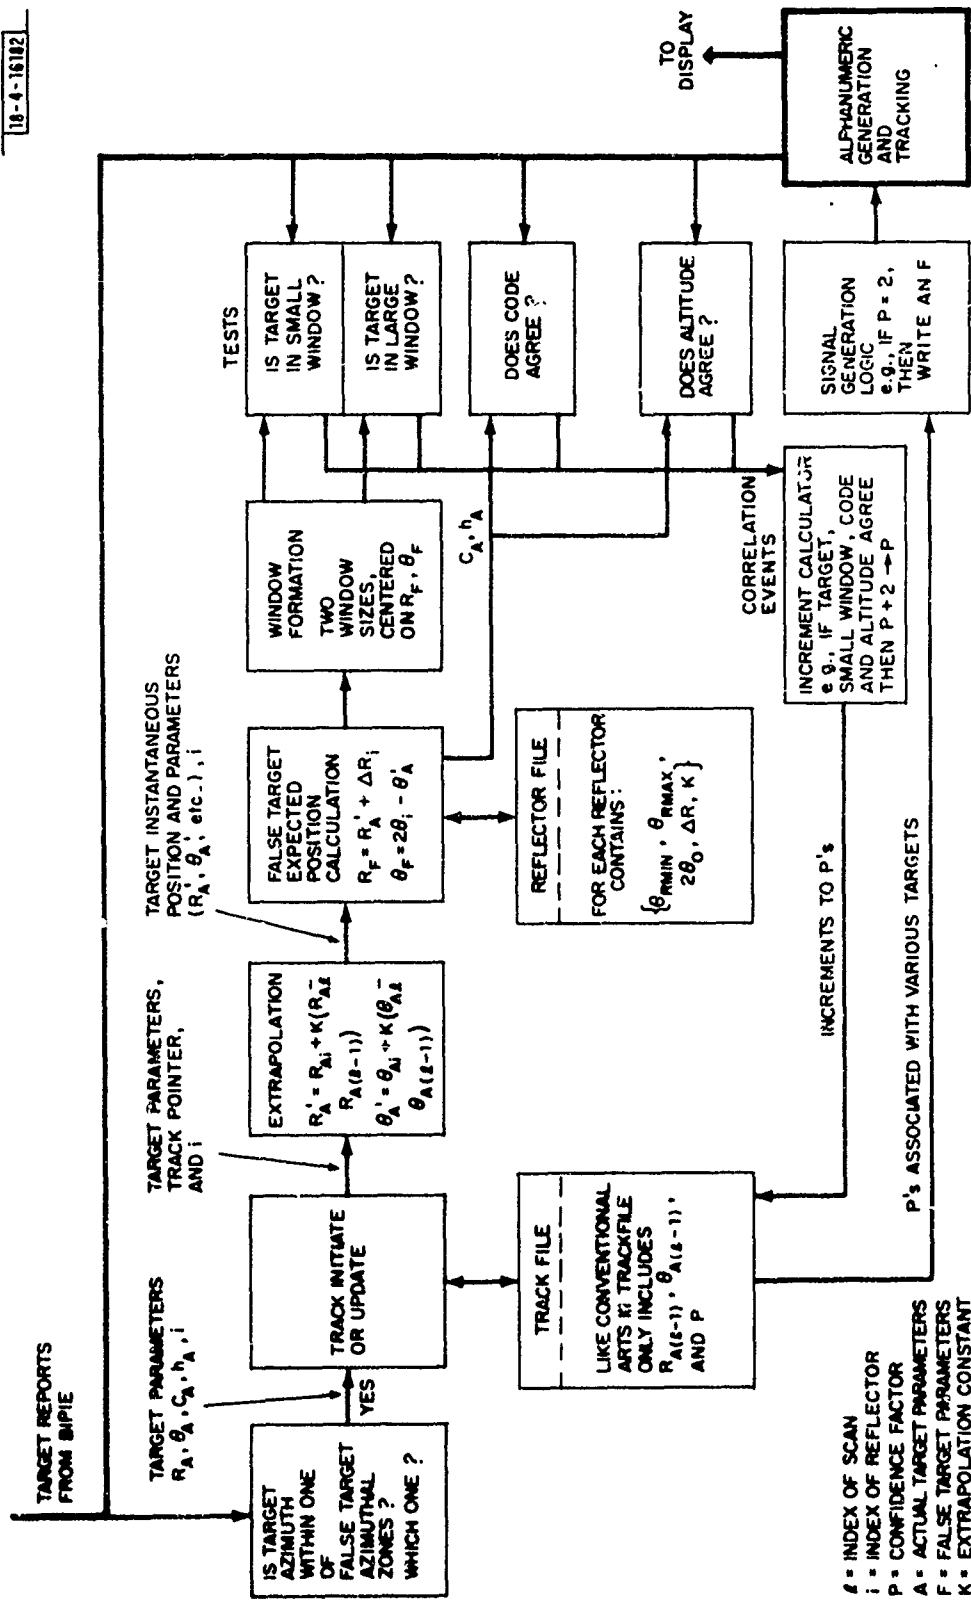
## SECTION 5

### A MORE SOPHISTICATED APPROACH

ARTS-III tracking involves both correlation and smoothing, and is intended in its present version primarily to keep data blocks properly positioned on the display, and to "coast" target symbols through short periods where aircraft replies are lost. It appears necessary to employ some elements of the ARTS-III tracking process, particularly the scan-to-scan correlation of target reports, in any false-target-elimination process which is more complex than the one discussed above.

The section presents a possible false-target-elimination procedure (see Fig. 7) which uses tracking to associate target declaration of a particular aircraft with one another. Many variations of this basic procedure are possible; it should be viewed as typical rather than preferred.

It should be noted initially that experience with ARTS-III reply, target declaration, and tracking data has demonstrated clearly that declared target position data yield far higher precision than tracker output data. The only source of noise of the type which tracking can filter out in the range measurement process is quantization; whenever a target is declared, one can be certain that its range is within 1/16 nmi of the proper value. Tracked positions often deviated by more. Thus, only positions associated with target declarations are employed in what follows.



l = INDEX OF SCAN  
 i = INDEX OF REFLECTOR  
 p = CONFIDENCE FACTOR  
 A = ACTUAL TARGET PARAMETERS  
 F = FALSE TARGET PARAMETERS  
 K = EXTRAPOLATION CONSTANT

Fig. 7. False target elimination algorithm employing tracking.

A brief example is appropriate here. Assume that a particular reflector at an azimuth of  $120^\circ$  is oriented in such a way that it illuminates a wedge of airspace centered about  $30^\circ$ . Assume further that an aircraft has just flown into the illuminated area, and that his present and past positions are as follows:

present:	29 nmi, $29^\circ$
last scan	28 nmi, $28^\circ$
previous scan	27 nmi, $27^\circ$

and so forth,

Here,  $\delta R$  and  $\delta \theta$  are obviously 1 nmi and  $1^\circ$ , and the position of the aircraft extrapolated ahead to the instant the radar points at the reflector is simply (29.25 nmi,  $29.25^\circ$ ), since that occurs one-quarter scan after the legitimate target is detected.

Given the luxury of being able to wait for the target report following false target occurrence, it would be possible to develop a similar process using interpolation rather than extrapolation. This would, of course, result in greater accuracy, since it would account for changes in aircraft heading made subsequent to the target declaration preceding false target occurrence. However, the degree of difference appears to be so small as to be outweighed by the disadvantage of having to wait several seconds after the occurrence of a false target before being able to decide that it's false.

In a manner similar to that used in the simpler procedure, the instantaneous position determined here is used to determine the position of a "window," which is again searched as the antenna azimuth passes through it for target

reports agreeing in code and altitude. In this case, though, the use of correlation that results from tracking allows "softer" decisions to be made; in keeping with this, perhaps two concentric windows should be used. Whenever a potentially false target occurred within these windows, a parameter would be established in the track file corresponding to the actual aircraft in question. This parameter would be similar to the track firmness parameter used in the present ARTS-III tracker, and would be incremented or decremented from scan to scan as confidence in the decision that the target is false grows; depending on its value on a particular scan, the symbology used to identify the false target might vary.

For example, the confidence parameter might be 3 bits long (8 levels), and be set initially to zero. Two window sizes might be used, say  $1/8$  nmi by  $1^\circ$ , and  $1/4$  nmi by  $2^\circ$ . Occurrence of a target agreeing exactly in code and altitude within the smaller window might increment the parameter by 2; a target within the larger window agreeing in code and altitude might increment it by 1; a target in the smaller window agreeing in code but not in altitude might increment it by one. Presence of a target agreeing neither in code nor altitude might not increment it at all. The absence of any target in either window might decrement the parameter by two. Thus, four declarations in a row, each agreeing in code and altitude, and each within the smaller window (implying correlation in range, azimuth, velocity, and heading between the actual aircraft and suspect false target) would suffice to drive the confidence parameter to its maximum value. More sporadic occurrence of a false target would hold its confidence parameter to a lower value.

The value of the confidence parameter would be used to determine the display symbology associated with the suspect false target. For example, a level of one or two might cause it to be tagged with a blinking "F." When the level reaches three or four, the symbol might no longer blink. A level of five or six might cause it to be tagged with a data block stating "CONFIRMED FALSE." In the future, when the decoded beacon video that is displayed on ARTS-III is available to ARTS-III for more sophisticated processing, a higher confidence level might result in the elimination of the false target video from the display.

## SECTION 6

### RECOMMENDATIONS AND CONCLUSIONS

While the more sophisticated approach to false target elimination described in Section 5 appears promising as a long-range general solution to the problem of multipath reflections, especially at sites with extremely high traffic levels, it is probably more appropriate to attempt to implement the simpler fix described in Section 4 initially at Albuquerque. That task should be simple, straightforward, and relatively inexpensive.

Since false target problems appear to be abundant throughout the FAA Secondary Surveillance Radar system, it is worthwhile to consider more widespread implementation of this corrective software. At some future time, when corrective action is considered for sites with extremely high traffic densities and false target incidences, or when higher quality output is required to support increasing levels of automation, it might be worthwhile to consider a more sophisticated approach, such as the one described in Section 5.