SURVEILLANCE PROCESSING IN THE MODE S SENSOR

J.L. Gertz

Massachusetts Institute of Technology, Lincoln Laboratory, USA

INTRODUCTION

The principal function of the Mode S sensor (1), an evolutionary upgrade to the current ATCRBS (Air Traffic Control Radar Beacon System) sensor, is the output of one report per aircraft per antenna scan. This report contains the current aircraft position (range and azimuth), the identity code of its transponder, and the altitude code as supplied by its encoding altimeter.

This information is derived from the aircraft transponder replies received at the sensor in response to interrogations transmitted by the sensor. For aircraft equipped with Mode S transponders, a single scheduled interrogation, directed only to that aircraft, elicits a single coding-protected reply containing both identity code and altitude code. For aircraft equipped with ATCRBS transponders, a sequence of interrogations alternately elicit replies containing un-protected identity code or altitude code from all aircraft in the antenna mainbeam.

From this description, it is clear that a Mode S aircraft report can be constructed directly from the single reply. Surveillance processing, defined as functions that perform scan-to-scan correlation and tracking, are required in general only to predict the next scan position of the aircraft. This information is needed for the proper scheduling of the next interrogation.

ATCRBS reports constructed from the aircraft replies, on the other hand, can have a number of deficiencies. The more common such problems are:

1. Either the identity code or altitude code or both can have bits declared either in error or with low confidence by the reply processor due to garbling of overlapping replies.

2. False alarm reports not corresponding to aircraft can be generated from fruit replies (responses to other sensors' interrogations) or reflection replies.

3. Multiple reports for an aircraft can be generated due to incorrect correlation of replies caused by errors in range, azimuth, or code determination.

Surveillance processing for ATCRBS aircraft is tasked with correcting these problems prior to report output to the controllers or other users. It does this by correlating raw target reports with existing track files, and using the information in these files derived from prior scan reports to correct, complete, or reject erroneous reports.

This paper presents the major algorithms contained within the Mode S sensor ATCRBS surveillance processing function. It then presents experimental results that demonstrate their effectiveness. Full details of surveillance processing can be obtained by reference to (2) or (3).

CORRELATION OF DISCRETE ATCRBS AIRCRAFT

Aircraft with ATCRBS transponders have their identity codes assigned during flight by controllers, and they may change several times during a flight. Two types of ATCRBS codes exist: discrete codes, which are intended to be unique within a control area, and non-discrete codes (such as 1200 in the US), which are assigned to all aircraft within a given flight class.

Discrete ATCRBS aircraft can generally be correlated through use of identity code alone, subject only to a "reasonableness test" on positional difference between target report and track. Should a complication exist, such as the code being garbled or non-unique, or the reasonableness test fail, the more complex procedures described next are employed.

CORRELATION OF NON-DISCRETE ATCRBS AIRCRAFT

Target-to-track correlation for non-discrete ATCRBS aircraft can be quite complicated, especially when several target reports are spatially close to a single track, or several tracks are near a single report (or even an intertwined collection of spatially close reports and tracks). Furthermore, even when only one report is near a given track, it is possible that the report and track represent different aircraft, and thus should not be correlated.

The correlation process for non-discrete reports is broken down into two parts. Association is defined as the identification of all possible feasible report and track pairings. Thus any association could constitute a valid match. Correlation is then defined as the selection of the best set of matches of targets and tracks whenever alternatives exist.

ASSOCIATION TESTS

Association of an ATCRBS report to an ATCRBS track is achieved whenever the three-dimensional (range ρ, azimuth θ, altitude h) positional difference between them is sufficiently small. Identity code agreement cannot be an association requirement as aircraft can change their codes; however, the degree of code match does determine the permissable amount of positional deviation.

Range and azimuth agreement are scored by placing three concentric range-azimuth (ρ-θ) zones around the predicted track position. Zone 1 encompasses measurement errors possible for an aircraft flying straight, zone 2 includes positions achievable by an aircraft performing a standard turn, and zone 3 covers unusual aircraft maneuvers.

Identity code and altitude matches are scored on an agree, potentially agree, or disagree scale. "Agree" for code is exact match; for altitude it is within 100 feet per second difference. "Potentially agree" means that agreement could be achieved if a single decoded bit were reversed, and thus a reply processing error may have ocurred. Report and track pairs with altitude disagreement are rejected at this point unless code swapping is suspected (see below).

Pairings agreeing on code and altitude and within zone 1 or zone 2 are unconditionally associated; those in zone 3 not agreeing on both code and altitude are rejected; others are further tested by the velocity reasonableness test defined below before an association decision is made.

Code Swapping

There are three situations where the reply correlation logic can create two reports for an aircraft, one containing the correct identity code and an incorrect (or non-existent) altitude, and the other containing the correct altitude and an incorrect (or non-existent) identity code. The first occurs when two aircraft are at about the same range and the reply correlation logic creates two target reports, each with the identity code of one aircraft and the altitude of the other aircraft.

The second instance of incorrect reply correlation occurs when a fruit reply is received just prior to the first real aircraft reply, while the third instance occurs when the first aircraft reply is decoded improperly by the reply processing hardware. In either of these cases, one report will contain a correct code and an incorrect one (caused by fruit or error), while a second report will contain the correct other code only.

Code swapping corrects all of these errors by physically interchanging the identity codes between the two reports. In the first case two proper reports result, while in the other cases one proper report and one rejected report are produced.

Velocity Reasonableness Test

The intent of the Velocity Reasonableness Test is to determine the likelihood of a current target report being part of the same report sequence as that represented by a given track. Two velocity vectors are employed in the test.

The first is $\vec{w}$, which is the last known velocity for the aircraft under track.

The second vector is $\vec{v}$, which would be the actual current velocity of the aircraft if the report in fact corresponds to it. The test basically judges the reasonableness of the required aircraft velocity change from $\vec{w}$ to $\vec{v}$. The vector comparison that constitutes the velocity reasonableness test is accomplished in two parts: angle and magnitude. The first part of the test is successfully passed if the angle difference is sufficiently small, that is, if,

$$\frac{\vec{w} \cdot \vec{v}}{|\vec{w}| \times |\vec{v}|} \geq -(f-1) * P_1 \qquad (1)$$

where f is the number of scans since the last report and P_1 is a parameter. Thus, a reversing motion is forbidden for a steady-state track.

The magnitude test checks for situations in which the velocity increase exceeds a reasonable limit. The association passes this part of the test whenever:

$$\frac{|\vec{v}|}{|\vec{w}+\vec{\varepsilon}|} \leq f \times P_2 \qquad (2)$$

where P_2 is another parameter. The vector $\vec{\varepsilon}$ is the measurement error vector in the direction of motion and serves to make the test conservative. The dual test, of $\vec{v}$ being too small, is not made; the angle test partially covers this case, and the error term would lead to automatic success in most situations.

CORRELATION SELECTION

Once the above association process has identified all acceptable target report to track pairings and corrected code swap errors, it is the function of the correlation process to select the "correct" matches from this set. In the large majority of cases, only one report will associate with a track, and that track will associate with only that one report. In other cases, however, a choice will exist, and hence selection criteria are required.

The primary selection mechanism is the utilization of the Quality Score computed for each association. This score, described in detail below, evaluates the degree of positional and attribute match between a given report and track pairing. Its value is zero for a perfect match.

The two simple cases of correlation selection are denoted by m tracks on one report and one track on n reports. Each of these situations is resolved by selecting for correlation the pairing with the lowest Quality Score. Ties are broken via the Deviation Score, described below.

The remaining correlation case is the intertwined m-on-n situation: m tracks and n reports form a closed association set (although not every report associates with every track in general). The "optimum" resolution of this situation is defined to be the set of selected pairings that minimizes the total of the quality scores for all reports and tracks. Correlated reports or tracks are assigned the quality score of their selected pairing, while a "penalty" quality score is assigned to leftover (non-correlated) reports or tracks.

Quality Score

Figure 1 presents in detail the items entering into the Quality Score as well as the manner in which each of these items is evaluated and the individual scores for each possible result. The final Quality Score for the association, as indicated in the figure, is the octal concatenation of the component test scores. (The penalty quality score mentioned above is set at octal 50000000.)

Since it is impossible for the score of one component test to carry over into the digit of the next one, this Quality Score is actually an implementation of a multi-stage decision algorithm. That is, if two associations exist for a track, the one chosen will be the one with the lower code-zone (digit 7) score, even if that association rated lower on all other criteria. If the associations tie on this criterion, however, the decision will be based on the next item, etc. Because all the decision item scores are combined into a single number, however, a single comparison will automatically implement the entire test hierarchy, selecting the winning association on the basis of the first non-tied decision stage.

Deviation Score

It is quite possible that the Quality Scores of two associations will be identical. For example, reports from two closely spaced general aviation aircraft, both reporting a code of 1200 and having no encoding altimeters, would often produce the same score relative to any track. The intent of the Deviation Score is to break such ties by taking into account the geometric difference between the track and target positions.

The Deviation Score doesn't merely reflect the distance between the positions; rather it indicates the likelihood of the aircraft under track being at the position represented by the target report. In particular, the scoring rules employ the fact that changes in aircraft speed from scan to scan are unlikely, most changes in aircraft velocity being caused by turns.

The Deviation Score (D) assigned to each point is computed as the product of two factors: one (f_1) that penalizes absolute distance from the predicted track position and the second (f_2) that penalizes deviations from the turning locus (approximated as a straight line segment). The two vectors needed for this computation, as shown in Fig. 2, are labelled $\vec{d}$ and $\vec{t}$. The former represents the deviation of the report relative to the predicted track position, while the latter is a unit vector in the direction of the approximated turning locus.

The penalty factor for absolute distance between target and track is defined to be:

$$f_1 = \left|\frac{\Delta\rho}{\epsilon_\rho}\right| + \left|\frac{\rho\Delta\theta}{\rho\epsilon_\theta}\right| \qquad (3)$$

where ϵ_ρ and ϵ_θ are the 3σ report measurement errors. The factor that rates the direction of this deviation does so by comparing its components in the directions parallel and perpendicular to the turning locus. That is:

$$f_2 = \frac{C_{perp}}{C_{par}} \qquad \text{where:} \qquad (4)$$

$$C_{par} = \vec{d} \cdot \vec{t}, \quad C_{perp} = \sqrt{|\vec{d}|^2 - C_{par}^2} \qquad (5)$$

Thus, deviations due to turns ($C_{perp} \approx 0$) are penalized very little compared to those requiring along-track accelerations.

TRACK INITIATION

ATCRBS tracks are automatically initiated when a pair of uncorrelated reports are found on successive scans that appear to have come from the same aircraft. To satisfy this criterion, ATCRBS reports must agree or potentially agree on both identity code and altitude. In addition, the physical separation of the two reports must be sufficiently small that a real aircraft could have traversed the distance in one scan.

Not all uncorrelated reports enter into the track initiation process. Under user control, various categories of reports that are judged not likely to be due to real aircraft can be eliminated (see below). The remaining uncorrelated reports are compared with those from the previous scan. If one or more matches are found, the report is used to start new tracks; otherwise, the report is added to the uncorrelated report buffer for comparison with subsequent scan reports.

Linked Tracks

The algorithm described above permits one report to initiate more than one new track. Since any one report can only correspond to one aircraft, it is clear that in such cases extraneous tracks have been formed. Although the proper track of the set is not known at initiation time, it will become evident on a subsequent scan. Thus, when one track of the set is correlated and the others coasted, these latter ones are dropped at once to prevent erroneous future correlations.

Whenever a current scan report initiates more than one track, these tracks are linked together. If only one track is formed, this track is linked to the tracks (if any) previously formed by the last scan report. This pair of rules guarantees that all tracks in a linked set have one report in common, and thus that only one can be real.

TRACK UPDATE

Each ATCRBS track has the information in its track file updated once per scan. If the track was correlated with a target report, the position and velocity predictions and the identity code and altitude values will all be modified according to the new data provided by that report. This report, in turn, will then be improved by using the many-scan composite information available in the track file. Uncorrelated tracks, on the other hand, are merely coasted ahead one scan by using the velocity estimate contained in the track file.

In the normal situation, the track position and velocity predictions are made by interpolating ahead the last two target data points in ρ, θ coordinates. This type of tracker, known as a 2-point interpolator or an $\alpha=1$, $\beta=1$ $\alpha\beta$ tracker, is sufficiently accurate when fed monopulse Mode S data to provide the predictions required for target to track correlations. In addition it provides the desirable feature of immediate sensitivity to turns. A very rudimentary form of turn detection (see below) is added to this tracker to prevent fatal track deviations when potentially spurious data points are encountered.

When an aircraft flies near the sensor, the errors inherent in simple ρ, θ tracking become sufficiently large that target-to-track correlation could no longer be supported. Thus an improved method of tracking is required. Two alternative methods are possible: second (or higher) order ρ, θ tracking and coordinate converted x, y tracking. Both of these methods are utilized.

Code Improvement

Once the best estimates of the track's identity code and altitude have been created by merging the new report information with that from previous scans, it is possible to improve these fields in the target report. The rule for the identity code is the simpler. If the report identity code agrees with that of the track in all mutual high confidence bits, the report identity code can be upgraded to that of the track. Thus incomplete or garbled identity codes can be completed prior to output.

Since altitudes can change from scan to scan, this bit-merging algorithm cannot be utilized for the altitude code. However, two improvements can be made. First, if the report altitude code has some low confidence bits, but the decoded altitude derived from treating all bits as correct produces a value near that of the track, this decoded altitude can be placed in the report. Second, if the report altitude code had no low confidence bits, but its decoded value is significantly different from that of the track, the report altitude is returned to Gray code form and marked as garbled. This rule catches altitude reply bit decoding errors and prevents spurious altitudes from being reported in the output.

Turn Detection

The normal 2-point interpolation tracking accepts the position measurement of each report as the current smoothed position. To prevent a single erroneous correlation or a measurement error situation from seriously deviating a track trajectory, smoothing beyond the track's zone 1 association box is not permitted for well-behaved straight-flying tracks.

Instead, when a suspect situation is encountered, the track is smoothed in the offending coordinate(s) only to the limit of the zone 1 box. Then, should the next correlating target report again fall outside of the zone 1 association box in the same direction, full smoothing is utilized on that scan and is maintained for the duration of the aircraft turn.

This process thus implements a very simple turn detection mechanism. The first report in a turn is treated as suspect, but once the turn is confirmed, the data points are followed fully.

DATA EDITING

In any ATCRBS surveillance processing system, some of the target reports created by reply correlation do not correspond to the positions of real aircraft. The main categories of such false alarms are reflection false targets, fruit targets, and split targets. Various algorithms are included in the surveillance processing functions to either mark these reports as false or to edit them out of the system output stream. The effectiveness of these data editing algorithms on real data is presented in the last section of the paper.

Reflection false targets are generally caused by the reflection of interrogations and replies off buildings or other structures, causing an apparent aircraft location behind the reflector. With discrete ATCRBS aircraft, the presence of two targets with the same unique code flags such an occurrence; the longer range report can then be rejected in most cases. The algorithm for detecting reflection false alarms when the identity code is garbled or not unique is discussed in the next section.

It is possible for two fruit replies from different aircraft to coincidentally agree in range and azimuth and thus produce a fruit report. These reports will generally consist of one identity reply and one altitude reply, as code agreement would be required for replies of the same mode to correlate. Thus, a report based on one identity reply and one altitude reply is dropped by data editing if it does not correlate with an existing track.

Various system defects can cause the reply sequence from an aircraft to be split (i.e., separated into two or more target reports.) Range splits are usually caused by transponder delay differences for identity and altitude replies. Thus when two reports are found that agree in azimuth and are close in range, one with identity code replies only and the other with altitude code replies only, the reports are reconstructed into one.

Azimuth splits are usually caused by tailing of the monopulse values of the reply sequence. This can occur when the reply elevation angle or frequency is sufficiently different from the values used for calibration that the calibration curve slope doesn't match the reply slope. The report from the edge-of-beam replies will then almost always look like and be treated as a fruit report.

Finally, reply decoder high confidence bit errors can cause code splits. As these are fairly common, the code swapping mechanism described above was created to correct the errors. The leftover code swap reports, after resolution, are then discarded.

Reflection False Targets

The geometrical situation that exists when a reflection false target is produced is depicted in Fig. 3. The angle θ and range ρ are contained in the suspect target report, while the reflector distance d and orientation angle ϕ are parameters that apply to reflectors at the sensor site. The unknown values that must be calculated are thus the range ρ' and azimuth θ' of the aircraft generating the false target. The set of equations that are used to compute θ' and ρ' are presented in the figure.

Ideally, if the report is indeed false, a track will be found whose position is very close to the calculated point and whose code and altitude agree perfectly with the report. Unfortunately, this ideal state is often not encountered.

This problem has been resolved by defining two sets of match criteria. A candidate target report is called false if a track is found that is within a small box of the computed position and that agrees on both identity code and altitude with the report. Otherwise, a report is called possibly false if the best track found is only within a larger box of the computed position, or only potentially agrees on either identity code or altitude. If no track agrees to either degree on all three attributes of position, identity code, and altitude, the report is called real. To prevent real tracks from being suppressed in error, a track is not labelled false until two false reports have been received; possibly false reports merely delay the final decision.

RESULTS

The major improvements that result from using surveillance processing on ATCRBS reports are code improvement and false alarm rejection. To see the effectiveness of these efforts, data recorded at Washington National Airport by an early version of the ATCRBS part of the Mode S sensor was analyzed. The numbers presented herein represent 100 radar scans.

Table 1 demonstrates the effectiveness of the identity code improvement process. It shows the percent of the reports that had ungarbled identity codes before and after surveillance processing. Crossing tracks are defined to exist when two aircraft are within two miles and 4° of each other, which is the mutual garble region. Clearly a significant increase in code reliability has been achieved.

The effectiveness of the data editing algorithms in detecting false alarm reports is illustrated by Table 2. As seen, 12.8% of all reports produced by reply correlation have been rejected; without surveillance processing, these reports would have been output to air traffic control. Of the remaining reports, only 2% were found to be extraneous, i.e. not representing aircraft, by after-the-fact analysis.

Figure 4 pictorially demonstrates the display improvement provided by surveillance processing. The first part of the figure shows all reports output by reply correlation for a 20 mile by 20 mile region for the 100 scans. The second part of the figure provides a picture of the reports rejected as due to reflections or other false alarm effects. Note that several fairly long reflection tracks indicating apparent aircraft are included in this data. These tracks would have presented problems to ATC if not detected. Finally, the last part of the figure presents the real correlated reports output by surveillance processing, the set of reports that would be input to air traffic control users. Comparing this display to the first one, it is clear that a substantial improvement in data quality has been achieved.

REFERENCES

1. Orlando, V.A., and Drouilhet, P.R., 1986, "Mode S Beacon System: Functional Description", Report No. DOT/FAA/PM-86/19, U.S. Government.

2. Gertz, J.L., 1977, "The ATCRBS Mode of DABS", Report No. FAA-RD-76-39, U.S. Government.

3. Department of Transportation, Federal Aviation Administration, 1984, "Mode Select Beacon System (Mode S) Sensor", Specification No. FAA-E-2716, U.S. Government.

TABLE 1 - Reports with ungarbled identity code

TRACKS	REPLY CORRELATOR OUTPUT	SURVEILLANCE OUTPUT
All	92.8%	97.5%
Crossing	76.5%	91.9%

TABLE 2 - Report categorization

Reply Correlator Output:	11321 Reports
Reflections:	317 - 2.8%
False Alarms:	1135 - 10.0%
Surveillance Processor Output:	9869 Reports
Aircraft:	9671 - 98.0%
Extraneous:	198 - 2.0%

Octal Digit and Conditions	Score
Digit 7 - zone and code	
zone = 1, code agree	0
zone = 2, code agree	1
zone = 1, code disagree	2
zone = 2, code disagree	3
zone = 3, code agree	4
Digit 6 - number of replies	
3 or more (identity and altitude)	0
2 of same type	0
1 of each type	1
1 reply	2
Digit 5 - code agreement	
agree, all bits high confidence	0
agree, some bits low confidence	1
no code	2
potentially agree	3
disagree and:	
some bits low, track code in transition	4
all bits high, track code in transition	5
some bits low, track code steady	6
all bits high, track code steady	7
Digit 4 - altitude agreement	
$\Delta h < 500$ feet	0
$\Delta h = 600$ feet	1
$\Delta h = 700$ feet	2
$\Delta h = 800$ feet	3
$\Delta h = 900$ feet	4
$\Delta h = 1000$ feet	5
$\Delta h > 1000$ feet	6
Digit 3 - track validity	
track established, $\rho > \rho_{valid}$	0
track established, $\rho < \rho_{valid}$	1
new track, $\rho > \rho_V$	2
new track, $\rho < \rho_V$	3
Digit 2, 1, 0 - deviation score computed when required as tie breaker	

Quality Score = $(d_7 d_6 d_5 d_4 d_3 d_2 d_1 d_0)_8$

Figure 1 Quality Score determination

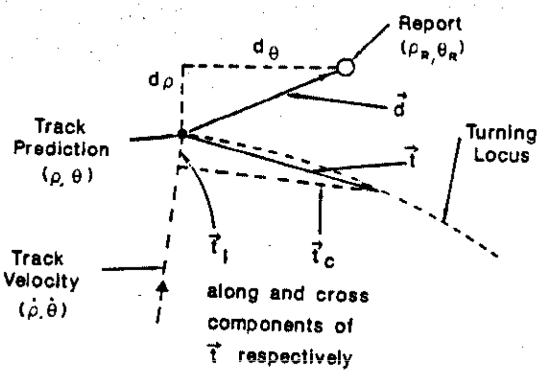

$$\vec{d} = (d_\rho,\ d_\theta) = ((\rho_R - \rho),\ \rho(\theta_R - \theta))$$

$$\vec{t} = \vec{t}_c - \vec{t}_l$$

$$\vec{t}_c = (t_{c\rho},\ t_{c\theta}) = (-A\rho\dot\theta,\ A\dot\rho)$$

$$\vec{t}_l = (t_{l\rho},\ t_{l\theta}) = (B\dot\rho,\ B\rho\dot\theta)$$

Where: $A = \frac{1 - \cos w}{w}$, $B = 1 - \frac{\sin w}{w}$

turn rate = w radians per scan

Figure 2 Deviation Score vectors

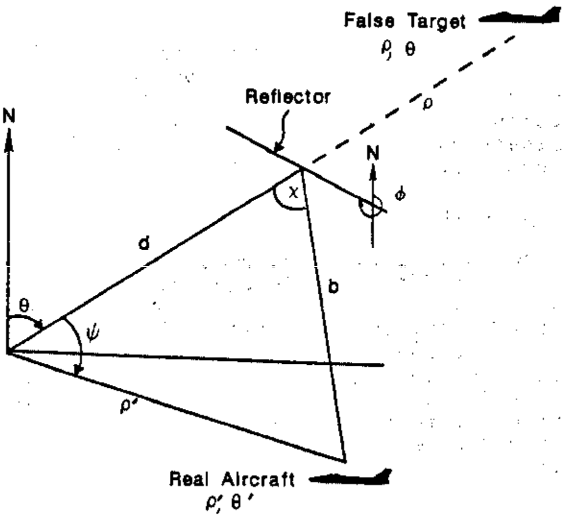

$$b = \sqrt{(\rho - d)^2 - h^2} \quad \text{(if arg < 0, target real)}$$

$$\chi = 540° - 2\,[\phi - \theta]$$

$$\rho' = \sqrt{b^2 + d^2 - 2db\cos\chi}$$

$$\psi = \sin^{-1}[b\sin\chi/\rho']$$

$$\text{if } b^2 > (\rho')^2 + d^2,\quad \psi = \begin{cases} 180-\psi & \text{if } \psi > 0 \\ -(180+\psi) & \text{if } \psi < 0 \end{cases}$$

$$\theta' = \theta + \psi$$

Figure 3 Computation of real aircraft position

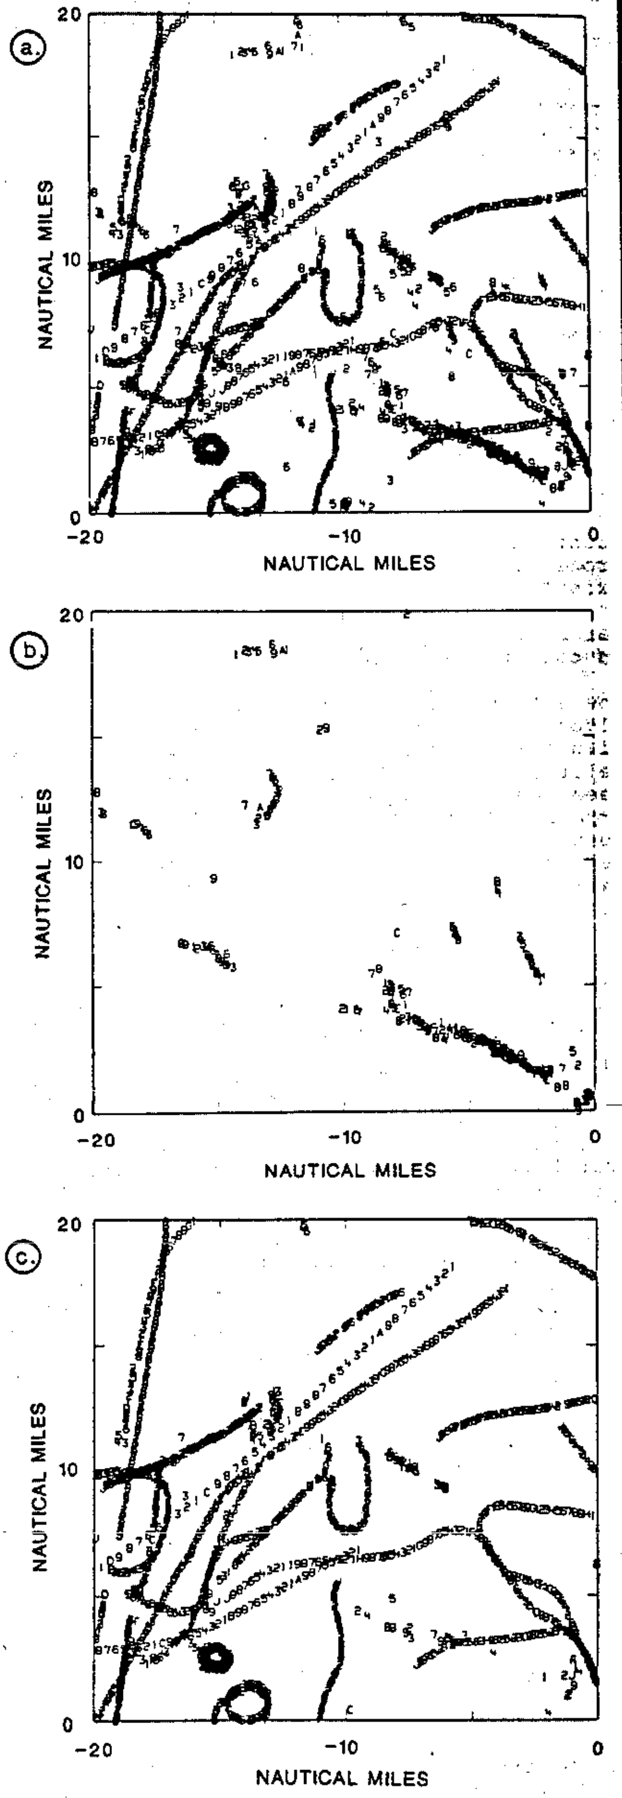

Figure 4 (a) Reply correlation output report
(b) Reflection and false alarms
(c) Correlated surveillance outputs