REPORT-TO-TARGET ASSIGNMENT IN MULTISENSOR MULTITARGET TRACKING*

Dr. Thomas Kurien
ALPHATECH, Inc.
111 Middlesex Turnpike
Burlington, MA 01803

Captain Martin E. Liggins, II
Rome Air Development Center
Griffiss Air Force Base
Rome, NY

ABSTRACT

This paper discusses two issues that need to be addressed for extending a practical single sensor multitarget tracking algorithm to the case of tracking with multiple sensors. These issues correspond to the compensation of the algorithm for asynchronous reporting rates from multiple sensors, and the interpretation and implementation of an nscan approximation technique for assigning reports from multiple sensors to targets. Procedures that account for these issues and have been demonstrated using real data are also provided.

1. INTRODUCTION

The function of a multitarget tracking algorithm is to track all targets within the surveillance volume of interest based on reports from surveillance sensors. Tracking a target typically involves the estimation of its position and velocity. Surveillance sensors provide reports on a scan-by-scan basis. The measurements contained in each sensor report depend on the type of sensor: active sensors such as a radar typically provide range, angle, and range-rate measurements; passive sensors such as an aero-acoustic or IR sensors typically provide only angle measurements.

To track multiple targets, the tracking algorithm has to first correlate the reports corresponding to each target from one scan to the next. The measurements from these reports are then used to estimate the position and velocity of the target. The key problem in multitarget tracking is the correlation of the reports for each target from one scan to the next. Factors which make this correlation process difficult are the presence of multiple maneuvering targets, the presence of false reports, missed detections for targets, random initiation and termination times for targets, and sensor measurement noise. Our objective is to show how a single sensor multitarget tracking algorithm, designed to address each of these factors, has been extended to the case of tracking with multiple sensors.

To satisfy this objective, we have organized the paper as follows: Section 2 will review the single sensor multitarget tracking algorithm and identify two issues which need to be addressed for the case of tracking with multiple sensors; Section 3 will discuss the first issue which is that of compensating for the asynchronous reporting rates of multiple sensors; and Section 4 will discuss the second issue which is that of correlating the multiple sensor reports for each target. These two sections will also specify procedures to account for these issues. Section 5 provides results derived from real data which demonstrate the effectiveness of these procedures.

2. REVIEW OF SINGLE SENSOR MULTITARGET TRACKING ALGORITHM

The single sensor multitarget tracking problem has been studied extensively and several algorithms have been proposed in the past [1]. Some of these algorithms, especially those used in operational systems, are based on heuristic rules formu-

lated using intuition and experience from actual surveillance scenarios. Since these algorithms are not based on a comprehensive model of the system they fail to handle situations other than those considered during the design.

Other single sensor multitarget tracking algorithms are based on the principle of statistical estimation theory [2-4]. These approaches define a precise mathematical model for the system, formulate the equivalent mathematical problem, and develop the algorithm to compute the optimal solution. These approaches satisfy most situations, but they generally have large computational requirements. To cut down the requirements of the optimal solution algorithm to manageable levels, carefully designed heuristic rules may be incorporated within the optimal solution without any significant loss in performance. In fact, we have designed and demonstrated such an algorithm. We will review the mathematical approach and the heuristics used to develop this single sensor multitarget tracking algorithm. The reader may refer to [4] for a detailed description of the algorithm.

2.1 OPTIMAL ALGORITHM**

We have used the mathematical framework of Hybrid State Estimation to formulate the solution methodology for the multitarget tracking problem. The general hybrid state model consists of both continuous-valued states (X) and discrete-valued states (Z). Measurements (Y) related to the hybrid state are used to evaluate an optimal (minimum-mean-squared-error or maximum-a-posteriori) estimate of the hybrid state. Variables in multitarget tracking can be identified with the generic hybrid model as follows: the state (position and velocity) of all targets constitute the continuous-valued state; indicators for individual target histories (birth time, dynamic model history, missed detection history, etc.) and indicators for individual reports in each scan (associated with a target, false alarm, etc.) constitute the discrete-valued state; and the noisy measurements of range, angle, and range-rate from a sensor constitute the measurements.

The methodology for evaluating the optimal estimate of the continuous-valued state is illustrated in Fig. 1. Sensor reports are grouped into several sets, each set representing a different interpretation of all the reports. Each interpretation corresponds to a particular value for the discrete-valued state i.e., it will consist of one set of feasible target histories indicating the assignment of individual sensor reports to targets in each scan. For each value of the discrete-valued state an optimal estimate of the continuous-valued state conditioned on this discrete-valued state (i.e., $p(X/Z, Y)$) may be evaluated. The normalized measurement residuals from these optimal estimators may then be used to evaluate the probabilities of the discrete-valued states $P(Z/Y)$. Finally, the probability density function of the continuous-valued state $p(X/Y)$ may then be evaluated from the probability density function $p(X/Z, Y)$ and the probabilities of the discrete-valued state $P(Z/Y)$. The equations for evaluating these probability density functions and probabilities are provided in [4].

* Research supported by Rome Air Development Center under Contract No. F30602-85-C-308.

** Even though we will discuss the approach for evaluating the optimal solution to the multitarget tracking problem in the context of a single sensor, the approach is valid for the multisensor case too.

Figure 1. Evaluation of the Optimal Estimate of Hybrid State.

Postulation of all the feasible discrete-valued states and the computation of their probabilities are the functions that make the evaluation of the optimal estimate computationally complex. Each discrete-valued state has to identify a history for each target and the false reports in each scan. In practice, it is convenient to construct these discrete-valued states recursively from one scan to the next. Target dynamic models link the target reports from one scan to the next. The approach that we have used is to correlate the sensor reports from one scan to the next using these models, then construct the discrete-valued states based on these correlated reports. This two step approach for constructing the discrete-valued states is referred to as the *track-oriented* approach.

The track-oriented approach constructs a target tree for each postulated target based on the target's dynamic model and the sensor reports. The root of each target tree represents the birth of the target, and the branches represent the different dynamic models which may be used by the target and the various reports with which the target may be associated in subsequent scans. A trace of successive branches from the root to a leaf of the tree corresponds to a potential *track* for the target. For convenience in notation, we refer to the target trees as *track hypotheses*, and the discrete-valued states as *global hypotheses*.

Track hypotheses are constructed recursively from one scan to the next. At each scan, new branches for existing tracks in each target tree are formed as follows: first new branches are formed corresponding to the different dynamic models for the target; each of these branches is then expanded to account for feasible associations with sensor reports in the scan. This is illustrated in Fig. 2 where we have first split the target track to account for a constant velocity model (S), a maneuver model (M), and a target termination (χ); we have then split the constant velocity track to account for associations with reports R_1 and R_2 , and to account for a missed detection ($\bar{D}$). A global hypothesis is formed by combining tracks from different target trees picking at most one track from each target tree. Global hypotheses are thus maintained as lists of pointers to the target tracks as indicated in Fig. 3. Likelihoods of the global hypotheses are then evaluated from the likelihoods of the tracks. The likelihoods in either case are proportional to the probabilities, and the equations for evaluating these likelihoods are provided in [4].

2.2 PRACTICAL ALGORITHM

The track-oriented approach represents a systematic method for constructing the global hypotheses of the hybrid system modeling the multitarget tracking problem. It can be used to construct the optimal solution to the multitarget tracking problem if computational resources can support the postulation of all global hypotheses. However, computers used in today's surveillance systems do not have these resources. In order to

Figure 2. Construction of New Branches for Existing Track.

Figure 3. Illustration of Track Hypotheses and Global Hypotheses.

construct a practical algorithm, all unlikely global hypotheses have to be eliminated. Several such elimination techniques have been incorporated in our algorithm. The use of multiple sensors for target tracking has an impact on one of these techniques -- the *nscan* approximation technique. In the remainder of this subsection, we will discuss the *nscan* approximation technique used in our single sensor tracking algorithm.

We have shown that at every scan, the track-oriented approach grows new branches for each existing branch of each target. These new branches represent the various possible hypotheses for the different dynamic models for the target and the different possible associations of the target with reports in that scan. In reality, we know that each target should have only one branch corresponding to the true dynamic model for the target and the correct report association.* For simplicity in notation, we will denote the selection of the correct branch at each scan as the *report-to-target assignment* problem even though it includes selection of both the dynamic model and the sensor report. This assignment of reports to targets has to be based on information contained in the measurements of all reports in all scans. One approach for assigning the reports to targets in a scan is to use information unique to that scan (this is commonly referred to as a *single hypothesis algorithm*). However, since a dynamic model links reports from one scan to the next, additional information for assigning reports to targets

* For clarity in the discussions, we assume that a MAP estimate of the continuous-valued state is evaluated.

in a particular scan may be extracted from subsequent scans of reports. An algorithm which makes report-to-target assignments in a particular scan based on information from subsequent scans is referred to as a multiple hypotheses algorithm. The optimal algorithm requires that all available scans of reports be processed prior to making any assignment; the computational requirements of this optimal algorithm, however, grow at an exponential rate. In practice, it is feasible to examine only a finite number of subsequent scans of reports: examination of such a finite but variable number n of subsequent scans is denoted as the *nscan approximation technique*.

A multitarget tracking algorithm that includes the nscan approximation technique has to perform two major steps at each scan. The first major step (step A) is to grow all feasible branches using track hypotheses from the previous scan and sensor reports in the current scan. This is equivalent to postulating multiple report-to-target associations in the current scan. The next major step (step B) is to identify the most likely set of branches for each target in the scan n scans earlier, and prune away the unlikely branches. This is equivalent to resolving the multiple report-to-target associations made in the scan n scans earlier. In fact, the process of selecting the most likely assignment of reports to targets, may be viewed as the nscan approximation procedure.

For the single sensor multitarget tracking algorithm, the steps in the nscan approximation procedure are as follows:

1. Form global hypotheses as combinations of target branches, then identify the most likely global hypothesis.
2. Identify the tracks from each target which are included in the most likely global hypothesis. These represent the set of most likely tracks for each target.
3. Identify the report in the scan n scans earlier, associated with the most likely track of each target.
4. For each target, prune all tracks that do not use the report identified in step 3.

It should be pointed out that this nscan approximation procedure is different from the nscan memory filters defined by Singer [5]. The objective of the nscan memory filters defined by Singer was to merge tracks postulated by individual sensors after examining $(n+1)$ scans of track data. In contrast, the objective of the nscan approximation procedure defined here is to resolve sensor report-to-target assignments after processing n more scans of sensor data.

In the next two sections we will discuss two issues which need to be addressed for extending this single sensor tracking algorithm to the case of tracking with multiple sensors. The first issue is the compensation required to account for the asynchronous reporting rates of the multiple sensors and its impact on the postulation of report-to-target associations in step A. The second issue is the interpretation of the nscan approximation for the case of multiple sensors and its impact on the procedure for resolving report-to-target associations. In each case we will first illustrate the problem, and then show how the tracking algorithm can be modified to account for the problem.

3. COMPENSATION FOR ASYNCHRONOUS REPORTING RATES OF MULTIPLE SENSORS

As we have mentioned in Section 1, surveillance sensors generate target report data on a scan-by-scan basis. The time interval between scans is denoted as the *scan period* and the times at which the individual scans of sensor data are reported

to the tracking algorithm are denoted as the *scan times*. Some sensors (e.g., aero-acoustic sensors) generate all the target reports in a scan at one point in time since all targets are perceived by the sensor at the same point in time. In this case, the *report times* for individual reports in a scan are the same as the scan time. Other sensors (e.g., search radars) generate reports for targets as the radar beam sweeps across the targets, and the times at which the individual targets are perceived are related to the target azimuths relative to the sensor. In this case, the report times for individual reports in a scan may not be the same as the scan time.

For the case of single-sensor tracking, even if the report times do not correspond to the scan time, processing all reports using the scan time will not introduce any significant errors. This is due to the fact that the difference between the scan time and the report time for a particular target remains essentially the same from one scan to the next. The exception arises when a target crosses the scan reference line, in which case the difference in these times for that target changes by a full scan period. In practice, the number of targets that cross the scan reference line is only a small fraction of the total number of targets. Hence, tracking algorithms using returns from a single sensor have satisfactory performance even if they do not compensate for the difference between the scan and report times. Note that the problem associated with targets crossing the scan reference line may be eliminated by dividing a scan into sectors, and processing targets on a sector-by-sector basis. When targets in a particular sector are processed, reports from that sector and the adjoining sectors are also included for potential associations with these targets [6].

For the case of multisensor tracking, however, the situation is quite different. In general, the various sensors will have different scan periods and each will be reporting asynchronously relative to the others. The difference between the report times and the scan time may change from one scan to the next for most targets. The questions that need to be answered in this case are: Do we continue to process sensor reports on a scan-by-scan basis? If we do, then how do we compensate for the non-constant difference between the report times and the scan time from one scan to the next for each target?

The answer to the first question is that we have to continue processing targets on a scan-by-scan basis. This allows us to make the key assumption that each target will give rise to at most one report per scan. This assumption enables us to formulate the report-to-target-assignment problem and implement the straightforward solution defined in subsection 2.1.

Having decided on processing reports on a scan-by-scan basis, there are two methods to account for the difference between report time and scan time. The obvious method is to predict each track to the time of individual reports. This will yield the optimal procedure for step A. However, it will also cause each postulated track to have a unique time tag corresponding to the time of the report that last updated it. The data structure for tracks would have to include a slot for saving this time tag. This method would also require that each track be predicted to the times of all individual reports in each scan.

The alternate method is to predict the reports to the scan times. All targets can then be predicted and updated at the same point in time, i.e., the scan time. Since the dynamic model to be used for predicting the report to the scan times will be a function of the target that gave rise to the report, one should predict the reports based on models for all the postulated targets. In practice, however, only the models for the targets for which the report lies within the gate need to be considered. For the gating operation itself, we do not need to account for the

individual report times. The error in the gate position caused by not accounting for the report times will be proportional to difference between the scan time and the report time. By choosing the scan time to be the middle of the scan interval, this error could have a maximum value of half the scan period. It is straightforward to enlarge the gate size to account for this maximum error.

4. NSCAN APPROXIMATION FOR MULTISENSOR TRACKING

For the case of single sensor tracking, only targets within the coverage region of that sensor can be tracked by the tracking algorithm. Consequently, the interpretation of the nscan approximation is to assign reports to targets in a particular scan after processing n subsequent scans of reports from that sensor alone. For the case of multisensor tracking on the other hand, all sensors do not have the same coverage region and all the tracked targets may not be within the coverage region of all sensors. As a result, the interpretation of the n subsequent scans becomes ambiguous: does it correspond to a total of n subsequent scans, or does it correspond to n subsequent scans from each sensor?

Neither of the above interpretations agrees with the principle behind the nscan approximation. Consider the first interpretation where we count a total of n scans from all sensors. With this interpretation, let us examine a scenario where some of the targets lie within the coverage region of one sensor — say sensor A — in a particular scan, and then remain outside the coverage regions of all sensors that give rise to the n subsequent scans. For this scenario, the tracking algorithm will be forced to assign reports from sensor A to these targets without the benefit of any added information. This corresponds to the use of less than n scans of report information in making report-to-target assignments.

Consider the second interpretation where we count n scans of reports from each and every sensor. With this interpretation, let us examine a scenario where some of the targets lie within the coverage region of several sensors whose scans are interleaved. For this scenario, the tracking algorithm will be forced to wait for n scans of report data from each sensor. This corresponds to the use of more than n scans of report information in making report-to-target assignments.

The problem with the previously stated interpretations is that both of them view the nscan approximation as a technique for allocating sensor reports to targets on a scan-by-scan basis. The correct interpretation is to view the nscan approximation as a technique for allocating a sensor report to a target on a target-by-target basis. Using this interpretation, a sensor report is assigned to a target after processing n subsequent scans of sensor reports for which that target is detected regardless of which sensors provide the reports. Note that this causes different reports in a scan to be assigned to targets at different points in time: in contrast, for single sensor tracking all reports in a scan are assigned to targets at the same point in time. The nscan approximation procedure using this interpretation is provided in the Appendix.

5. RESULTS

As we have mentioned in subsection 2.2, a practical multitarget tracking algorithm uses several heuristics to reduce the computational requirements of the optimal algorithm. The large computational requirements associated with exercising the optimal algorithm on real data prevents us from evaluating empirical results for the optimal algorithm or even the optimal algorithm in conjunction with just a single heuristic. Consequently, we had to use an algorithm which incorporates

these other heuristics to evaluate the two procedures discussed in Sections 3 and 4. The presence of these heuristics will prevent us from evaluating the effectiveness of each of the two procedures in isolation. We can minimize the effects of these other heuristics, however, by normalizing the results.

Data was collected from two civilian Air Traffic Control radars separated by a distance of over 100 nautical miles. Each radar has a scan period of about 10 seconds. We have confined attention to reports in the region between the two radars where they have common coverage. The average number of false reports per scan in this region was 14. Reports from both radars comprise range and azimuth measurements. Aside from compensating the azimuth measurements of both radars for North offsets, we did not have to compensate the measurements for any other calibration errors.

To evaluate the benefit of compensating for the difference between the report time and the scan time, a test target was flown around the vicinity of the line joining the two radars. We compared the errors in the estimated speed of the test target for a period of about 15 minutes with and without the compensation procedure discussed in Section 3. The speed evaluated by the navigational computer on board the test target was used as truth. For the region between the two radars, the difference between the report and scan times for a target may have a value as large as 2.5 seconds. However, for the trajectory flown by the test target, these differences are much smaller. Even for this trajectory, the use of the compensation procedure improved the average error in the estimated speed by 10 percent. The associated computational requirements (measured by CPU time) increased by 13.6 percent.

To evaluate the nscan approximation procedure, targets of opportunity were tracked in the region between the two radars for the same 15 minute period. We examined the effect of the value of n on the continuity of tracks postulated for these targets. The targets of opportunity consisted of commercial airliners and light aircraft. A total of 58 such targets were observed with an average number of 32 targets per radar scan. Target responses received by the radar beacon systems (which are collocated with the two radars) were used to establish the truth tracks. Track continuity for a postulated target was measured by the number of track segments for the target. The number of track segments for a target represents the number of times the track for the target gets dropped and then reinitialized; the larger the number of track segments, the poorer the track continuity.

Figure 4 shows two graphs. The first graph indicates the variation in the average number of track segments for all the postulated targets as a function of the value of n in the nscan approximation procedure. It can be seen that as the value of n increases from zero to three, the number of track segments decreases linearly from 2.02 to 1.42 (30 percent decrease); however, a further increase in the value of n from three to four results in only a marginal decrease (1 percent decrease) in the number of track segments. The second graph indicates the variation of the CPU time (which measures the computational requirements of the nscan approximation procedure) as a function of the value of n . These CPU times have been normalized with respect to the CPU time required where $n = 0$. The plot shows that the CPU time increases linearly with the value of n . If there were no other heuristics in the tracking algorithm, the CPU time would increase exponentially with the value of n ; however, these other heuristics keep the growth in the computational requirements linear.

The improvement in tracking performance resulting from the use of a multiple hypothesis procedure ($n > 0$) compared to

Figure 4. Effect of value of n on the Tracking Algorithm.

a single hypothesis procedure ($n = 0$) can be clearly seen from Fig. 4. The figure also indicates how the choice of n can be used to control the trade-off between the tracking performance and the computational requirements. Our studies indicate that a choice of n between two and three provides good performance and requires computational resources which may be easily satisfied with current computers.

6. SUMMARY

We have reviewed a single sensor multitarget tracking algorithm which has been developed on a firm mathematical basis and which has been successfully demonstrated using real data. We have identified two practical issues that need to be addressed for extending this algorithm to the case of tracking with multiple sensors. These are the issues of compensating for asynchronous reporting rates from multiple sensors, and the implementation of an nscan approximation technique for assigning reports from multiple sensors to targets. We have defined procedures which account for these issues, and provided results using real data which demonstrate the efficacy of these procedures.

Acknowledgements

The authors thank Dr. R. B. Washburn, Mr. M. A. Gerber, and Ms. S. L. Stutsman for supporting the preparation of this paper. Dr. Washburn and Mr. Gerber reviewed a draft of the paper and provided suggestions which improved the clarity of the paper. Ms. S. L. Stutsman helped in evaluating the results provided in Section 5.

REFERENCES

1. Blackman, S.S., Multiple Target Tracking with Radar Applications, Artech House, 1986.
2. Bar-Shalom, Y., "Tracking Methods in a Multiobject Environment," IEEE Transactions on Automatic Control, Vol. AC-23, August 1978, pp.618-626.
3. Reid, D.B., "An Algorithm for Tracking Multiple Targets," IEEE Transactions on Automatic Control, Vol. AC-24, December 1979, pp. 843-854.
4. Kurien, T., and R.B. Washburn, "Multiobject Tracking Using Passive Sensors," Proc. of the 1985 American Control Conference, Boston, MA, June 19-21, 1985, pp. 1032-1038.
5. Singer, R. A., R.G. Sea, and K.B. Housewright, "Derivation and Evaluation of Improved Tracking Filters for Use in Dense Multitarget Environments," IEEE Transactions on Information Theory, Vol. IT-20, No. 4, July, 1974, pp. 423-432.
6. Nagarajan, V., V. Hanuma Sai, and G. K. Chaturvedi, "A New Approach to Scan-to-Scan Correlation and Its Implementation," Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing, April 14-16, 1983, Boston, MA, pp. 711-714.

APPENDIX

Nscan Approximation Procedure for Multisensor Tracking

Let:

T denote the set of all targets postulated by the tracking algorithm at a particular scan,

S_t denote the set of current and past n scan numbers for which the target t lies within the coverage region of any sensor, and

T_r denote the set of targets which are updated in the current scan, and for which a unique report has to be assigned to the target in one of the previous n scans.

The steps in the nscan approximation procedure for multisensor tracking are as follows:

1. Form global hypotheses and identify the most likely global hypothesis.
2. Identify the ordered subset of scan numbers S_r from $S(S = \cup \{S_t\})$ which represent the smallest scan numbers for the set of targets T_r .
3. Identify the largest scan number S_{lr} in S_r .
4. Identify the subset of targets $T_{sr} \in T_r$ for which S_{lr} is the smallest scan and at least one report is associated with these targets in that scan.
5. Identify the track hypotheses included in the most likely global hypothesis and designate the set of reports associated with the targets T_{sr} in scan S_{lr} by R_{sr} .
6. Resolve targets in T_{sr} at scan S_{lr} by pruning all branches for each target in T_{sr} which do not use the reports in R_{sr} at scan S_{lr} .
7. Resolve other targets in T at scan S_{lr} by pruning all branches that use the reports in R_{sr} at scan S_{lr} . (Target t is considered resolved in scan S_{lr} if cardinality of the set of reports associated with the target at scan S_{lr} is less than or equal to 1.)
8. Identify all scans in S that have been resolved. (Note that a scan is resolved when all targets in T have been resolved in that scan.)
9. Remove all the resolved targets from T_r , and remove all resolved scans from S_t .
10. Repeat steps 3 through 9 for the unresolved targets in T_r , until T_r is empty.