12th International Conference on Information Fusion
Seattle, WA, USA, July 6-9, 2009

A Robust Method for Computing Truth-to-Track Assignments

Mark Silbert
Air 4.5.3.3
NAVAIR
Patuxent River, MD.
mark.silbert@navy.mil

Abstract – *Multi-sensor, multi-target tracking is the process of tracking one or more targets given periodic measurement reports of the target locations from one or more sensors. When there are many targets and/or the target density is high, this can be a complex problem. Because of the complexity, it is useful to measure how well the multi-target tracking system is performing. There are many ways to measure the performance but usually no one single measure can encompass the overall performance. Much work has been done on defining useful metrics such as tracking accuracy, track purity, and track continuity. Before any of these metrics can be applied, one must first determine the identity of the target tracks, i.e., the target being tracked by each track. Determining which track corresponds to which target is called the truth-to-track assignment problem. In the past, this assignment has been accomplished using local cost or global cost minimization techniques. While these techniques are useful when using real data, we will show that the assignments it produces can be less than ideal. In the cases where the measurement reports can be identified, such as in computer simulation, we will describe a much more powerful technique, called Report-Identity Based (RIB) assignment that is much more robust and objective. Furthermore, RIB assignment can be applied equally well to both report-to-track and track-to-track tracking systems. The paper will describe how it can be used for all types of tracking systems.*

Keywords: Multi-target tracking performance, Multi-sensor tracking performance, truth-to-track assignment, report-identity based assignment, data fusion tracking performance.

1 Introduction

Multi-sensor, multi-target tracking is the process of tracking one or more targets given periodic measurement reports of the targets. For these types of problems, the true number of targets is usually unknown. When there are many targets and/or the target density is high (*i.e.*, the targets are in close proximity), this tracking can be a complex problem. Because of the complexity, it is useful to measure how well the multi-target tracking system is performing. There are many ways to measure the performance but usually no one single measure can encompass the overall performance. Much work has been done on defining useful metrics, to include such metrics as track accuracy, track purity, and track continuity. Before nearly all metrics can be applied, one must first determine the identity of the target tracks resulting from the tracking system. Namely, for each track produced by the tracking system, one must determine which target is being tracked by that track. Determining which track corresponds to which target is called the *truth-to-track assignment problem*. In the past, these assignments have been accomplished using local cost or global cost minimization techniques. We will describe these techniques and show that in certain situations, the resulting assignments are much less than ideal. It is important to understand the quality of the assignment process. Since most metrics rely on the truth-to-track assignments, if the assignments are wrong, the resulting metrics are likely to be wrong, or at least, suspect at best. This paper will describe the potential problems with these assignment methods and then describe a new technique, called *Report-Identity Based (RIB)* assignment which does not suffer from these problems. We will show how the RIB assignment works and how it results in more believable and objective assignments.

1.1 Two General Types of Multi-Sensor Tracking

When developing a multi-sensor tracking system, one must decide what data from the sensors will be merged. There are two general approaches, as shown in Figure 1. The first approach involves collecting the measurement reports from each sensor and developing the combined tracks, effectively treating all the sensors as one sensor. This type of multi-sensor tracking is called *centralized tracking* or *report-to-track (RTT) tracking*. The second approach is to develop a report-to-track tracker for each sensor that each yields sensor tracks. These sensor tracks are then given to a track merger function that combines the sensor tracks from each tracker that are believed to be tracking the same target. This type of tracker is called *distributed tracking* or *track-to-track (TTT) tracking*. TTT tracking systems are also referred to as *track fusion* systems. In either tracking approach, we refer to the resulting tracks as merged tracks or combined tracks. The performance of the tracking system is based on how well the combined tracks accurately follow the targets being detected by the sensors. Thus, to assess the

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **JUL 2009**	2. REPORT TYPE	3. DATES COVERED **06-07-2009 to 09-07-2009**

4. TITLE AND SUBTITLE **A Robust Method for Computing Truth-to-Track Assignments**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **NAVAIR, , ,Patuxent River,MD**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM002299. Presented at the International Conference on Information Fusion (12th) (Fusion 2009). Held in Seattle, Washington, on 6-9 July 2009. U.S. Government or Federal Rights License.

14. ABSTRACT
Multi-sensor, multi-target tracking is the process of tracking one or more targets given periodic measurement reports of the target locations from one or more sensors. When there are many targets and/or the target density is high, this can be a complex problem. Because of the complexity, it is useful to measure how well the multi-target tracking system is performing. There are many ways to measure the performance but usually no one single measure can encompass the overall performance. Much work has been done on defining useful metrics such as tracking accuracy, track purity, and track continuity. Before any of these metrics can be applied, one must first determine the identity of the target tracks, i.e., the target being tracked by each track. Determining which track corresponds to which target is called the truth-to-track assignment problem. In the past, this assignment has been accomplished using local cost or global cost minimization techniques. While these techniques are useful when using real data, we will show that the assignments it produces can be less than ideal. In the cases where the measurement reports can be identified, such as in computer simulation, we will describe a much more powerful technique, called Report-Identity Based (RIB) assignment that is much more robust and objective. Furthermore, RIB assignment can be applied equally well to both report-to-track and track-totrack tracking systems. The paper will describe how it can be used for all types of tracking systems.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Public Release**	**7**	

performance of the tracking system, one must first determine which target should be assigned to each combined track. The difficulty in making these assignments correctly, which is the truth-to-track assignment (TTA) problem, will be explained in the next section.

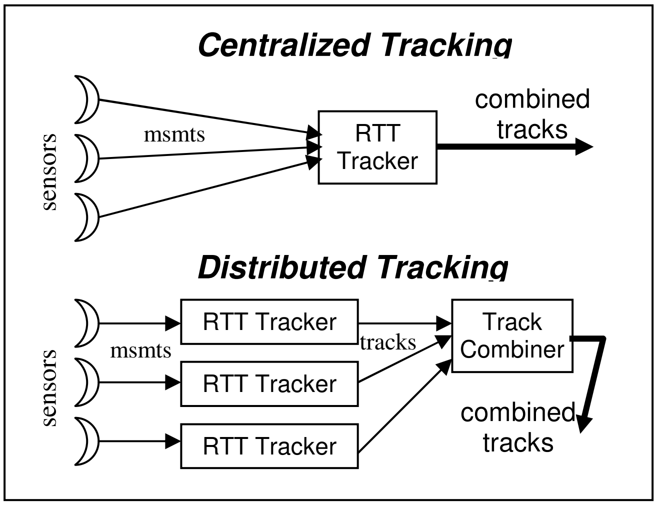

Figure 1. Comparison of centralized tracking to distributed tracking for multi-sensor tracking systems.

2 Truth-to-Track Assignment

To better understand the complexity of the truth-to-track assignment problem, consider the problem illustrated in Figure 2. It shows four targets labeled A, B, C, and D and five resulting tracks labeled 1 through 5. For each track, we must determine which target is being tracked. For this example, the first problem is that there are more tracks than there are targets. Additionally, many of the tracks appear ambiguous. Consider track 1. Initially, it is not clear whether it is associated with target A or B, but by the end of the track, it appears to be tracking B more than A. Should we assign the track entirely to target B? It doesn't seem correct that we should. Track 2 appears to be entirely a false track, so it should not be assigned to any target. Track 3 initially appears to be a false track but then appears to be tracking B. Therefore, it seems we should not assign the entire track to B. Track 4 is clearly tracking C initially but then appears to jump to tracking target D. Should track 4 be assigned to C or D? It seems appropriate that we should initially assign the track to target C and then transition it to target D. But now we need to decide at what point in time the transition of the assignment should occur. Finally, track 5 is tracking the beginning part of target C. So we see that multiple tracks could be tracking the same target. This example illustrates just some of the issues that make the truth-to-track assignment a difficult problem. It should be clear, however, that because of the many ways the tracks can change, the assignments need to be re-calculated frequently over the duration of the scenario. In many cases, the assignments need to be re-calculated as often as after every track update to account for all the possible assignment changes. In the next sections, we will describe how the current approaches deal with these situations and the problems that arise from these approaches.

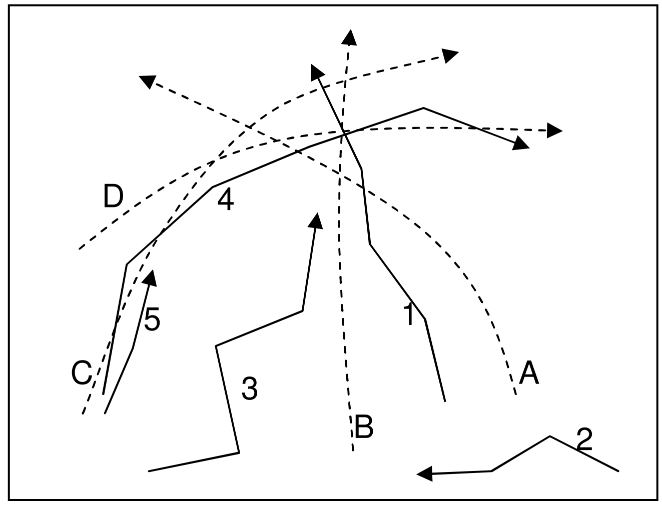

Figure 2. The true target paths are shown as dashed lines, labeled A through D. The resulting tracks are shown as solid lines, labeled 1 through 5. Which tracks should be assigned to which target?

2.1 TTA Using Local Cost Minimization

The simplest and most obvious approach to determining the truth-to-track assignments is by using a *local cost minimization* (LCM) technique such as *nearest neighbors*. In this approach, for a given point in time, each track is assigned to the nearest target. To determine the nearest target, we use ground truth to determine where all the targets were at that time. The process of moving the targets through time to a specified common time is called *time-alignment*. The distance or cost of assigning each track to each of the time-aligned targets is then computed. Each track is then assigned to the target that yields the smallest cost. Gating is usually applied to ensure that no track is assigned to a distant target. If no target falls within the gating distance, the track is labeled a false track.

While the LCM method is efficient and straightforward, it has serious drawbacks. First, since each track assignment is based on the closest target, it will yield overly optimistic estimates of the tracking accuracy. Second, if the sensor tracks are *impure*, i.e., contain reports from detecting other targets or false-alarms, the assignments will likely be unstable over time. Third, if there are several targets nearby, and several tracks to be assigned, there is nothing to ensure the tracks are distributed "fairly" among these targets. Namely, the assignment process should avoid the situation where only a few of the targets end up with all the tracks.

The LCM method for TTA has another important problem that we'll see is not overcome by the global cost minimization methods either. The LCM-based approach can be used to assign sensor tracks, or combined RTT/TTT tracks. However, in the case of a TTT tracking system, the assignments determined for the sensor tracks could differ from the assignments determined for the combined tracks. This differing of assignments is illustrated in Figure 3. The figure shows three targets labeled A, B, and C, all moving together eastwardly. The targets' true paths are shown as light dashed lines. There are two sensors tracking these targets. Sensor 1 has two tracks, S11 and S12, shown as solid lines. Sensor 2 also has two tracks, S21 and S22, shown as heavy dashed lines. Using nearest-neighbors, S11 would be assigned to A, S21 and S12 would be assigned to B, and S22 would be assigned to C. However, suppose the TTT combiner merged S11 and S21 into combined track T1 and merged S12 and S22 into combined track T2. Using nearest-neighbors on the combined tracks would assign T1 to A and T2 to C. But T1 includes S21 which was assigned to B. Similarly, T2 includes S12 which was also assigned to B. Thus the combined track assignments of S12 and S21 differ from their sensor track assignments. Because of the many problems with LCM methods, they are generally not used.

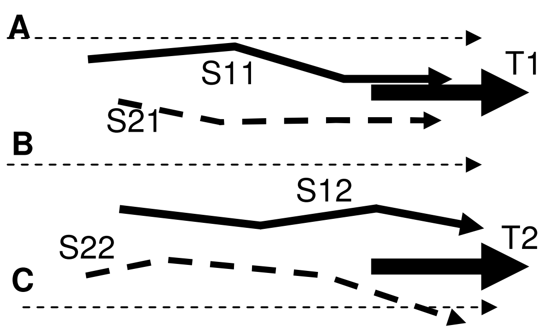

Figure 3. Three targets, A, B, C and two sensors. Each sensor has two tracks. Sensor 1 has sensor tracks S11 and S12. Sensor 2 has sensor tracks S21 and S22. S11 and S21 are merged to form combined track T1. S12 and S22 are merged into combined track T2. Sensor track assignments differ from the combined track assignments.

2.2 TTA Using Global Cost Minimization

Global cost minimization (GCM) is a much more sophisticated method used to determine truth-to-track assignments. Since this method addresses many of the problems with the LCM method, it has become the method of choice for many tracking performance evaluation systems. Furthermore, it is a very effective method when evaluating tracking systems where real data, as opposed to simulated data, is used. We will show, however, that the GCM method introduces a new set of problems.

Whereas the LCM method allows multiple tracks to be assigned to one target, the GCM method enforces a one-to-one assignment.[1] The GCM method begins very similarly to the LCM method. The targets are first time-aligned. Then the distance (or cost) of each track to each of the time-aligned targets is computed. The cost assignments are stored in a table called a cost matrix, where c_{ij} is the cost of assigning track *i* to target *j*. The GCM method then finds the truth-to-track assignments such that the overall cost is minimized, subject to the following constraints:

- Every track is either uniquely assigned to one target or is designated as a "false track".
- Every target is either uniquely assigned to one track or is designated as "not tracked".

	False	Tgt_1	Tgt_2	Tgt_3	…	Tgt_m
No		$\mathbf{c_{01}}$	$\mathbf{c_{02}}$	$\mathbf{c_{03}}$	**…**	$\mathbf{c_{0m}}$
Trk_1	$\mathbf{c_{10}}$	$\mathbf{c_{11}}$	$\mathbf{c_{12}}$	$\mathbf{c_{13}}$	**…**	$\mathbf{c_{1m}}$
Trk_2	$\mathbf{c_{20}}$	$\mathbf{c_{21}}$	$\mathbf{c_{22}}$	$\mathbf{c_{23}}$	**…**	$\mathbf{c_{2m}}$
Trk_3	$\mathbf{c_{30}}$	$\mathbf{c_{31}}$	$\mathbf{c_{32}}$	$\mathbf{c_{33}}$	**…**	$\mathbf{c_{3m}}$
. . .	**…**	**…**	**…**	**…**	**…**	**…**
Trk_n	$\mathbf{c_{n0}}$	$\mathbf{c_{n1}}$	$\mathbf{c_{n2}}$	$\mathbf{c_{n3}}$	**…**	$\mathbf{c_{nm}}$

$$\textbf{\textit{Minimize}} \quad \sum_{\forall ij} c_{ij} x_{ij} \qquad x_{ij} \in \{0,1\}$$

$$\textbf{\textit{Subject to}} \quad \sum_{j=0}^{m} x_{ij} = 1 \quad \textbf{\textit{for i=1,…n}}$$

$$\sum_{i=0}^{n} x_{ij} = 1 \quad \textbf{\textit{for j=1,…m}}$$

Figure 4. Cost matrix and constraints. Each c_{ij} is the cost of assigning track *i* to target *j*. Cost c_{i0} is cost for labeling track *i* a false track. Cost c_{0j} is the cost for assuming target *j* is not being tracked. x_{ij} is 1 if track *i* is assigned to target j, and 0 otherwise.

The cost matrix and the GCM constraints are shown in Figure 4. In order to determine the overall (global) cost, a cost must also be established for designating a track as a false track and for designating a target as not tracked. These

[1] The LCM method can easily be extended to enforce one-to-one assignments by only selecting from the targets that have not already been assigned. This approach gives rise to a class of non-optimal methods called greedy methods. While these methods are typically efficient and may produce good results, they suffer from the undesirable feature that the ordering of how the assignments are made matters.

false-track and not-tracked costs are somewhat problematic. Although they are needed for the GCM methods, there are no formal means for determining them. As a result, these costs are determined through some ad hoc approach.

Since GCM enforces a one-to-one assignment between targets and tracks, it addresses the LCM problem of assigning tracks to targets that are close together. One-to-one assignment also helps simplify many of the metric calculations including track accuracy. However, since GCM methods computes sensor track assignments and combined track assignments (for TTT systems) independently, it still suffers from the differing assignment problem, mentioned earlier.

The GCM also raises new problems. First, the computations required to find the truly minimal global cost grows exponentially with the number of tracks and targets. Efficient algorithms (e.g., JVC, Lagrangian Relaxation, and Munkres) have been developed to find near-minimal solutions in much less time, but the process is still computationally expensive.[2] Second, finding this minimum cost solution provides no assurance that the resulting assignments are correct. In fact, there is no reason to believe, in general, that the lowest cost is tied in any way to the correct assignments. Keep in mind that this lowest cost assignment depends on how the ad hoc false-track and not-tracked costs were computed as well. Third, by enforcing one-to-one assignments, it is possible in certain cases to get very poor TTAs, especially when false tracks exist. An unfortunately placed false track can cause a cascade of wrong assignments. This is illustrated in Figure 5. The figure shows three targets labeled A, B and C. The outer circle around each target is the gate size. Only those tracks within the gate can be assigned to the target. The tracks are denoted as diamonds, labeled *a*, *b* and '*?*'. The '*?*' denotes a false track. The correct assignments are track *a* to target A, track *b* to target B, and the false track unassigned. However, since GCM methods impose a one-to-one assignment, it will do what it can to get the three tracks assigned to the three targets. In this example, it can only be accomplished by forcing the false track to be assigned to target A, which then causes track *a* to be assigned to target B, and track *b* assigned to target C.

Enforcing one-to-one assignments causes another problem with GCM methods. Suppose there are m targets and n tracks. In many cases, tracking systems generate more tracks than there are targets, so n is often larger than m. When this occurs, GCM methods will select at most m tracks to assign and the rest of the tracks are simply ignored. It does not consider the case that there may be targets with multiple tracks. Furthermore, GCM will pick the best (i.e., closest) of the redundant tracks which unwarrantedly biases the accuracy metrics, just as the LCM methods just discussed.

The last problem with GCM methods stems from their computational complexity. Recall that because of the dynamics of the tracking problem, the assignments should be calculated quite frequently over the duration of the scenario. However, since GCM methods are computational, the number of times that the assignments are calculated must be substantially reduced. Thus, a set of assignments could persist longer than they should yielding incorrect measures of performance.

Now that the deficiencies of the LCM and GCM methods are understood, we will now introduce the report-identity based method and show that it does not suffer from these problems.

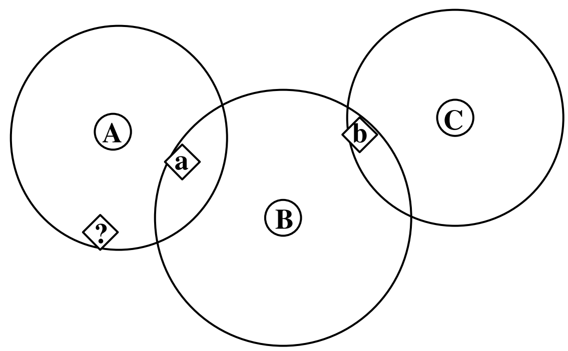

Figure 5. Illustrating the cascade effect of wrong assignments that could occur using GCM methods. A, B, and C are the true target locations. The circles around the targets are the gates size. The tracks are designated as diamonds. The track labeled "?" is a false track. The correct assignments are track *a* to target A and track *b* to target B.

2.3 Report Identity Based (RIB) Assignment

Instead of basing the truth-to-track assignments on some type of cost minimization, the report identity based method uses the identity of the track constituents. The simple idea is that the identity of a sensor track is based solely on the identity of the sensor measurement reports that make up the track. This means, of course, that we must know the identity of each report. The identity of a report is the target that was truly detected. In simulation, this information is usually known. When it is not known, it may still be deducible using the (known) measurement error of the sensor to determine which target was most likely detected. If more than one target falls within the measurement error, we simply include all those targets as the possible identity of the report. The RIB assignment can handle multiple identities for a measurement report. For the remainder of this paper, we'll assume the identity of each measurement report is known, even if it has multiple identities.

[2] When the number of target and tracks is very large where only near minimal costs are found, GCM approaches could suffer from the undesirable feature that the ordering of the assignments matters, just as greedy methods.

It should be clear that if all the measurement reports that make up the sensor track are from detecting the same target, then the identity of the sensor track is simply that target. These types of tracks are called *pure* tracks. A sensor track is called *impure* if it is made up of measurement reports that do not all have the same target identity. Figure 6 illustrates a pure and an impure track. We will explain how the RIB assignment method handles impure tracks.

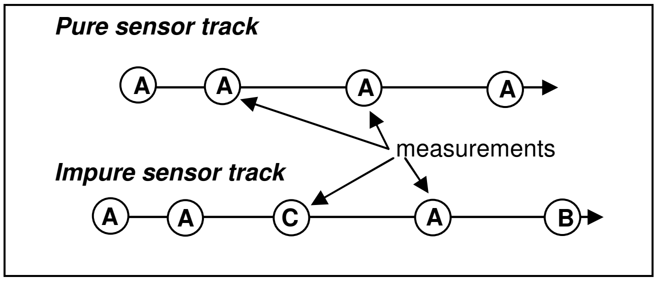

Figure 6. A pure and impure sensor track.

2.3.1 Handling Impure Sensor Tracks

Consider the example in Figure 7. The figure shows a sensor track composed of four measurement reports. The first and last reports are from unambiguously detecting target A. The second report could be from either target A or B and the third report could be from targets A, B or C. A straightforward scheme for deciding on the identity of the sensor track would be to simply count up the "votes" and assign it to the target with the majority of the votes. In this case, target A would have the majority of the votes, so the sensor track would be assigned to target A.

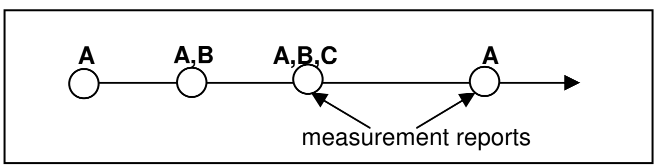

Figure 7. An impure sensor track with measurement reports with multiple identities.

While this voting approach seems reasonable, it does suffer from a key problem: although all the votes were treated equally, the ordering of the votes (i.e., measurement reports) matters. Recent votes should carry more weight than older votes. This problem can be corrected using any reasonable time-weighting function. A good method for time-weighting the votes, is to think of the votes as having a "half-life." The half-life is that exponentially decaying amount of time until the vote is worth half. For example, if we believe a vote decays to half of a vote in 5 minutes, then solving for decay constant, λ in equation 1, yields the rate in equation 2.

$$v = e^{-\lambda t} \tag{1}$$

$$\lambda = \frac{\ln 2}{5} = \tfrac{1}{5}\ln 2 \approx 0.1386 \ min^{-1} \tag{2}$$

Thus the weight of a vote reduces by half as the report ages every 5 minutes with a decay constant of 0.1386 per min. The assignment of the sensor track to a target is then accomplished by adding the time-weighted votes and employing majority rule.

Notice that unlike the LCM/GCM methods, there is no need to time-align if the identities of the measurement reports are already known.

2.3.2 Assigning Combined Tracks from TTT Tracking Systems

Both the LCM and the GCM assignment methods treat the assignment of sensor tracks and the assignment of combined tracks independently. They simply ignore the TTAs for the sensors tracks when determining the TTAs for combined tracks. The RIB assignment method ensures the targets assigned to the combined tracks are consistent with the sensor track assignments. This consistency occurs because the combined track assignments are based on the sensor track assignments.

2.3.3 Handling Impure Combined Tracks from TTT Tracking Systems

As discussed earlier, in a TTT tracking system, the combined tracks are composed of one or more sensor tracks. If a combined track is made up of sensor tracks that are pure and all are assigned to the same target, then it is trivial to determine the assignment of the combined track: it is simply that target. However, a dilemma occurs when determining the assignment of a combined track made up of one or more impure sensor tracks. Should the assignment of the combined track be based on the identities of the sensor tracks, or on the identities of the measurement reports that make up the sensor tracks? This problem is illustrated in Figure 8. The figure shows three impure sensor tracks that are grouped together to form a combined track. To simplify the problem, we'll ignore the time-weighting for now and assume all reports are equally weighted. The first sensor track, with nine votes for target A and three for target B, would be assigned to target A. The second sensor track would be assigned to target B since it has three votes for B, but only one vote for A. The third sensor track, with four votes for B but only two votes for A, would also be assigned to target B. So if we base the assignment of the combined track using the majority at the sensor track level, we would assign the combined track to B (since two votes for B versus one vote for A). However, if the report votes are added together, then target A get 12 votes while B gets only ten votes. Thus, if the voting for the combined track is based on the report votes, the combined track would be assigned to target A. A contradiction occurs. Which target should the combined track assignment to? Note that if the voting is at the sensor track level, then we are implicitly treating all sensor tracks as equally weighted. But it is unfair to weight the sensor tracks equally since some tracks may have longer histories than others. A track with more history should carry more weight. In the example in Figure 8, the first sensor track has much more history than the other two sensor tracks. A sensor track with more history has more measurement reports. Thus, if the voting is based on the reports, the sensor tracks with longer history will naturally get more weight. Therefore, voting at the report level is used to determine both the sensor track and the combined track assignments. Using the identity of the reports to determine the identity of all tracks is why the assignment method is called Report Identity Based (RIB) assignment.

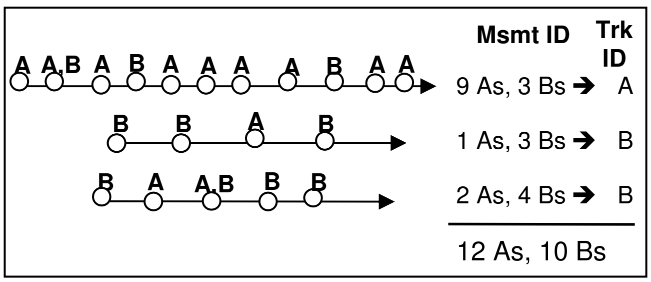

Figure 8. Three impure sensor tracks merged into one combined track.

3 Benefits of the RIB Assignment

There are many benefits of the RIB assignment method over the LCM/GCM approaches, particularly for track-to-track tracking systems. In this section, some of the lesser obvious benefits of using the RIB assignment method for TTT tracking systems will be discussed.

3.1 Avoiding the Ripple Effect

We begin this section re-visiting the problem discussed earlier and shown in figure 5. Recall we showed that an unfortunately placed false track could case a rippling of incorrect assignments using GCM methods. Since the RIB assignment method is based on the identity of the reports that make up the sensor track, the false track would be recognized as such and thus, would not be assigned to any target. Similarly, the remaining two tracks would be assigned to the correct targets, using the identity of each sensor track, instead of some distance or cost. No rippling of incorrect assignments occur using RIB method.

3.2 Detecting Inconsistent Combined Tracks

One of the additional benefits of the RIB assignment method is the ability to check whether a grouping of sensor tracks into a combined track is consistent. Just because it is possible to assign a combined track to a target, does not mean that the sensor tracks should have been grouped together in the first place. For example, suppose a TTT tracking system grouped a sensor track made up of only reports from detecting target A with another sensor track made up of only reports from detecting target B. Using the RIB assignment method, the combined track would be assigned to either A or B. But there is something fundamentally wrong with this combined track. It shouldn't have been grouped at all. The LCM/GCM methods do not detect this problem and as a result would not discuss it. The RIB assignment method, on the other hand, can tell that this grouping is inconsistent and can report it. Thus, the RIB assignment adds another dimension to measuring the performance of track-to-track tracking systems. At a minimum, a grouping of sensor tracks must have a target in common for the grouping to be consistent. Tighter rules for consistency can be used to ensure the sensor tracks assignments share a common target as well. In any case, the RIB assignment method can easily check for this occurrence.

3.3 Avoiding False Rewards

As previously mentioned, LCM and GCM methods ignore the sensor track assignments when assigning the combined tracks. This means any feature data that is associated with the sensor track (which can help determine the track's identity) is also ignored. This is especially problematic in track-to-track tracking systems. Since there is no consistency checking in those methods to ensure that the identity of the combined track agrees with the identity of the sensor tracks that make up the combined track, a target could actually be assigned to a combined track that is made up of sensor tracks that aren't even tracking that target. This is illustrated in Figure 9.

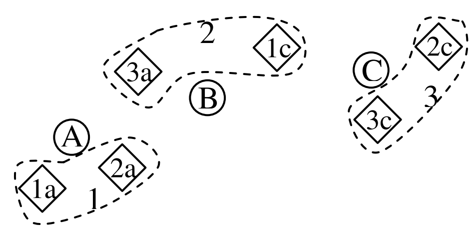

Figure 9. Giving credit for tracking target B when it's not even detected.

The figure shows three targets A, B, and C. There are three sensors but each of them are only tracking A and C. Sensor 1 has tracks 1a and 1c, tracking targets A an C respectively. Sensor 2 has tracks 2a and 2c, tracking targets A and C respectively. Likewise, sensor 3 has tracks 3a and 3c, tracking targets A and C respectively, as well. Target B is undetected by all three sensors. Suppose the TTT combiner (mistakenly) formed the sensor tracks into the three combined tracks (shown as dashed loops). Using the RIB assignment method, combined track 1 would be assigned to target A and combined track 3 would be assigned to target C. Combined track 2 (composed of sensor tracks 3a and 1c) would be recognized as an inconsistent grouping, but could assign it. It would be assigned to either A or C, depending on their relative weighting, but certainly not to target B. Using GCM (or LCM) methods, combined track 2 would be assigned to target B even though it's composed of sensor tracks for targets A and C. The resulting TTAs using GCM methods would make it appear that the tracking system is tracking all three targets. The TTT tracking system made a mistake and yet the GCM assignment would end up rewarding for it. Worse yet, if a competing TTT tracking system did correctly merge the sensor tracks into combined tracks, the GCM assignment would make it appear inferior to the first one because its coverage would be lower (since it would show not tracking B).

3.4 Dealing with Sensor Bias

One of the presumed benefits of GCM assignment methods is that they can handle some sensor biases. The idea is that if a sensor has a bias, then all its sensor tracks will be shifted about equally. And since the GCM methods seek to reduce the overall cost, the minimal cost will coincide with the correct assignments. This argument, however, breaks down when dealing with multiple sensors in a TTT tracking system. Consider the example shown in Figure 10. The figure shows two sensors each tracking two targets labeled A and B. The sensor tracks from sensor 1 are labeled 1a and 1b. The sensor tracks from sensor 2 are labeled 2a and 2b. Both sensors have a bias. The bias from sensor 1 is northeasterly and the bias from sensor 2 is southwesterly. Since GCM methods find the overall minimum cost, it will correctly assign the sensor tracks to the proper targets. But suppose a TTT tracking system, unaware of the biases, fails to merge these sensor tracks concluding the two sensors are each tracking two different targets. The TTT tracking system would then report all four sensor tracks as four combined tracks, each containing a single sensor track. The combined tracks are labeled 1, 2, 3, and 4. Using a GCM method to assign the four combined tracks would result in assigning combined track 1 to target B and combined track 4 to target B (combined tracks 2 and 3 would be ignored). But as you can see from the figure, this is exactly reversed from the assignments of the single sensor tracks that each is composed from. Worse yet, since the two selected combined tracks are closer to the incorrectly assigned target, it appears the TTT tracking system has increased the tracking accuracy of the targets. Here again, the GCM assignment has falsely rewarded a faulty tracking system.

The RIB assignment, on the other hand, would correctly assign combined tracks 1 and 3 to A and combined tracks 2 and 4 to B, thereby avoiding this problem.

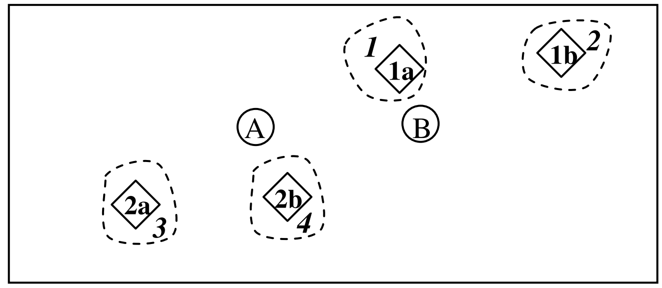

Figure 10. GCM methods causing reward for bias error.

4 Conclusions

As explained, the first step to evaluate the tracking performance of a multi-sensor, multi-target tracking system is to have the correct truth-to-track assignments. We introduced the RIB assignment method as a more correct and robust means to compute these assignments. It was shown that the RIB assignment method does not suffer from the problems that the standard LCM and GCM methods for assignments have. Furthermore, the RIB method is elegantly simple and efficient. Although the approach is primarily geared toward simulation environments where the identity of the measurement reports is usually available, the method can be applied to real data situations as well, provided the measurement errors of the sensors can be characterized. It was shown that in addition to providing robust assignments, there were other key advantages. The RIB assignment method can deal with impure tracks and ambiguous measurement report identities. Furthermore, the method ensures consistency between the sensor track assignments and the combined track assignments in track-to-track tracking systems. A new dimension in consistency checking for these TTT tracking systems was also explained. Overall, the RIB assignment method will help ensure that whatever metrics are used to evaluate the tracking performance, they will be based on objective, more believable assignments.

References

[1] Blackman, S. and Popoli, R., *Design and Analysis of Modern Tracking Systems*, Artech House, Boston, 1999.

[2] Bar-Shalom, Y., Li, X. R., and Kirubarajan, T., *Estimation with Applications to Tracking and Navigation*, Wiley, Boston, 2001.

[3] Silbert, M. and Burroughs, A. "Measuring the Performance of Track-to-Track Fusion Systems," In *Proceedings of Aerospace Conference*, Volume 3, IEEE, March 2004.

[4] Klein, L., *Sensor and Data Fusion*, SPIE – The International Society for Optical Engineering, 2004.