Methods of Multisensor Tracking

C. M. Rogers
QinetiQ Ltd

Abstract Multisensor multitarget tracking is a complex problem that has only recently received much attention. Several approaches for combining information between sensors may be taken, consisting of various association and fusion architectures. This paper examines measurement and track association and fusion, and the inherent advantages and disadvantages, and concludes that track association techniques combined with measurement fusion are very promising but need significant further research to gain the maximum advantage of the multisensor capability.

1. Introduction

The aim of a tracking system is to take measurements of an object using some form of sensor and to use them to estimate parameters of interest, such as the position and velocity of the object. Tracking is often performed with a Kalman filter, which provides the optimum (lowest variance) estimate of parameters given gaussian processes. The track formed by the filter typically has a much lower variance than the variance of the raw measurements, and allows predictions to be made about the object parameters at some future time.

Increasingly it is becoming popular to use multiple sensors in a tracking system. In this way it is possible to overcome the shortcomings of some sensors with the benefits of others (Radar providing range and range rate information to an infrared system is one example), and also to achieve multiple measurements of parameters (Radar can also measure angle). The information from the separate sensors then needs to be combined.

In a multisensor tracking system, several approaches may be taken for combining information from individual sensors to provide a better estimate of the parameters to be tracked[1]. The first of these is the measurement fusion approach, where measurements from each sensor are combined to form centralised "system" tracks. A second possible approach is track fusion, where each sensor forms its own local tracks, and then these tracks are combined to form centralised system tracks. The question then arises as to how to chose which information from each sensor should be combined.

2. Measurement Fusion Approach

Consider first the measurement fusion approach. A sensor j produces a number N_j of measurements. It is assumed that there are already in existence a number M of system tracks, whereby M is not necessarily the same as N_j. Choosing which information to combine is then the familiar problem of measurement to track allocation. The term allocation is used here instead of association as the measurement is allocated to that track and used in some way to continue that track (Figure 1).

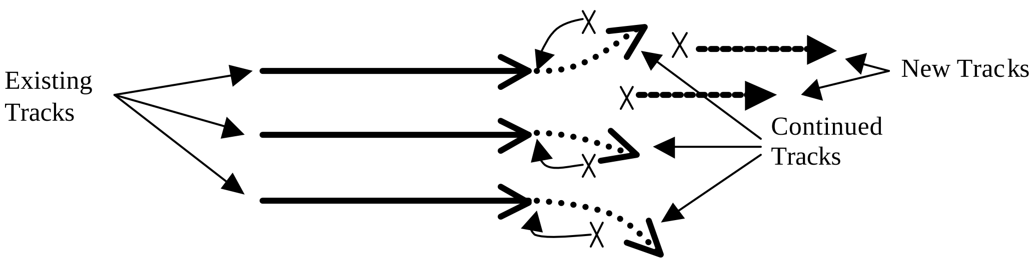

Figure 1 Measurement (X) to track allocation

Many methods have been derived for performing this measurement to track allocation process, with the current most popular being numerous variations on probabilistic data association filters (PDAF's) and multiple hypothesis tracking filters. Key desirable features for the methods are the ability to initiate, continue and delete tracks and fast computation. These features are readily available when only one sensor is providing measurements[2], but when multiple sensors are used, the complexity of incorporating such features increases dramatically.

2.1 Multisensor Measurement Fusion Approach

In a multisensor approach, it is necessary to associate measurements with each other before allocating them to the system tracks (Figure 2). This is not the same as associating the measurements directly with the tracks, as a problem arises when a new track is started from a measurement from one sensor (Figure 3). If this happens, all the measurements from the other sensors need to be re-evaluated to see whether they should be allocated to this new track in preference to the previously allocated old track.

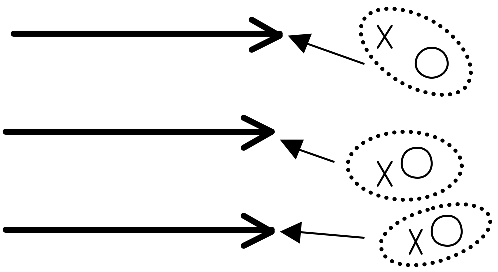

Figure 2 Measurement (*X*) to measurement (*O*) association, then allocation

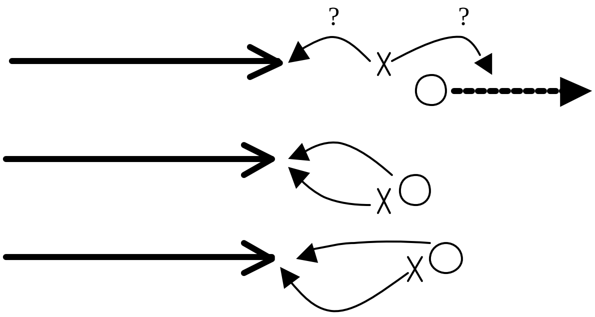

Figure 3 Measurement (*X*) to track allocation, then measurement (*O*) to track allocation

As a general rule, the more noisy information is, the more difficult it is to determine associations[3]. As measurements are inherently more noisy than tracks, measurement association becomes more difficult than in the measurement to track allocation problem with a single sensor. This can be seen diagrammatically if noise is represented by a circle around a measurement. With more noise it is likely that more circles will overlap and so more associations are possible, leading to less certainty in the final decision (Figure 4).

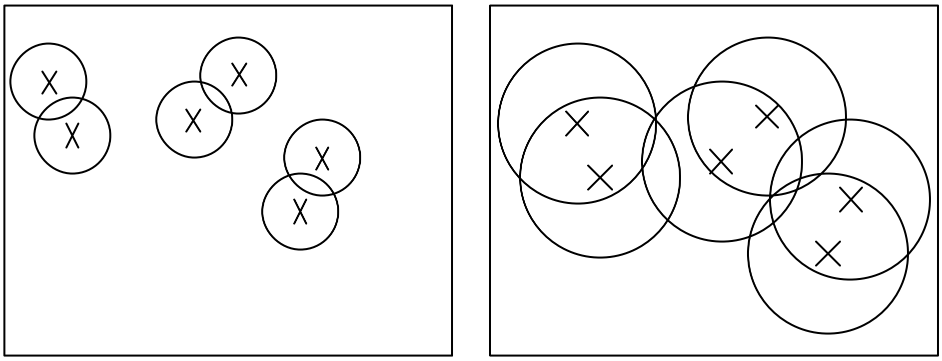

Figure 4 Low noise simple association and high noise complex association

Another problem arises if measurements are associated incorrectly, as the system track will be corrupted by the incorrectly allocated measurement. This corruption will propagate along the track for as long as the track exists, even if measurements from the correct object are allocated after this point. It is impossible to remove the corruption unless the track is deleted and then started again (Figure 5). The re-initialisation of the track will lose the benefit of the track forming, namely reduced noise, so the track accuracy will be compromised.

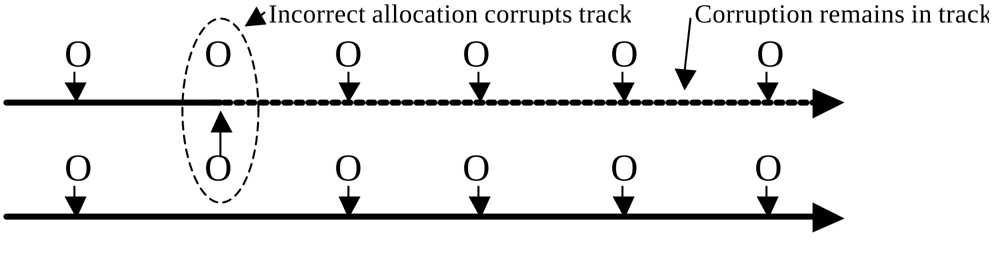

Figure 5 Track corruption through incorrect allocation

However, if these problems can be overcome, and as there is no loss of information in the process, measurement fusion does give the most accurate means of tracking.

3. Track Fusion Approach

Consider now the track fusion approach. Measurement allocation is still performed, but this time within a particular sensor, meaning that the problems and complexities inherent in multisensor measurement allocation are removed. Due to the relatively low track noise, it is more likely that measurements will be correctly allocated to tracks than when measurements are associated with measurements as mentioned in section 2.1. Once each sensor has tracks, these are associated with each other by some means, potentially the same means used for the measurement to measurement association process[4]. Associated tracks are then fused to provide the system tracks (Figure 6).

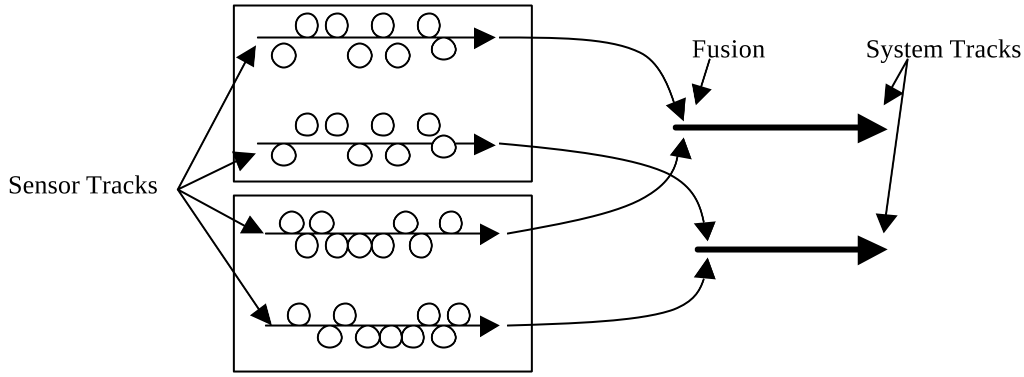

Figure 6 Track fusion approach

The main disadvantage in this approach is that information is lost in the sensor track forming stage so that the system tracks are not as accurate as they would be had they been made up from the correct measurements directly.

However, as the tracks are less noisy than the raw measurements the association process becomes easier. There is less uncertainty as to the association between tracks. Also, if the system track retains no history, but is recreated after each association process from the fusion of the sensor tracks, there is no worry about permanently corrupting it with incorrect measurements. If tracks are incorrectly associated at one time step, the corruption of the system track only remains until the tracks are correctly associated at a later time step. There is no residual corruption in the system track (Figure 7).

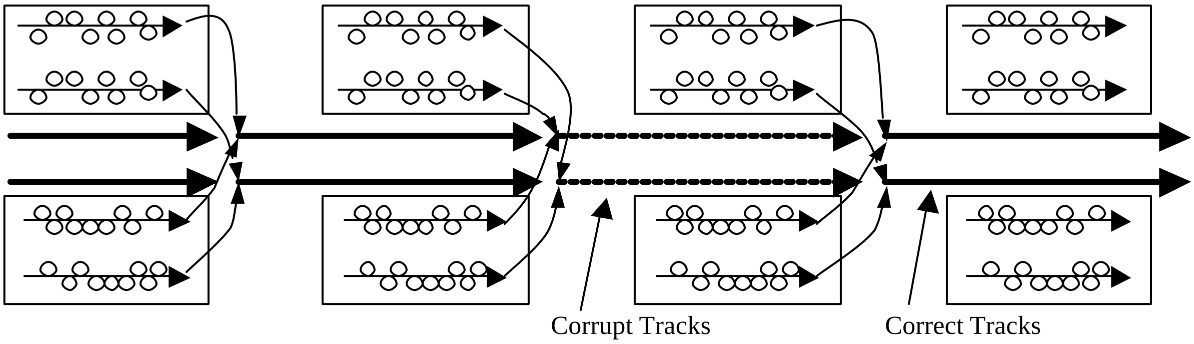

Figure 7 Effect of track corruption in track fusion approach

A third option exists, whereby it is possible to retain the association accuracy given by track association with the tracking accuracy of measurement fusion. Measurement to track allocation is performed local to the sensors as before so that each sensor maintains its own tracks. These track are then associated with each other through some method. Once the track association has been performed, rather than fusing the sensor tracks to form system tracks, the measurements that have been allocated to the associated tracks are fused through measurement fusion (Figure 8).

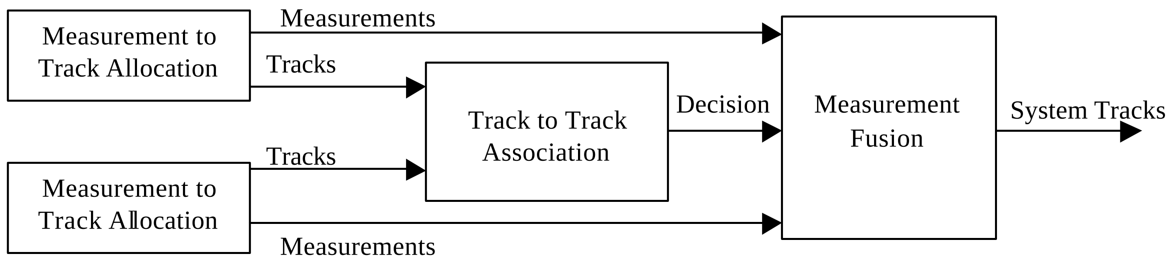

Figure 8 Track to track association with measurement fusion

It is therefore possible to retain the best of both worlds whilst removing some of the problems from each individual method. However, the potential for lasting system track corruption remains if, for example, measurements are incorrectly allocated to tracks within the sensor or tracks are incorrectly associated. If the latter occurs, it can be relatively easily detected. This is done by keeping a record of the last association result and comparing it with the current association result. If the previously associated tracks are deemed to no longer associate then it is clear that at some point the tracks have been incorrectly associated. It may then be necessary to re-initialise the system track with a track fused track, whereby it is assumed that the most current association is correct, when an error is detected. It is important to note that this re-initialisation is not as serious as in the measurement fusion approach. As the sensor tracks have good accuracy, good system track accuracy will be maintained through fusion. If the measurement allocation process is incorrect it is very difficult to detect.

4. Conclusions

The advantages and disadvantages of three methods for choosing and combining information from multiple sensors for the purpose of tracking multiple objects have been discussed.

It has been shown that whilst measurement fusion provides the most accurate estimate of parameters to be tracked, it is also the most complex method for which to implement some association process and arguably the most unreliable for association due to the noisy nature of the measurements.

This is in contrast to the track fusion approach which takes advantage of the low noise tracks from each sensor for the association process and allows system track corruption to be limited to the time between association steps if no system track history is retained. However, parameter estimate accuracy is lost as information is lost at the sensor track forming stage.

A third option of track association combined with measurement fusion has been presented which utilises the advantages of both previous methods, maintaining the tracking accuracy of measurement fusion with the association accuracy of track association. Whilst it can suffer from system track corruption, in many cases this can be detected and resolved.

There have been many studies made about measurement association and fusion, but far less research has been undertaken into track association and fusion. It is believed that it is necessary for this to be done to obtain methods for achieving the maximum advantage from the use of multiple sensors.

In order to test these various methods a test bed is required. This will need to provide a useful environment in which to test the association procedures and the tracking accuracy of the tracking and fusion system. It will probably be necessary to simulate realistic sensor measurements and to implement tracking filters. It is anticipated that the main thrust of this research will be towards track association, but comparisons to measurement association will need to be made.

References

[1] Hall D L, Llinas J; An introduction to multisensor data fusion; Proceedings of the IEEE, vol 85, no 1; Jan 1997

[2] Musicki D, Evans R, Stankovic S; Integrated probabilistic data association; IEEE Transactions on automatic control, vol 39, no 6; June 1994

[3] Zhou B, Bose N K; Multitarget tracking in clutter: fast algorithms for data association; IEEE Transactions on aerospace and electronic systems, vol 29, no 2; April 1993

[4] Stone L D, Williams M L, Tran T M; Track to track association and bias removal; Proceedings of SPIE, vol 4728; 2002

The author is an EngD student at the Dept. of Electronic and Electrical Engineering, University College London

The work described in this paper has been funded by the EGC Domain of the MOD CRP.

The author also wishes to acknowledge the contributions of colleagues within QinetiQ.