

Online Sensor Registration

Jaco Vermaak¹, Simon Maskell², and Mark Briers^{1,2}

¹Cambridge University Engineering Department, Cambridge, UK.

²QinetiQ Ltd., Malvern Technology Centre, St Andrews Road, Malvern, Worcs, UK.
jv211@eng.cam.ac.uk, s.maskell@signal.qinetiq.com, m.briers@signal.qinetiq.com

Abstract—In a multi-sensor scenario, accurate data fusion is best achieved by processing the measurements from all the sensors at a fusion node to produce tracks. However, inaccuracies in the position and/or rotation of the sensor can lead to “ghost” tracks, particularly when the sensors are not co-located. This paper presents a framework which models the uncertainty over the sensors’ registration parameter (e.g. position and rotation) and discloses an unscented implementation technique (other methods based on particle filters can be accommodated within our framework), where each sensor self-localises using targets of opportunity. The aim is to solve the sensor registration problem whilst adding minimal overhead to an existing tracker, which is facilitated by making the standard assumption that the state of the joint target factorises over the individual targets.

TABLE OF CONTENTS

1	INTRODUCTION
2	A GENERAL FRAMEWORK FOR MTT PROBLEMS
3	A FRAMEWORK FOR SENSOR REGISTRATION
4	MIXTURE (UNSCENTED) KALMAN FILTER
5	CONCLUSIONS
6	ACKNOWLEDGEMENTS

1. INTRODUCTION

The objective of this paper is to develop a framework in which to perform sensor registration in the context of Multi-Target Tracking (MTT)[2]. Sensor registration, also known as “calibration” in the vision community[3], involves determining, for example, the location and orientation of all the sensors alongside the target states. The constraint we place on such a framework is that it should have as little as possible impact on existing MTT strategies. Put differently, it is required that the proposed framework should combine with existing MTT strategies with no or little modifications to the MTT algorithms. We defer simulation results to a future jour-

nal publication (in preparation), which will apply our algorithm to the self-calibration of a set of webcams being used to track “markerless” people (in a computer vision context), and so we merely suggest a mechanism (namely an unscented approximation) to performing inference. The novelty of our approach is deemed obvious from existing techniques within the scientific literature; see references [5], [7], [8] for example. With respect to reference [5], what we are doing is (implicitly) equivalent to this method of inflating the measurement noise to reflect the uncertainty over the parameters of the model, and then estimating these parameter values. However, using our approach, this uncertainty can reduce with time as the data provides more information about the value of the parameter.

We will first present a general framework that describes most existing MTT tracking algorithms, and then show how this framework can be extended to accommodate the sensor registration problem.

We will use the following notation in what follows:

- N is the number of targets.
- x_t^i is the state of the i -th target at time t .
- $x_t^{1:N}$ is the joint state of all the targets at time t .
- y_t are the measurements at time t .
- $y_{1:t}$ are the measurements from the first time step up to and including time t .
- a_t^i is the target to measurement association variable for the i -th target at time t .
- $y_t^{a_t^i}$ is the measurement that is associated to the i -th target at time t .
- θ_t is a vector unknown parameters that is comprised of the sensor parameters, that is, the joint state of all of the (multi) sensor parameters.

Since this paper intends to develop a probabilistic framework whose terms can be approximated using *any* standard data association technique and *any* standard filtering technique (including particle filters) we present a large amount of probabilistic detail. This is necessary to preserve the generality of the approach. The implementation details that we do provide in section 4 are

intended to ease the burden on the general reader and allow a more efficient implementation of the framework developed.

The format of the paper is as follows: in section 2, we provide a probabilistic overview of the MTT problem. In the section that follows, we discuss a mathematical framework for sensor registration and as stated provide implementation details, using the (mixture) unscented Kalman filter as an inference vehicle in section 4. Finally, we conclude in section 5.

2. A GENERAL FRAMEWORK FOR MTT PROBLEMS

For online MTT the distribution of interest is the posterior density function $p(x_t^{1:N}|y_{1:t})$, and the objective is to estimate this distribution recursively in time. Most practical MTT algorithms make the assumption that this distribution factorises over the individual targets, *i.e.*

$$p(x_t^{1:N}|y_{1:t}) \approx \prod_{i=1}^N p(x_t^i|y_{1:t}). \quad (1)$$

This assumption is largely historical, and facilitates a computationally feasible extension of single target tracking strategies, such as the Kalman Filter (KF), to the multi-target scenario. From an estimation perspective it leads to a significant reduction in complexity: searching the product of marginal spaces can be done more efficiently than searching the joint space. Practically, it means that the computational cost to perform MTT scales linearly with the number of targets (MTT data association algorithms permitting), which is important if the number of targets could be in the 10s, or possibly 100s. In contrast, treatment of the full joint posterior scales exponentially with the number of targets.

The recursive update of the target distribution proceeds in two steps: a prediction step and a filtering, or update, step. The prediction step proceeds independently for each of the targets as

$$p(x_t^i|y_{1:t-1}) = \int \underbrace{p(x_t^i|x_{t-1}^i)}_{\text{dynamics}} \underbrace{p(x_{t-1}^i|y_{1:t-1})}_{\text{previous filtering}} dx_{t-1}^i. \quad (2)$$

Due to the association uncertainty the filtering step cannot proceed independently for each of the targets. By introducing the association variable the updated target

posterior can be expanded as

$$\begin{aligned} p(x_t^i|y_{1:t}) &= \sum_{a_t^i} p(x_t^i, a_t^i|y_{1:t}) \\ &= \sum_{a_t^i} \underbrace{p(a_t^i|y_{1:t})}_{\text{association posterior}} \underbrace{p(x_t^i|a_t^i, y_{1:t})}_{\text{conditional filtering}}. \end{aligned} \quad (3)$$

Thus, the target posterior is a discrete mixture, with one mixture component for each valid association. The mixture components comprise the target posteriors conditional on the valid associations, and combine in proportion to the corresponding association posteriors. Using Bayes' rule and the standard state-space modelling assumptions the conditional filtering distribution can be developed as

$$\begin{aligned} p(x_t^i|a_t^i, y_{1:t}) &= \frac{p(y_t|x_t^i, a_t^i, y_{1:t-1})p(x_t^i|a_t^i, y_{1:t-1})}{\int p(y_t|s_t^i, a_t^i, y_{1:t-1})p(s_t^i|a_t^i, y_{1:t-1})ds_t^i} \\ &= \frac{c(y_t^{C(a_t^i)})p(y_t^{a_t^i}|x_t^i)p(x_t^i|y_{1:t-1})}{c(y_t^{C(a_t^i)}) \int p(y_t^{a_t^i}|s_t^i)p(s_t^i|y_{1:t-1})ds_t^i} \\ &= \frac{p(y_t^{a_t^i}|x_t^i)p(x_t^i|y_{1:t-1})}{\int p(y_t^{a_t^i}|s_t^i)p(s_t^i|y_{1:t-1})ds_t^i}. \end{aligned} \quad (4)$$

In the above $p(y_t^i|x_t^i)$ is the likelihood model for a target-originated measurement y_t^i , and $c(\cdot)$ is the clutter likelihood model. In the expansion above this model is assumed to apply to all the measurements in y_t that are not associated to the target in question. The indices of these measurements are denoted by $C(a_t^i)$. Thus, the conditional filtering distribution can be expressed as

$$p(x_t^i|a_t^i, y_{1:t}) = p(x_t^i|y_{1:t-1}, y_t^{a_t^i}) = \frac{\overbrace{p(y_t^{a_t^i}|x_t^i)}^{\text{target likelihood}} \overbrace{p(x_t^i|y_{1:t-1})}^{\text{prediction}}}{\underbrace{p(y_t^{a_t^i}|y_{1:t-1})}_{\text{predictive likelihood}}}, \quad (5)$$

where the predictive likelihood in the denominator is given by

$$p(y_t^{a_t^i}|y_{1:t-1}) = \int p(y_t^{a_t^i}|x_t^i)p(x_t^i|y_{1:t-1})dx_t^i, \quad (6)$$

and acts as the normalising constant for the conditional filtering distribution.

The quantity $p(a_t^i|y_{1:t})$ in (3) is the association posterior for the i -th target, which can be obtained by marginalising the joint association posterior, *i.e.*

$$p(a_t^i|y_{1:t}) = \sum_{a_t^{1:N}} p(a_t^{1:N}|y_{1:t}), \quad (7)$$

where a_t^{-i} denotes $a_t^{1:N}$ with the i -th component removed. The joint association posterior can, in turn, be expressed as

$$\begin{aligned}
p(a_t^{1:N}|y_{1:t}) &= \int p(x_t^{1:N}, a_t^{1:N}|y_{1:t}) dx_t^{1:N} \\
&= \int \frac{p(y_t|x_t^{1:N}, a_t^{1:N})p(x_t^{1:N}, a_t^{1:N}|y_{1:t-1})}{p(y_t|y_{1:t-1})} dx_t^{1:N} \\
&\propto \int \underbrace{p(y_t|x_t^{1:N}, a_t^{1:N})}_{\text{joint likelihood}} \\
&\quad \times \underbrace{p(x_t^{1:N}|y_{1:t-1})}_{\text{predictions}} \underbrace{p(a_t^{1:N})}_{\text{association prior}} dx_t^{1:N}.
\end{aligned} \tag{8}$$

In the above $p(a_t^{1:N})$ is the prior of the joint associations. Both the joint likelihood and the prediction distributions factorise straightforwardly over the targets, so that the expression above simplifies to

$$\begin{aligned}
p(a_t^{1:N}|y_{1:t}) &\propto p(a_t^{1:N})c(y_t^{C(a_t^{1:N})}) \\
&\quad \times \prod_{i=1}^N \int p(y_t^{a_t^i}|x_t^i)p(x_t^i|y_{1:t-1}) dx_t^i \\
&= p(a_t^{1:N})c(y_t^{C(a_t^{1:N})}) \prod_{i=1}^N p(y_t^{a_t^i}|y_{1:t-1}),
\end{aligned} \tag{9}$$

where the terms in the product are the predictive likelihoods in (6), and the clutter model is applied to all the measurements in y_t that are not associated to any of the targets under the joint association hypothesis $a_t^{1:N}$. The indices of these measurements are denoted by $C(a_t^{1:N})$.

If the clutter model, $c(\cdot)$ is assumed to be such that the likelihood factorises over the clutter measurements, then (9) can be written as a product over the targets of independent terms:

$$\begin{aligned}
p(a_t^{1:N}|y_{1:t}) &\propto p(a_t^{1:N}) \prod_{y_t' \in (C(a_t^{1:N}))} c'(y_t') \prod_{i=1}^N \int p(y_t^{a_t^i}|x_t^i)p(x_t^i|y_{1:t-1}) dx_t^i \\
&= p(a_t^{1:N}) \prod_{y_t} c'(y_t) \prod_{i=1}^N \int \frac{p(y_t^{a_t^i}|x_t^i)}{c'(y_t^{a_t^i})} p(x_t^i|y_{1:t-1}) dx_t^i
\end{aligned} \tag{10}$$

where if a_t^i implies a missed detection, the corresponding value of $c'(y_t^{a_t^i})$ is unity. The fact that this approximation to the likelihood factorises over the targets is exploited

by strategies for either efficiently finding the maximum likelihood estimate of $a_t^{1:N}$ using Global Nearest Neighbour (GNN) algorithms such as the auction algorithm[1], or efficiently calculating this enumeration[6]. By making such a factorisation, one can exploit redundancy within the problem (brought about by gating procedures) to create such efficient algorithms.

The MTT tracking framework is now fully specified. At any time step the framework first computes the prediction distribution for each of the targets independently according to (2). The valid joint associations (after gating) are then enumerated, and the association probabilities computed for each target according to (7)-(9). Finally, for each target, the conditional filtering distributions are computed according to (5), and combined in proportion to the association probabilities in (3) to obtain the new target posterior.

3. A FRAMEWORK FOR SENSOR REGISTRATION

We will assume that all the sensor parameters are collected in the parameter θ_t , which is allowed to vary with time. We will also assume that the dynamic models for the targets are independent of θ_t . The distribution of interest is now the posterior $p(x_t^{1:N}, \theta_t|y_{1:t})$. Similar to (1) we will assume that this distribution factorises as

$$p(x_t^{1:N}, \theta_t|y_{1:t}) \approx p(\theta_t|y_{1:t}) \prod_{i=1}^N p(x_t^i|y_{1:t}), \tag{11}$$

i.e. the individual targets and the parameter are mutually independent. In what follows we will first show how to compute the target state posterior $p(x_t^i|y_{1:t})$, and then show how the parameter posterior $p(\theta_t|y_{1:t})$ can be obtained.

State Filtering

We show here how to compute the target state posterior $p(x_t^i|y_{1:t})$, so that it is independent of the parameter θ_t .

The prediction step for the target states proceeds as before, *i.e.*

$$p(x_t^i|y_{1:t-1}) = \int \underbrace{p(x_t^i|x_{t-1}^i)}_{\text{target dynamics}} \underbrace{p(x_{t-1}^i|y_{1:t-1})}_{\text{previous filtering}} dx_{t-1}^i \tag{12}$$

since the target dynamic models are independent of the parameter.

By introducing the association variable and the parame-

ter the filtering distribution can be expanded as

$$\begin{aligned}
p(x_t^i | y_{1:t}) &= \int \sum_{a_t^i} p(x_t^i, a_t^i, \theta_t | y_{1:t}) d\theta_t \\
&= \int \left[\sum_{a_t^i} p(a_t^i | \theta_t, y_{1:t}) p(x_t^i | a_t^i, \theta_t, y_{1:t}) \right] p(\theta_t | y_{1:t}) d\theta_t.
\end{aligned} \tag{13}$$

The term in square brackets is exactly the discrete mixture in (3), where the conditioning on θ_t was implicit. Thus, the new filtering distribution is now a continuous mixture, with one discrete mixture of the form in (3) for each valid value of the parameter.

The implication of the continuous mixture in (13) on the computational complexity is potentially large: the valid joint associations have to be enumerated for each value of the parameter to compute the association probabilities for the discrete mixture. This is clearly impractical, even for discrete approximations of the parameter posterior. To overcome this difficulty we consider the approximation where the joint posterior of the association and parameter variables factorises, *i.e.*

$$p(a_t^i, \theta_t | y_{1:t}) \approx p(a_t^i | y_{1:t}) p(\theta_t | y_{1:t}). \tag{14}$$

With this approximation the target posterior can be written as

$$p(x_t^i | y_{1:t}) \approx \sum_{a_t^i} p(a_t^i | y_{1:t}) p(x_t^i | a_t^i, y_{1:t}), \tag{15}$$

with

$$p(x_t^i | a_t^i, y_{1:t}) = \int p(x_t^i | a_t^i, \theta_t, y_{1:t}) p(\theta_t | y_{1:t}) d\theta_t. \tag{16}$$

Under reasonable assumptions the integral above is much simpler to solve or approximate than the one in (13). Thus, in (15) we again have a discrete mixture for the target posterior, with one mixture component for each valid association.

The marginal association posterior, which is required for the approximation above, can again be obtained from the corresponding joint, as in (7), with the joint now given by

$$\begin{aligned}
p(a_t^{1:N} | y_{1:t}) &= \int p(x_t^{1:N}, a_t^{1:N}, \theta_t | y_{1:t}) dx_t^{1:N} d\theta_t \\
&= \int \frac{p(y_t | x_t^{1:N}, a_t^{1:N}) p(x_t^{1:N}, a_t^{1:N}, \theta_t | y_{1:t-1})}{p(y_t | y_{1:t-1})} dx_t^{1:N} d\theta_t \\
&\propto \int \underbrace{p(y_t | x_t^{1:N}, a_t^{1:N})}_{\text{joint likelihood}} \underbrace{p(x_t^{1:N} | y_{1:t-1})}_{\text{predictions}} \underbrace{p(\theta_t | y_{1:t-1})}_{\text{association prior}} \\
&\quad \times p(a_t^{1:N}) dx_t^{1:N} d\theta_t.
\end{aligned} \tag{17}$$

Taking account of the factorisation of the joint likelihood and prediction distributions over the targets, leads to the simplification

$$\begin{aligned}
p(a_t^{1:N} | y_{1:t}) &\propto p(a_t^{1:N}) c(y_t^{C(a_t^{1:N})}) \int p(\theta_t | y_{1:t-1}) \prod_{i=1}^N p(y_t^{a_t^i} | \theta_t, y_{1:t-1}) d\theta_t,
\end{aligned} \tag{18}$$

which is similar to (9).

What remains is the computation of the parameter posterior $p(\theta_t | y_{1:t})$, which is required for the marginalisation over the parameter in (13) and (16). We discuss this in the next section.

Parameter Filtering

We now show how to compute the parameter posterior, or parameter filtering distribution, $p(\theta_t | y_{1:t})$. The prediction step is based upon a time varying Markov process, as follows:

$$p(\theta_t | y_{1:t-1}) = \int \underbrace{p(\theta_t | \theta_{t-1})}_{\text{parameter dynamics}} \underbrace{p(\theta_{t-1} | y_{1:t-1})}_{\text{previous filtering}} d\theta_{t-1}. \tag{19}$$

By introducing the associations the parameter filtering distribution can be expanded as

$$\begin{aligned}
p(\theta_t | y_{1:t}) &= \sum_{a_t^{1:N}} p(\theta_t, a_t^{1:N} | y_{1:t}) \\
&= \sum_{a_t^{1:N}} p(a_t^{1:N} | y_{1:t}) p(\theta_t | a_t^{1:N}, y_{1:t}),
\end{aligned} \tag{20}$$

where $p(a_t^{1:N} | y_{1:t})$ is the joint association posterior in (18). By using Bayes' rule and the independence of the prediction distributions for the targets, the conditional

parameter filtering distribution above can be expressed as

$$p(\theta_t | a_t^{1:N}, y_{1:t}) = \frac{p(\theta_t | y_{1:t-1}) \prod_{i=1}^N p(y_t^{a_i} | \theta_t, y_{1:t-1})}{\int p(\phi_t | y_{1:t-1}) \prod_{i=1}^N p(y_t^{a_i} | \phi_t, y_{1:t-1}) d\phi_t} \quad (21)$$

Substituting this result into (20) yields

$$p(\theta_t | y_{1:t}) = \sum_{a_t^{1:N}} p(a_t^{1:N} | y_{1:t}) \frac{p(\theta_t | y_{1:t-1}) \prod_{i=1}^N p(y_t^{a_i} | \theta_t, y_{1:t-1})}{\int p(\phi_t | y_{1:t-1}) \prod_{i=1}^N p(y_t^{a_i} | \phi_t, y_{1:t-1}) d\phi_t}. \quad (22)$$

Since all the terms in the above expression can be computed straightforwardly one option is to solve this expression directly. However, the summation over all the valid joint hypotheses may be computationally cumbersome, and it would be convenient to find an approximation that allows us to factorise the above expression over the individual targets. We can achieve such an approximation if we assume that the joint association posterior factorises over the targets, *i.e.*

$$p(a_t^{1:N} | y_{1:t}) \approx \prod_{i=1}^N p(a_t^i | y_{1:t}), \quad (23)$$

and that the integral in the denominator of (22) can be approximated as

$$\begin{aligned} & \int p(\phi_t | y_{1:t-1}) \prod_{i=1}^N p(y_t^{a_i} | \phi_t, y_{1:t-1}) d\phi_t \\ & \approx \prod_{i=1}^N \int [p(\phi_t | y_{1:t-1})]^{1/N} p(y_t^{a_i} | \phi_t, y_{1:t-1}) d\phi_t. \end{aligned} \quad (24)$$

This approximation will be reasonably accurate if it is true that the product of predictive likelihoods that do not share the same parameter value is negligible. Whether this is a reasonable assumption has to be verified experimentally. By applying the approximations in (23) and (24) to the expression in (22), the parameter filtering distribution can finally be expressed as

$$p(\theta_t | y_{1:t}) \approx p(\theta_t | y_{1:t-1}) \times \prod_{i=1}^N \left[\sum_{a_t^i} \frac{p(a_t^i | y_{1:t}) p(y_t^{a_i} | \theta_t, y_{1:t-1})}{\int [p(\phi_t | y_{1:t-1})]^{1/N} p(y_t^{a_i} | \phi_t, y_{1:t-1}) d\phi_t} \right]. \quad (25)$$

We may also try an approximation where we approximate as independent of a_t^i the integral in the denomina-

tor of (25), in which case the parameter filtering distribution further simplifies to

$$p(\theta_t | y_{1:t}) \approx C p(\theta_t | y_{1:t-1}) \prod_{i=1}^N \left[\sum_{a_t^i} p(a_t^i | y_{1:t}) p(y_t^{a_i} | \theta_t, y_{1:t-1}) \right], \quad (26)$$

where C is the constant that normalises the expression to be a proper distribution. Clearly, by considering the targets one-at-a-time one can perform mixture reduction sequentially, thus reducing the computational burden naively associated with (26).

4. MIXTURE (UNSCENTED) KALMAN FILTER

In this section we sketch how the sensor registration framework presented in the previous section can be implemented using Gaussian mixtures in the spirit of the mixture (unscented) Kalman filter. The Unscented Transform (UT) has been used to conduct an approximation of the state transition and observation densities, though quasi-Monte Carlo sampling. The resulting filter, the *Unscented Kalman Filter* (UKF), considers a set of (*sigma*) points that are deterministically selected from the Gaussian approximation to $p(x_{t-1} | y_{1:t-1})$. These points are all propagated through the true models and the parameters of $p(x_t, y_t | y_{1:t-1})$ are then estimated from the transformed samples. The reader is referred to reference [4] for a recent review of standard unscented filtering techniques. Using this framework based upon other techniques (such as Extended Kalman filtering or particle filtering) is deemed to be a simple operation.

Preliminaries

In the derivations that follow we will need some results concerning Gaussians and Gaussian mixtures. We refer the reader to any standard basic probability text for proofs.

- Suppose that the random variables x and y are jointly Gaussian distributed, *i.e.*

$$p(x, y) = \mathcal{N}\left(\begin{bmatrix} x \\ y \end{bmatrix} \mid \begin{bmatrix} \mu_x \\ \mu_y \end{bmatrix}, \begin{bmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{xy}^T & \Sigma_{yy} \end{bmatrix}\right). \quad (27)$$

Then the marginal and conditional distributions of x are also Gaussian, and given by

$$p(x) = \mathcal{N}(x | \mu_x, \Sigma_{xx}) \quad (28)$$

$$p(x|y) = \mathcal{N}(x | \mu_{x|y}, \Sigma_{xx|y}), \quad (29)$$

with

$$\mu_{x|y} = \mu_x + \Sigma_{xy} \Sigma_{yy}^{-1} (y - \mu_y) \quad (30)$$

$$\Sigma_{xx|y} = \Sigma_{xx} - \Sigma_{xy} \Sigma_{yy}^{-1} \Sigma_{xy}^T. \quad (31)$$

A similar result holds for the marginal and conditional distributions of y .

- Suppose that the random variables x and y are jointly distributed according to a Gaussian mixture, *i.e.*

$$p(x, y) = \sum_k w_k \mathcal{N} \left(\begin{bmatrix} x \\ y \end{bmatrix} \middle| \begin{bmatrix} \mu_{x,k} \\ \mu_{y,k} \end{bmatrix}, \begin{bmatrix} \Sigma_{xx,k} & \Sigma_{xy,k} \\ \Sigma_{xy,k}^T & \Sigma_{yy,k} \end{bmatrix} \right). \quad (32)$$

Then the marginal and conditional distributions of x are also Gaussian mixtures, and given by

$$p(x) = \sum_k w_k \mathcal{N}(x | \mu_{x,k}, \Sigma_{xx,k}) \quad (33)$$

$$p(x|y) = \sum_k \tilde{w}_k \mathcal{N}(x | \mu_{x|y,k}, \Sigma_{xx|y,k}), \quad (34)$$

with

$$\mu_{x|y,k} = \mu_{x,k} + \Sigma_{xy,k} \Sigma_{yy,k}^{-1} (y - \mu_{y,k}) \quad (35)$$

$$\Sigma_{xx|y,k} = \Sigma_{xx,k} - \Sigma_{xy,k} \Sigma_{yy,k}^{-1} \Sigma_{xy,k}^T \quad (36)$$

$$\tilde{w}_k = \frac{w_k \mathcal{N}(y | \mu_{y,k}, \Sigma_{yy,k})}{\sum_l w_l \mathcal{N}(y | \mu_{y,l}, \Sigma_{yy,l})}. \quad (37)$$

A similar result holds for the marginal and conditional distributions of y .

Representation

We will represent both the target and parameter filtering distributions with Gaussian mixtures. We would like the operations at a particular time step to be closed with respect to this representation, *i.e.* starting from

$$p(x_{t-1}^i | y_{1:t-1}) = \sum_j \alpha_{j,t-1}^i \mathcal{N}(x_{t-1}^i | a_{j,t-1}^i, P_{j,t-1}^i) \quad (38)$$

$$p(\theta_{t-1} | y_{1:t-1}) = \sum_k \gamma_{k,t-1} \mathcal{N}(\theta_{t-1} | b_{k,t-1}, Q_{k,t-1}), \quad (39)$$

we would like the result of the operations at time t to be

$$p(x_t^i | y_{1:t}) = \sum_j \alpha_{j,t}^i \mathcal{N}(x_t^i | a_{j,t}^i, P_{j,t}^i) \quad (40)$$

$$p(\theta_t | y_{1:t}) = \sum_k \gamma_{k,t} \mathcal{N}(\theta_t | b_{k,t}, Q_{k,t}). \quad (41)$$

In the remainder of this section we will outline a strategy to achieve this.

Prediction

The joint prediction distribution for the i -th target state, target-originated measurement and the parameter can

be expanded as

$$\begin{aligned} p(y_t^i, x_t^i, \theta_t | y_{1:t-1}) &= p(y_t^i | x_t^i) \int p(x_t^i | x_{t-1}^i) p(x_{t-1}^i | y_{1:t-1}) dx_{t-1}^i \\ &\quad \times \int p(\theta_t | \theta_{t-1}) p(\theta_{t-1} | y_{1:t-1}) d\theta_{t-1} \\ &= \sum_j \sum_k \alpha_{j,t-1}^i \gamma_{k,t-1} p(y_t^i | x_t^i) \\ &\quad \times \int p(x_t^i | x_{t-1}^i) \mathcal{N}(x_{t-1}^i | a_{j,t-1}^i, P_{j,t-1}^i) dx_{t-1}^i \\ &\quad \times \int p(\theta_t | \theta_{t-1}) \mathcal{N}(\theta_{t-1} | b_{k,t-1}, Q_{k,t-1}) d\theta_{t-1}. \end{aligned} \quad (42)$$

The integrals in the above expression can be solved by the Kalman filter for linear Gaussian models. For general non-linear and non-Gaussian models, approximations can be obtained with the unscented Kalman filter. In either case the joint prediction distribution is a Gaussian mixture that can be written as

$$p(y_t^i, x_t^i, \theta_t | y_{1:t-1}) = \sum_j \sum_k \alpha_{j,t-1}^i \gamma_{k,t-1} \mathcal{N} \left(\begin{bmatrix} y_t^i \\ x_t^i \\ \theta_t \end{bmatrix} \middle| \begin{bmatrix} \hat{y}_{jk,t}^i \\ \hat{a}_{j,t}^i \\ \hat{b}_{k,t} \end{bmatrix}, \Upsilon_{jk,t} \right), \quad (43)$$

where

$$\Upsilon_{jk,t} = \begin{bmatrix} \hat{S}_{jk,t}^i & \hat{U}_{jk,t}^{iT} & \hat{V}_{jk,t}^{iT} \\ \hat{U}_{jk,t}^i & \hat{P}_{j,t}^i & 0 \\ \hat{V}_{jk,t}^i & 0 & \hat{Q}_{k,t} \end{bmatrix}.$$

Using the results in Section 4 many of the distributions of interest can be straightforwardly derived from this joint. We give some of these in the below.

- **State and parameter prediction.** The marginal prediction distributions for the target state and the parameter follow straightforwardly by marginalising (43) as

$$p(x_t^i | y_{1:t-1}) = \sum_j \alpha_{j,t-1}^i \mathcal{N}(x_t^i | \hat{a}_{j,t}^i, \hat{P}_{j,t}^i) \quad (44)$$

$$p(\theta_t | y_{1:t-1}) = \sum_k \gamma_{k,t-1} \mathcal{N}(\theta_t | \hat{b}_{k,t}, \hat{Q}_{k,t}). \quad (45)$$

Note that the prediction distribution for the parameter follows independently of which target it is combined with.

- **Gating.** For gating we require the distribution $p(y_t^i | y_{1:t-1})$. This can be obtained by marginalising (43)

over (x_t^i, θ_t) , leading to

$$p(y_t^i | y_{1:t-1}) = \sum_j \sum_k \alpha_{j,t-1}^i \gamma_{k,t-1} \mathcal{N}(y_t^i | \hat{y}_{jk,t}^i, \hat{S}_{jk,t}^i). \quad (46)$$

• **Filtering.** For state filtering based on a target-originated measurement y_t^i the required distribution is the state posterior $p(x_t^i | y_{1:t-1}, y_t^i)$. It can be obtained by first marginalising the joint in (43) over the parameter θ_t , and then computing the conditional on the measurement, leading to

$$p(x_t^i | y_{1:t-1}, y_t^i) = \sum_j \sum_k \tilde{\alpha}_{jk,t}^i \mathcal{N}(x_t^i | a_{jk,t}^i, P_{jk,t}^i), \quad (47)$$

with

$$a_{jk,t}^i = \hat{a}_{j,t}^i + \hat{U}_{jk,t}^i [\hat{S}_{jk,t}^i]^{-1} (y_t^i - \hat{y}_{jk,t}^i) \quad (48)$$

$$P_{jk,t}^i = \hat{P}_{j,t}^i - \hat{U}_{jk,t}^i [\hat{S}_{jk,t}^i]^{-1} \hat{U}_{jk,t}^{iT} \quad (49)$$

$$\tilde{\alpha}_{jk,t}^i = \frac{\alpha_{j,t-1}^i \gamma_{k,t-1} \mathcal{N}(y_t^i | \hat{y}_{jk,t}^i, \hat{S}_{jk,t}^i)}{\sum_l \sum_m \alpha_{l,t-1}^i \gamma_{m,t-1} \mathcal{N}(y_t^i | \hat{y}_{lm,t}^i, \hat{S}_{lm,t}^i)}. \quad (50)$$

• **Measurement prediction.** Several of the computations for sensor registration require the measurement prediction distribution conditional on the parameter, *i.e.* $p(y_t^i | \theta_t, y_{1:t-1})$. It can be obtained by first marginalising the joint in (43) over the target state x_t^i , and then computing the conditional on the parameter, leading to

$$p(y_t^i | \theta_t, y_{1:t-1}) = \sum_j \sum_k \tilde{\omega}_{jk,t}^i \mathcal{N}(y_t^i | y_{jk,t}^i, S_{jk,t}^i), \quad (51)$$

with

$$y_{jk,t}^i = \hat{y}_{jk,t}^i + \hat{V}_{jk,t}^{iT} \hat{Q}_{k,t}^{-1} (\theta_t - \hat{b}_{k,t}) \quad (52)$$

$$S_{jk,t}^i = \hat{S}_{jk,t}^i - \hat{V}_{jk,t}^{iT} \hat{Q}_{k,t}^{-1} \hat{V}_{jk,t}^i \quad (53)$$

$$\tilde{\omega}_{jk,t}^i = \frac{\alpha_{j,t-1}^i \gamma_{k,t-1} \mathcal{N}(\theta_t | \hat{b}_{k,t}, \hat{Q}_{k,t})}{\sum_l \sum_m \alpha_{l,t-1}^i \gamma_{m,t-1} \mathcal{N}(\theta_t | \hat{b}_{m,t}, \hat{Q}_{m,t})}. \quad (54)$$

In a similar vein to the above it is possible to obtain the full joint prediction distribution for all the targets and the parameter. This distribution takes the form

$$p(y_t^{1:N}, x_t^{1:N}, \theta_t | y_{1:t-1}) = \sum_J \sum_k \alpha_{J,t-1}^{1:N} \gamma_{k,t-1} \times \mathcal{N} \left(\begin{bmatrix} y_t^{1:N} \\ x_t^{1:N} \\ \theta_t \end{bmatrix} \middle| \begin{bmatrix} \hat{y}_{Jk,t}^{1:N} \\ \hat{a}_{J,t}^{1:N} \\ \hat{b}_{k,t} \end{bmatrix}, \begin{bmatrix} \hat{S}_{Jk,t}^{1:N} & \hat{U}_{Jk,t}^{1:NT} & \hat{V}_{Jk,t}^{1:NT} \\ \hat{U}_{Jk,t}^{1:N} & \hat{P}_{J,t}^{1:N} & 0 \\ \hat{V}_{Jk,t}^{1:N} & 0 & \hat{Q}_{k,t} \end{bmatrix} \right), \quad (55)$$

where $J = (j^1 \dots j^N)$ is the multivariate index ranging over the target mixture components, and the parameters

are given by

$$\alpha_{J,t-1}^{1:N} = \prod_{i=1}^N \alpha_{j^i,t-1}^i \quad (56)$$

$$\hat{y}_{Jk,t}^{1:N} = (\hat{y}_{j^1 k,t}^1 \dots \hat{y}_{j^N k,t}^N) \quad (57)$$

$$\hat{a}_{J,t}^{1:N} = (\hat{a}_{j^1,t}^1 \dots \hat{a}_{j^N,t}^N) \quad (58)$$

$$\hat{S}_{Jk,t}^{1:N} = \text{diag}(\hat{S}_{j^1 k,t}^1 \dots \hat{S}_{j^N k,t}^N) \quad (59)$$

$$\hat{P}_{J,t}^{1:N} = \text{diag}(\hat{P}_{j^1,t}^1 \dots \hat{P}_{j^N,t}^N) \quad (60)$$

$$\hat{U}_{Jk,t}^{1:N} = \text{diag}(\hat{U}_{j^1 k,t}^1 \dots \hat{U}_{j^N k,t}^N) \quad (61)$$

$$\hat{V}_{Jk,t}^{1:N} = (\hat{V}_{j^1 k,t}^1 \dots \hat{V}_{j^N k,t}^N). \quad (62)$$

As before, we can obtain any marginal or conditional of interest from this joint distribution. Of particular interest is the marginal posterior of the parameter $p(\theta_t | y_{1:t-1}, y_t^{1:N})$. This distribution can be obtained from the full joint in (55) by first marginalising over the target states $x_t^{1:N}$, and then conditioning on the measurements $y_t^{1:N}$, leading to

$$p(\theta_t | y_{1:t-1}, y_t^{1:N}) = \sum_J \sum_k \tilde{\gamma}_{Jk,t} \mathcal{N}(\theta_t | b_{Jk,t}, Q_{Jk,t}), \quad (63)$$

with

$$b_{Jk,t} = \hat{b}_{k,t} + \hat{V}_{Jk,t}^{1:N} [\hat{S}_{Jk,t}^{1:N}]^{-1} (y_t^{1:N} - \hat{y}_{Jk,t}^{1:N}) \quad (64)$$

$$Q_{Jk,t} = \hat{Q}_{k,t} - \hat{V}_{Jk,t}^{1:N} [\hat{S}_{Jk,t}^{1:N}]^{-1} \hat{V}_{Jk,t}^{1:NT} \quad (65)$$

$$\tilde{\gamma}_{Jk,t} = \frac{\alpha_{J,t-1}^{1:N} \gamma_{k,t-1} \mathcal{N}(y_t^{1:N} | \hat{y}_{Jk,t}^{1:N}, \hat{S}_{Jk,t}^{1:N})}{\sum_L \sum_m \alpha_{L,t-1}^{1:N} \gamma_{m,t-1} \mathcal{N}(y_t^{1:N} | \hat{y}_{Lm,t}^{1:N}, \hat{S}_{Lm,t}^{1:N})}. \quad (66)$$

Hypothesis Evaluation

Subsequent to measurement gating all the valid joint hypothesis can be enumerated. These then need to be substituted into (18) for evaluation, and normalised, to obtain the posterior joint association probabilities. Substituting the expressions for the parameter and measurement predictions in (45) and (51), respectively, into (18), leads to

$$p(a_t^{1:N} | y_{1:t}) \propto p(a_t^{1:N}) c(y_t^{C(a_t^{1:N})}) \times \sum_k \gamma_{k,t-1} \int \mathcal{N}(\theta_t | \hat{b}_{k,t}, \hat{Q}_{k,t}) \times \prod_{i=1}^N \left[\sum_j \sum_l \tilde{\omega}_{jl,t}^i(\theta_t) \mathcal{N}(y_t^{a_t^i} | y_{jl,t}^i(\theta_t), S_{jl,t}^i) \right] d\theta_t, \quad (67)$$

where we have made the dependency of the measurement prediction parameters on the parameter θ_t explicit. The integral in the above expression cannot be solved analytically, but if we draw sigma-points according to $\hat{\sigma}_{\theta_t, k} = \{\hat{w}_{k,t}^n, \hat{\theta}_{k,t}^n\} \sim \mathcal{N}(\theta_t | \hat{b}_{k,t}, \hat{Q}_{k,t})$, we can approximate it using the unscented approximation, leading to

$$p(a_t^{1:N} | y_{1:t}) \approx K p(a_t^{1:N}) c(y_t^{C(a_t^{1:N})}) \sum_k \sum_n \gamma_{k,t-1} \hat{w}_{k,t}^n \times \prod_{i=1}^N \left[\sum_j \sum_l \tilde{w}_{jl,t}^i(\hat{\theta}_{k,t}^n) \mathcal{N}(y_t^{a_t^i} | y_{jl,t}^i(\hat{\theta}_{k,t}^n), S_{jl,t}^i) \right], \quad (68)$$

where K is the constant that normalises the probabilities. It can be obtained as the reciprocal of the sum over all the unnormalised posterior joint association probabilities. Once the joint probabilities are computed, the marginal probabilities for the individual targets can be easily obtained by suitable marginalisation (summation).

State Filtering

We can now obtain the marginal state filtering distribution by substituting mixture state filtering distribution for a target-originated measurement in (47), into the expression in (15), leading to

$$p(x_t^i | y_{1:t}) = \sum_{a_t^i} \sum_j \sum_k p(a_t^i | y_{1:t}) \tilde{\alpha}_{jk,t}^i(y_t^{a_t^i}) \times \mathcal{N}(x_t^i | a_{jk,t}^i(y_t^{a_t^i}), P_{jk,t}^i),$$

where we have made the dependency of the parameters of the mixture distribution on the measurement explicit. This follows from the fact that $p(x_t^i | a_t^i, y_{1:t}) = p(x_t^i | y_{1:t-1}, y_t^{a_t^i})$. Thus, as required, the new state filtering distribution is again a Gaussian mixture. As is evident from the expression above the number of mixture components will tend to grow exponentially over time, so that mixture reduction techniques are required to make the approach computationally feasible.

Parameter Filtering

In a similar spirit to state filtering we can obtain the parameter filtering distribution by substituting the mixture parameter filtering distribution for a set of target-originated measurements in (63), into the expression in

(20), leading to

$$p(\theta_t | y_{1:t}) = \sum_{a_t^{1:N}} \sum_J \sum_k p(a_t^{1:N} | y_{1:t}) \tilde{\gamma}_{Jk,t}(y_t^{a_t^{1:N}}) \times \mathcal{N}(\theta_t | b_{Jk,t}(y_t^{a_t^{1:N}}), Q_{Jk,t}),$$

where we have made the dependency of the parameters of the mixture distribution on the measurements explicit. This follows from the fact that $p(\theta_t | a_t^{1:N}, y_{1:t}) = p(\theta_t | y_{1:t-1}, y_t^{a_t^{1:N}})$, under the assumption that clutter measurements contain no useful information for parameter estimation. Thus, as was the case for state filtering, the new parameter filtering distribution is again a Gaussian mixture. Here also the number of mixture components will grow exponentially over time, making mixture reduction techniques necessary.

5. CONCLUSIONS

In this paper we have presented a generic probabilistic framework that solves the sensor registration problem using targets of opportunity. The derivation of the framework ensures that the set of *existing* MTT algorithms can be utilised without *any* modification. In addition to this generic probabilistic exposition, we provide a brief overview of the mixture unscented Kalman filter, a novel filtering algorithm that could be employed in such scenarios (or more generally in other contexts). Simulation results have been deferred to a future journal publication; it is the intention of this paper to disclose the framework to the scientific community, with the authors noting that it is our belief that the integrity of the data association algorithm will determine the quality of the results¹, and not the framework presented.

6. ACKNOWLEDGEMENTS

The authors would like to thank the Royal Commission for the Exhibition of 1851 (www.royalcommission1851.org.uk). This research was supported by the UK MOD Corporate Research Programme through work funded by the Defence Technology Centre for Data and Information Fusion. We thank Dr Simon Godsill of Cambridge University Engineering Department for his contribution to this research.


¹That is not to say that the integrity of the data association algorithm will guarantee that the results will follow; we postulate that using a sophisticated data association algorithm (such as multi-frame assignment, for example) will give better results than a naive (PDAF based, for example) approach. Finding the correct compromise between computational complexity and algorithmic performance is the top priority. Results that compare the framework based upon different data association techniques are not necessarily informative about the framework in general.

REFERENCES

- [1] Y Bar-Shalom and W D Blair, editors. *Multitarget-Multisensor Tracking: Applications and Advances Volume III*. Artech House, Boston/London, 2000.
- [2] Y Bar-Shalom and X R Li. *Estimation and Tracking: Principles, Techniques, and Software*. Artech House Inc., 1993.
- [3] G Chesi, A Garulli, A Vicino, and R Cipolla. Estimating the fundamental matrix via constrained least-squares: a convex approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(3):397–401, 2002.
- [4] S J Julier and J K Uhlmann. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3):401–422, 2004.
- [5] M Mallick. Maximum likelihood geolocation using a ground moving target indicator (GMTI) report. In *Aerospace Conference Proceedings*, volume 4, pages 4–1561–4–1570, 2002.
- [6] S R Maskell, M Briers, and R Wright. Efficient hypothesis management. In *SPIE*, 2004.
- [7] N Nabaa and R H Bishop. A recursive solution to the sensor registration problem in a multiple sensor tracking scenario. In *SPIE, Signal and Data Processing of Small Targets*, volume 2759, pages 385–396, 1996.
- [8] N N Okello and S Challa. Joint sensor registration and track-to-track fusion for distributed trackers. *IEEE Transactions on Aerospace and Electronic Systems*, 40:808–823, July 2004.



Jaco Vermaak was born in South Africa in 1969. He received the B.Eng. and M.Eng. degrees from the University of Pretoria, South Africa in 1993 and 1996, respectively, and the Ph.D. degree at the University of Cambridge, U.K. in 2000. He has worked as a post-doctoral researcher at Microsoft Research Europe, Cambridge, U.K. between 2000 and 2002. He is currently employed as a senior research associate in the Signal Processing Group of the Cambridge University Engineering Department. His research interests include audio-visual tracking techniques, multi-media manipulation, statistical signal processing methods and machine learning. ∩∩



Simon Maskell holds PhD a first class degree with Distinction from Cambridge University Engineering Department. His PhD was funded by one of six prestigious Royal Commission for the Exhibition of 1851 Industrial fellowships awarded on the basis of outstanding excellence to researchers working in British industry. At QinetiQ, he is a lead researcher for tracking in the Advanced Signal and Information Processing group, ASIP. As such, he leads on several projects and coordinates ASIP's tracking research while also supervising a number of other researchers. Simon has authored a number of papers, as well as several technical reports and a patent. He has also been the leading force behind the development of a QinetiQ product to provide a generic solution to all of QinetiQ's tracking problems. ∩∩



Mark Briers joined QinetiQ after receiving a First Class (Hons) degree in Mathematics in July 2001. As part of his undergraduate studies, Mark was awarded the Dr. David Whiteman prize for “outstanding work” and was also presented with the Institute of Mathematics and its Applications (IMA) prize. In September 2003 he was awarded one of six prestigious Royal Commission for the Exhibition of 1851 Industrial fellowships, which funds his work at QinetiQ towards his PhD from Cambridge University Engineering Department. He has authored several international conference papers and several technical reports. ∩∩