

The Identity Management Kalman Filter (IMKF)

Brad Schumitsch, Sebastian Thrun, Leonidas Guibas, Kunle Olukotun
Computer Science Department
Stanford University
Email: schumitsch@,thrun@,guibas@cs.,kunle@stanford.edu

Abstract—Tracking posterior estimates for problems with data association uncertainty is one of the big open problems in the literature on filtering and tracking. This paper presents a new filter for the on-line tracking of many individual objects with data association ambiguities. It tightly integrates the continuous aspects of the problem (locating the objects) with the discrete aspects (the data association ambiguities). The key innovation is a probabilistic information matrix that efficiently encodes identity information by linking world entities with internal tracks of the filter, enabling it to maintain a full posterior over the system amid data association uncertainties. The filter scales quadratically in complexity, just like a conventional Kalman filter, avoiding the exponential explosion of the association space. We derive the algorithm formally and present large-scale simulation results.

I. INTRODUCTION

Tracking many objects is of great importance in many real-world applications, including flight control systems, large scale sensor networks [6], camera networks, and so on [2]. One of the key challenges in this domain pertains to the fact that multi-object tracking combines both *continuous* and *discrete* inference: continuous, because the objects that are being tracked move through continuous spaces, and discrete because of the inherent data association problems of assigning measurements and object identities to individual tracks.

The literature has extensively studied both aspects, although, as we argue, with unsatisfactory results for the discrete aspects of the problem. The continuous aspect of the problem is usually handled using Kalman filters [21] or the various nonlinear variants thereof [7], [11], [8], [22]. However, the discrete data association is still not solved in ways that scale satisfactorily to large systems.

It is easily seen that the full posterior under unknown data association grows exponentially. This is because at each fixed time step there might exist $O(N!)$ different ways of associating the measurement vector entries with the objects in the environment (N denotes the number of objects). A common approximation is the popular multi-hypothesis KF [15], [1], which selects a fixed subset of the most likely K among those hypotheses. Unfortunately, MHKFs are subject to the curse of dimensionality. It is easy to show that the appropriate number K may grow exponentially in N , in order to meet any reasonable error bound. This is because each of those hypotheses corresponds to a specific presumed data association assignment for all the objects in the filter. For the same reason, the popular paradigm of particle filters [14], [5] suffers from this exponential blow-up, because each particle is based on a fixed, non-probabilistic commitment for the data association variables.

This paper proposes a new Kalman filter for tracking N objects, which retains a full posterior over data associations yet avoids this exponential complexity. The representation of this tracker is factored into a continuous Kalman filter, and a discrete data association matrix. Thus, the *identity management Kalman filter (IMKF)* maintains three parameters: a conventional mean μ and covariance Σ , but also a novel quadratic data association matrix Ω .

The IMKF bypasses the problem arising from unknown data associations by maintaining internal “tracks.” Each such track might correspond to multiple objects in the environment, at different points in time. The probabilistic correspondence between the internal tracks and the objects in the environment is encoded in the identity association matrix Ω . This matrix maintains a posterior in *information form*, which is a non-normalized logarithmic form of the data association probabilities. The advantage of this form is that (under mild approximations) it encodes the otherwise exponential distribution over the exponentially many ways to assign internal tracks to objects in the physical world, into a much simple quadratic matrix.

In deriving the IMKF, we make a number of approximations to escape the otherwise inherent exponential complexity of the true posterior. The key approximation lies in the fact that we use a maximum likelihood data association track for updating the continuous KF parameters; however, the discrete data association probabilities are then updated in the full Bayesian posterior form. Another approximation occurs when updating the data association matrix. Here we approximate a set of marginals using Jensen’s inequality, which is essential for avoiding an exponential expansion of the posterior. While these approximations are significant, they appear not to harm the filter’s ability to generate useful full posteriors over the objects being tracked under data association uncertainty. As experimental results illustrate, our filter effortlessly scales to large values of N (e.g., $N = 3,000$, which makes for a Kalman filter with 12,000 dimensions), a size where previous filters fail.

The identity management Kalman filter (IMKF) is related to a rich body of prior work. The idea of using association matrices in information form goes back to [20], [17]; however, neither of these papers address the continuous tracking problem, and as a result the discrete methods are somewhat *ad hoc*. It is also related to the Joint Probabilistic Data Association filter (JPDA), which encapsulates data association ambiguity in its continuous filter updates (which IMKF does *not* do). Unfortunately, if there is a large amount of ambiguity in




Figure 1. Motion Path of 4 objects with 3 Data Association Ambiguities

the data association decisions, JPDA estimates will have an inflated uncertainty in the continuous state estimates of the objects; the same is the case for the particle filter version of JPDA [16]. We find that the resulting posteriors quickly become unusable in tracking applications with large identity uncertainties. Very recent work uses MCMC algorithms to give a provable approximation algorithm for JPDA [12], but such techniques are still quite expensive computationally. The IMKF is the first method that scales to large N and can cope with massive data association uncertainty, yet maintains a focused posterior estimate.

Although we do not address the distributed computation aspects of the IMKF in this paper, an especially attractive feature of the filter in such settings, such as wireless sensor networks, is that it permits the local accumulation of identity evidence along each track by distributing the columns of the identity matrix into the network, obviating the need for expensive normalizations in the style of [19]. See [20], [17] for further discussion.

Finally, it should be noted that there exist alternative, non-filter methods for computing tracking posteriors with multiple objects [13], [4]. These methods require that all data be stored and are hence not executable online.

The remainder of this paper is organized as follows. We begin with an intuitive description of the IMKF in Section II. Section III reviews a standard and a maximum likelihood Kalman filter. Section IV introduces the theoretical-IMKF (TIMKF), an inefficient filter that illustrates the main idea of the IMKF. Next in section V we describe the IMKF itself. Experimental results are provided in Section VI, followed by the paper's conclusion.

II. INTUITIVE DESCRIPTION

To describe the benefits of the identity management Kalman filter (IMKF) intuitively consider the situation in Figure 1.

We are tracking $N = 4$ objects. At time $t = 0$ there is no ambiguity in our system, and the position of objects A, B, C, and D are known. At time $t = 1$ we encounter two data association ambiguities – between objects A and B, as well as C and D. That is, the measurements received at $t = 2$

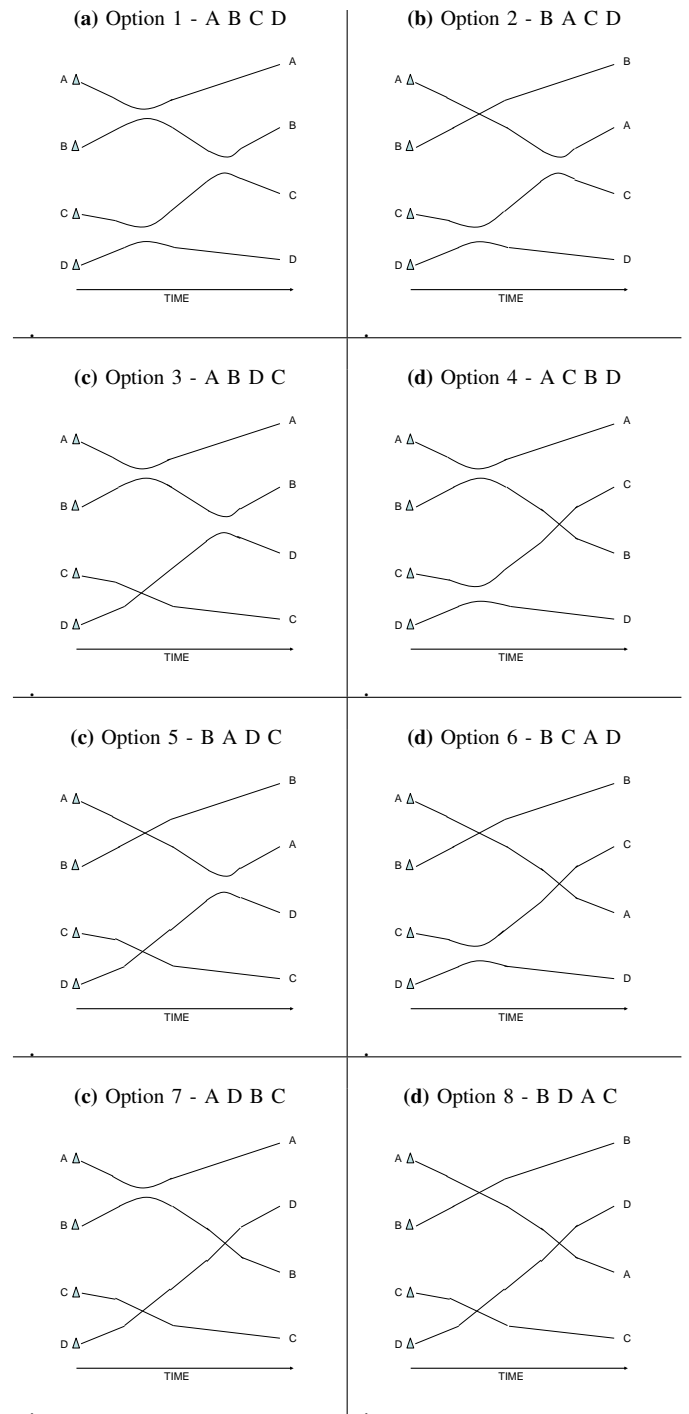


Figure 2. All 8 motion paths are consistent with measurements through $t = 2$. An MHT maintains all 8 options and selects option 6 as correct after the measurements at $t = 3$

are equally consistent with objects A and B crossing paths as well as with not crossing paths (for example, the measurements might not have contained any identity information, and the position information was consistent with both models of motion). The paths of objects of C and D are similarly ambiguous.

At time $t = 2$ a 3-rd data association ambiguity occurs between the middle two objects (notice these object identities

are unknown). At time $t = 3$, measurements containing identity information indicate the locations of objects C and A. These measurements contain enough information to resolve all ambiguity.

A multiple hypothesis tracker (MHT) correctly resolves the ambiguity by storing $2^3 = 8$ distinctive hypotheses. After $t = 1$ the following possibilities are equally likely. [ABCD, BACD, ABDC, BADC]. After $t = 2$ the following 8 possibilities are possible. [ABCD, BACD, ABDC, BADC, ACBD, BCAD, ADBC, BDAC] (see figure 2). After the measurements at $t = 3$ are taken into account, only the hypothesis BCAD is consistent, and the MHT concludes this is the correct state.

The IMKF provides the power of the MHT, while maintaining an $O(N^2)$ complexity, as opposed to $O(N!)$. Just like the Kalman filter, the identity management Kalman filter (IMKF) maintains N tracks. These tracks are parameterized by a $\{\Sigma, \mu\}$ pair. And just like the KF, the IMKF integrates continuous pose/velocity estimated into the most likely track. However, the key difference is that the IMKF also maintains an identity management matrix. This matrix is illustrated in Figure 1c. Each entry in this matrix specifies the “amount of evidence” we have that a specific track corresponds to a specific object (in a yet-to-be-specified form). In the beginning, we might have strong information that track 1 corresponds to object A, track 2 to object B. 3 to C, and 4 to D. This is expressed by the matrix

$$\Omega(0) = \begin{pmatrix} 10 & 0 & 0 & 0 \\ 0 & 10 & 0 & 0 \\ 0 & 0 & 10 & 0 \\ 0 & 0 & 0 & 10 \end{pmatrix}. \quad (1)$$

At time t_1 data association ambiguities obscure the paths of tracks 1 and 2, as well as 3 and 4. The IMKF responds by “averaging” the corresponding tracks together (the exact formula for this averaging is given below).

$$\Omega(1) = \begin{pmatrix} 9.3 & 9.3 & 0 & 0 \\ 9.3 & 9.3 & 0 & 0 \\ 0 & 0 & 9.3 & 9.3 \\ 0 & 0 & 9.3 & 9.3 \end{pmatrix} \quad (2)$$

At time t_2 data association ambiguities cause more ambiguity. The IMKF responds by “averaging” tracks 2 and 3

$$\Omega(2) = \begin{pmatrix} 9.3 & 8.6 & 8.6 & 0 \\ 9.3 & 8.6 & 8.6 & 0 \\ 0 & 8.6 & 8.6 & 9.3 \\ 0 & 8.6 & 8.6 & 9.3 \end{pmatrix}. \quad (3)$$

At time t_3 measurements strongly indicate the location of objects C and A. The IMKF efficiently incorporates this information by simply adding a number (in this case 5) to the corresponding entry in Ω . The exact calculation of this number is discussed below. Thus

$$\Omega(3) = \begin{pmatrix} 9.3 & 8.6 & \mathbf{13.6} & 0 \\ \mathbf{9.3} & 8.6 & 8.6 & 0 \\ 0 & \mathbf{13.6} & 8.6 & 9.3 \\ 0 & 8.6 & 8.6 & \mathbf{9.3} \end{pmatrix}. \quad (4)$$

The most likely track-to-identity assignment is retrieved from Ω by selecting the N entries, one from each column and

row, that have the largest combined sum. In (4) these entries are in bold. Notice this indicates the assignment [B C A D], the correct assignment. The IMKF has correctly resolved the 3 data association ambiguities.

Notice that data association ambiguity strategies that are entirely local to the track level will not completely resolve the ambiguities from our example. For example, consider the strategy of each track keeping a distribution of its identity. After the data association ambiguity at time t_1 , track 1 could be A or B. Since track 1 is not measured at t_3 , the measurement at t_3 does not directly resolve this ambiguity.

III. STANDARD KALMAN FILTERS

A. Basic Problem

The problem is to track N objects in a closed world.

We let $x_i(t)$ denote the position of the i -th target at time t . We assume the target objects obey the following linear motion model:

$$x_i(t+1) = F_i(t)x_i(t) + G_i(t)u_i(t) + v_i(t), \quad (5)$$

where $F_i(t)$ and $G_i(t)$ are matrices and $v_i(t)$ is a random vector with covariance matrix $V_i(t)$.

At each time step, we receive N measurements of the form:

$$z_i(t) = H_i(t)x_i(t) + w_i(t), \quad (6)$$

where $H_i(t)$ is a matrix and $w_i(t)$ is a random vector with covariance matrix $W_i(t)$.

We assume a 1-dimensional state-space per object, no state-space cross-correlation between the objects being tracked, $u_i(t) = 0 \forall i, t$, and all of F , G , V , H , and W are constant over time and across objects. Extensions to more general settings are straightforward.

To simplify notation, we will sometimes drop the time argument when it is clear from context.

B. Kalman Filter with Known Data Association

The well-known Kalman Filter maintains a $[\mu_i, \Sigma_i]$ pair for every object being tracked. These two parameters implicitly define the following Gaussian pdf

$$p(x_i) = \det(2\pi\Sigma_i)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}(x_i - \mu_i)^T \Sigma_i^{-1} (x_i - \mu_i)\right\}. \quad (7)$$

As evident from above μ_i is the mean of this Gaussian, and Σ_i is the covariance matrix.

The Kalman filter is updated recursively by the following well-known update equations [10]:

$$\mu_i(t) = \bar{\mu}_i(t) + K_i(t)(z_i(t) - G_i\bar{\mu}(t)), \quad (8)$$

$$\Sigma_i(t) = (I - K_i(t)H_i)\bar{\Sigma}_i(t), \quad (9)$$

$$\text{with } \bar{\mu}_i(t) = F_i\mu_i(t-1), \quad (10)$$

$$\text{and } \bar{\Sigma}_i(t) = F_i\Sigma_i(t-1)F_i^T + V_i \quad (11)$$

$$\text{and } K_i(t) = \bar{\Sigma}_i(t)H_i^T(H_i\bar{\Sigma}_i(t)H_i^T + W_i)^{-1} \quad (12)$$

In the above I is the Identity Matrix and K is known as the Kalman gain.

C. Kalman Filter with Maximum Likelihood Data Association

In this section, we introduce data association uncertainty into the tracking problem. We then describe a variant of the Kalman Filter which uses the well-known maximum likelihood data association algorithm.

Above, we implicitly assumed the correspondence between measurements and objects was known. In general, there are $N!$ different ways N objects can be associated with N measurements. We call this unknown correspondence between objects and measurements the *data association* problem.

We represent these $N!$ different data association possibilities by $N!$ permutation matrices which we number C_1 through $C_{N!}$. The ij -th element of $C_k(t)$ is 1 if and only if in the data association represented by C_k measurement $z_j(t)$ originated from object $x_i(t)$. Additionally, we let $k(j)$ represent which object measurement j originated from under data association C_k . And, we let $c_j(t)$ represent which object measurement $z_j(t)$ originated from at time t .

It is convenient to use vector notation to refer to all N objects at once. We let $X(t) = [x_1(t) \dots x_N(t)]^T$, and $Z(t) = [z_1(t) \dots z_N(t)]^T$.

We assume that the entities each have a signature and the measurements contain some identity information. We let $z_i^I(t)$ be the identity information for measurement $z_i(t)$ and $Z^I(t)$ be the vectored version. We assume the distributions $p(z_i(t)|c_i(t))$ are known for all i . Note, this implies we also know $p(Z^I(t)|C_i(t))$.

We let $p(C_k(t))$ be the probability that the data association $C_k(t)$ is correct. This probability can be calculated by:

$$\begin{aligned} p(C_k(t)) &\propto p(Z(t)|C_k(t), X(t))p(Z^I(t)|C_k(t), X(t), Z(t)) \\ &\propto p(Z(t)|C_k(t), X(t))p(Z^I(t)|C_k(t)). \end{aligned} \quad (13)$$

At the same time $p(Z(t)|C_k(t), X(t))$ can be calculated as the product of the N relevant single measurement probabilities.

$$p(Z(t)|C_k(t), X(t)) = \prod_{n=1}^N p(z_n(t)|c_{k(n)}(t), x_{k(n)}(t)) \quad (14)$$

Similarly,

$$p(Z^I(t)|C_k(t), X(t)) = \prod_{n=1}^N p(z_n^I(t)|c_{k(n)}(t), x_{k(n)}(t)) \quad (15)$$

Since multiplication is communicative, the probability of any data association is proportional to a product of N terms. Thus we have (dropping the time index t)

$$p(C_k) \propto \prod_{n=1}^N p(z_n|c_{k(n)}, x_{k(n)})p(z_n^I|c_{k(n)}, x_{k(n)}) \quad (16)$$

Each term $p(z_n|c_{k(n)}, x_{k(n)})$ can be efficiently calculated, since all relevant distributions are Gaussians.

$$\begin{aligned} p(z_n|c_{k(n)}, x_{k(n)}) &\propto \exp(d_m(z_n, \mu_{k(n)})) \\ \text{with } d_m(z_n, \mu_{k(n)}) &= -\frac{1}{2}(z_n - H_{k(n)}\mu_{k(n)})(W + H\Sigma H')^{-1} \\ &\quad (z_n - H_{k(n)}\mu_{k(n)})^T \end{aligned} \quad (17)$$

The quantity $d_m(z_n, \mu_{k(n)})$ is called the Mahalanobis distance.

It is computationally advantageous to take the logarithm of (16), which both changes the product to a sum and removes the exp in (17). Hence

$$\log(p(C_k)) = \text{const} + \sum_{n=1}^N (d_m(z_n, \mu_{c_n}) + \log p(z_i|c_{k(n)})) \quad (18)$$

One standard way to handle data association uncertainty is to assume the most likely data association is correct. Finding the most likely data association can be done in $O(N^3)$ using a bipartite graph matching algorithm [3]. Let $C_*(t)$ represent the most likely data associations at time t .

The Kalman Filter update equations using the most likely data association are identical to those presented in the last section, except the the measurement vector must be permuted. Below we present the update equations for the maximum likelihood KF in vector form: $\mu(t) = [\mu_1(t) \dots \mu_N(t)]^T$ and $\Sigma(t) = \text{diag}([\sigma_1(t) \dots \sigma_N(t)]^T)$. F , G , V , H , and W are all similarly extended.

$$\mu(t) = \bar{\mu}(t) + K(t)(C_*(t)Z(t) - G\bar{\mu}(t)), \quad (19)$$

$$\Sigma(t) = (I - K(t)H)\bar{\Sigma}(t), \quad (20)$$

$\bar{\mu}(t)$, $\bar{\Sigma}(t)$, and $K(t)$ are analogous to the non-vectored versions.

IV. INEFFICIENT THEORETICAL IMKF (TIMKF)

The maximum likelihood data association filter has a drawback – it could have selected the wrong data association. This makes the MLKF (as well as multiple-hypothesis tracker variants that keep the K best data associations) inherently brittle.

The Identity-Management Kalman Filter (IMKF) is a real-time algorithm for tracking that overcomes this fragility. In this section, we explain a computationally infeasible IMKF, which we call a theoretical IMKF (TIMKF). In the next section, we develop an efficient version of this filter.

The key change to our state representation is this: instead of each $[\mu \ \Sigma]$ pair representing the distribution of a single fixed entity, each pair could represent *any* of the entities. The IMKF implicitly maintains a distribution over the $N!$ ways these $[\mu \ \Sigma]$ pairs could be mapped to real object identities. We call this uncertainty between entities in the filter and object identities in the world the *identity mapping* problem.

We represent the $N!$ possible identity mappings at time t by the $N!$ unique permutation matrixes, indexed as M_1 to $M_{N!}$. The ij -th element of $M_k(t)$ is 1 if and only if $[\mu_j(t), \Sigma_j(t)]$ define the distribution of $x_i(t)$. The TIMKF

maintains $p(M_1(t))$ through $p(M_{N!}(t))$ explicitly. Clearly, $\sum_i p(M_i(t)) = 1, \forall t$.

Formally, the TIMKF maintains the parameters $[\mu, \Sigma, p(M_1), \dots, p(M_{N!})]$. The TIMKF implicitly maintains the following pdf (a mixture of $N!$ gaussians):

$$p(X(t)) = \sum_{i=1}^{N!} p(M_i(t)) S_i(t), \quad (21)$$

$$\text{where } S_i(t) = \det(2\pi\Sigma(t))^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}(X(t) - M_i(t)\mu(t))^T M_i(t)\Sigma(t)^{-1} M_i^T(X - M_i(t)\mu(t))\right\}. \quad (22)$$

These changes have added complexity to the Kalman Filter. However, by explicitly modeling the identity mapping problem, the TIMKF can tolerate data association uncertainty. We explain below.

A. TIMKF Exact Posterior after Measurement Update

As the TIMKF's maintained posterior is a mixture of, at most, $N!$ gaussians, the TIMKF is not able to maintain the true posterior (which grows exponentially in time). The TIMKF must thus make an approximation at every time step. We derive the TIMKF updates recursively. That is, we start by assuming the distribution the TIMKF maintains at time $t-1$ is correct. Due to data association ambiguity, the true distribution (which the TIMKF cannot maintain) at time t will have $N!$ times as many mixture components as at time $t-1$.

In deriving the TIMKF update equations, we will first explicitly calculate the exact distribution of $X(t)$, which we denote as $p_e(X(t))$, assuming the TIMKF was correct at time $t-1$ (this intractable distribution will in general have $N!^2$ mixture components). Then we will find a good way to approximate this exact distribution with the parameters the TIMKF can maintain.

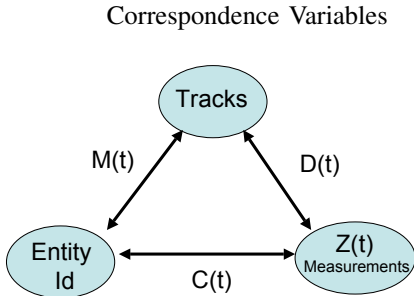


Figure 3. Our 3 correspondence variables. $C(t)$ is the tradition measurement-to-identity data association. $M(t)$ relates the internal IMKF tracks to the entities and is called the identity mapping. $D(t)$ gives the internal track-to-measurement correspondence. Specifying any two of the three determines the third.

In deriving $p_e(X(t))$, it is convenient to consider an alternative way of thinking about data association. We can consider instead *track-to-measurement* associations. As the name suggests, a track-to-measurement association is a correspondence between internal tracks of the TIMKF ($[\mu_i \Sigma_i]$ pairs) and measurements. For every identify mapping there is a different

data association mapping that leads to the same track-to-measurement association. Obviously, there are $N!$ different track-to-measurement associations. We refer to them as $D_1(t)$ to $D_{N!}(t)$. The probability that the k -th track-to-measurement association is correct is $p(D_k(t))$.

Recall from equation (21) that the TIMKF implicitly maintains an $N!$ mixture distribution of X (with the Gaussian corresponding with each identity mapping weighted by its probability). Also recall that at each time step there are $N!$ different data association possibilities. However, given the assumed identify mapping $P_i(t-1)$ was correct as well as the track-to-measurement association $D_j(t)$, the measurement update is simply the standard Kalman Update given in section III-C. We define $R(i, k, t)$ as the distribution arising from updating $S_i(t-1)$ with the Standard KF update assuming track-to-measurement association $D_j(t)$.

Thus, the exact distribution of $X(t)$ after incorporating $Z(t)$ is:

$$p_e(X(t)) = \sum_i \sum_k p(M_i(t-1)) p(D_k(t)) R(i, k, t). \quad (23)$$

This exact distribution has $N!^2$ mixture components, and the TIMKF can only maintain $N!$ mixture components. The ideal TIMKF parameters at time t would minimize the distance between the true distribution and the maintained distribution in some metric of interest.

For tractability reasons, we split TIMKF update problem into two portions. We first select the new positions for the tracks (updating μ and Σ), and then choose the best identity mapping given these tracks.

B. TIMKF μ and Σ update

To update the tracks, the TIMKF uses the standard KF update assuming the most likely data association.

$$\mu(t) \leftarrow \text{updated as eq. (19)}, \quad (24)$$

$$\Sigma(t) \leftarrow \text{updated as eq. (20)}. \quad (25)$$

Recall that μ and Σ implicitly define S_1 though $S_{N!}$ by equation (22). The remaining problem is to select the proper identity mapping distribution, $[p(M_1(t))$ though $p(M_{N!}(t))]$.

C. TIMKF Identify Mapping Update

The maintained distribution of the TIMKF at time t will be a mixture of $N!$ gaussians. The means and variances of these gaussians are defined by μ and Σ , whose update was given in the previous section. As discussed above, the true updated distribution is a n^2 mixture distribution. The identity-mapping update at time t selects the weights on the maintained distribution so it approximates the true updated distribution.

To calculate a probability of a particular identity mapping at time t (i.e. the "weight" on one of the maintained mixtures), we take a weighted sum of how similar each of the $N!^2$ mixtures of the updated true distribution are to the maintained mixture of this particular mapping. We let $\Psi(i, j, k, t)$ represent how much the mixture $R(j, k, t)$ contributes to identity mapping i at time t . Thus:

$$p(M_i(t)) = \sum_k \sum_j p(D_k(t))p(M_j(t-1))\Psi(i, j, k, t)$$

with $\Psi(i, j, k, t) = p(M_i(t)|D_k(t), M_j(t-1))$ (26)

To compute $\Psi(i, j, k, t)$ we assume the corresponding mixture of the true distribution $R(j, k, t)$ is a "measurement" of all N entities with mean μ and variance Σ . We then calculate the probability of this measurement, assuming the entities were distributed according to $S_i(t)$. We compare this to the likelihood of the "measurement," assuming the entities are distributed by each other possible $S_j(t)$.

This calculation is of the same form as the data association calculation of section II. We let $s_i(n, t)$ be the distribution of n th entity of $S_i(t)$ and $r(n, j, k, t)$ be the distribution of the n th entity of $R(j, k, t)$. We let $d_m(s_q(n, t), r(n, j, k, t))$ be the Mahalanobis distance between $s_q(n, t)$ and $r(n, j, k, t)$. Then, suppressing t .

$$\log \Psi(i, j, k) = \frac{\sum_n d_m(s_i(n), r(n, j, k))}{\sum_q \sum_n d_m(s_q(n), r(n, j, k))}. \quad (27)$$

D. Identity-Mapping Refinement

Data association ambiguity causes uncertainty in the position of the objects being tracked. (The true posterior grows exponentially every time step). Identity information in the measurements helps reduce this uncertainty.

The TIMKF uses this identity information to update the identify mappings. Using Bayes rule, the update is:

$$p(M_i(t)) \leftarrow \tilde{Y}^{-1} p(Z^P(t)|M_i(t))p(M_i(t)), \quad (28)$$

$$\text{with } \tilde{Y} = \sum_j p(Z^P(t)|M_j(t))p(M_j(t)). \quad (29)$$

V. IMKF

Algorithm 1 Summary of IMKF

$\mu(t) \leftarrow$ Standard MLKF update
 $\Sigma(t) \leftarrow$ Standard MLKF update
 Calculate T by examining plausible data associations
 $\Omega(t) = \exp T \log \Omega(t-1)$ /* Identity-Mapping Update */
 $\Omega(t) \leftarrow \Omega(t) + \log(\xi(t))$ /* Mapping Refinement */

The TIMKF is intractable for two main reasons. One, it requires explicitly maintaining a distribution over the $N!$ identity mappings, and thus requires $O(N!)$ storage. Two, the update rules eq (26) and eq (27) require considering $O(N!)$ computation. Below, we show how the IMKF overcomes these challenges.

A. Overcoming Storage Intractability

To overcome the storage intractability of the TIMKF, the IMKF stores the track-to-identity probabilities in the $N \times N$ information matrix $\Omega(t)$. The idea of an information form

filter is not new [17], [20]. $\Omega(t)$ implicitly stores the identity mapping probabilities ($M_1(t)$ through $M_{N!}(t)$) by:

$$p(M_i(t)) = \frac{1}{Y} \exp \text{tr} M_i(t)^T \Omega(t), \quad (30)$$

$$\text{where } Y = \sum_{j=1}^{N!} \exp \text{tr} M_j(t)^T \Omega(t). \quad (31)$$

Intuitively, the ij -th element of $\Omega(t)$ represents the strength of our evidence that track i is associated with identity j . This was illustrated in the introductory example. Obviously, $\Omega(t)$ with only N^2 terms cannot express every distribution over the $N!$ identity-mappings. However, it is often a very good approximation, please see [17] for details.

B. Overcoming Computation Intractability

Eq (26) and (27) give the identity-mapping update rule explicitly for all $N!$ possible mappings. As discussed above, the IMKF stores the identity mapping in information form in the $\Omega(t)$. In [17], it is shown that under a Jensen approximation, the update of $\Omega(t)$ is:

$$\Omega(t) = \exp T \log \Omega(t-1) \quad (32)$$

Where $T(t)$ is a $N \times N$ transition matrix. In words, the ab th entry of $T(t)$ describes how much the marginal of the identify of track a at time $t-1$ will contribute to the marginal of track b at time t .

Each element of $T(t)$ is a sum of the relevant identity-mapping probabilities, (given by eq (27)). The summation over $N!$ different data associations is made tractable by only considering plausible data associations. These associations are found by examining the simple circuits on a certain directed graph. See [18].

Calculating $\Psi(i, j, k)$ exactly requires $O(N!)$ computation, due to calculating the denominator of (30). However, in general $\Psi(i, j, k)$ can be approximated quickly because the Mahalanobis distance between the distributions $s_i(n)$ and $r(n, j, k)$ is often small. In particular, recall that both $s_i(n)$ and $r(n, j, k)$ arise from a KF update. If they share neither a common prior nor a common measurement, the Mahalanobis distance between the posteriors is likely small. See [18] for details.

C. IMKF Identity-Mapping Refinement

The Identity-Mapping Refinement step can be implemented exactly (see [17] for derivation.) We let $\xi(t)$ be a $N \times N$ matrix with the ij -th element equal to the probability that the identify measurement associated with track i originated from object j . Then:

$$\Omega(t) \leftarrow \Omega(t) + \log(\xi(t)). \quad (33)$$

Again intuitively, the IMKF refinement update utilizes the identity information in the measurements to update the posterior to reduce the uncertainty added by data association ambiguities.

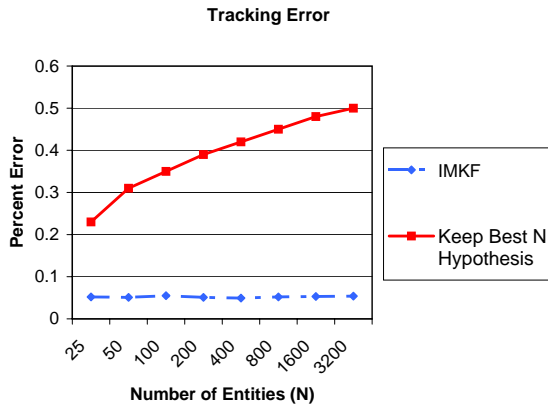


Figure 4. Scaling properties of our filter

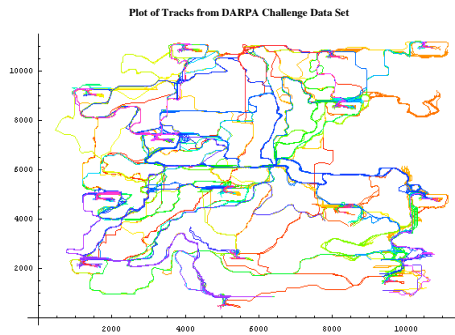


Figure 5. Tracks from the DARPA Challenge Data

VI. RESULTS

We tested the scaling of our approach on synthetic data. Here we significantly outperform a particle-filter based approach.

We also applied our algorithm to a recent DARPA-designed tracking problem, which is being used in the DARPA ACIP program as the main challenge problem. This problem contains over 3,000 moving objects moving in a realistic fashion. The entities being tracked do not follow a simple motion model, but rather move in a complex environment completing useful tasks. For example, a simulated delivery truck starts at a warehouse and moves to a retail store following roads. We choose this data set because it enabled us to test our approach on a very large data set $N = 3,000$ that was much better than a random simulation. Results given are averaged over 50 runs.

For each entity tracked, we kept both positional and velocity state-information (thus we had a 4 dimensional state-space). We set F to assume a constant velocity model, with a relatively large noise term (V) to account for the turns and accelerations actual entities take periodically.

We perform quite well on this data. We consider an object

Urban Environment of DARPA Challenge Data




Figure 6. Illustration of Urban Environment in DARPA Challenge Data




Figure 7. Comparative results on Challenge Data

to be correctly tracked if a track's closest ground-truth entity (measured by the L2 metric in position only) matches the most likely identity of this track. When tracking 200 entities, we have an error of 4%, lower than the 10.5% for the MHT with 200 hypotheses, and much lower than the standard kalman filter, which has above a 90% error rate. The last figure shows that the Identity Management Kalman Filter has low error even when tracking 3,000 entities.

The results on the DAPRA challenge data seem to validate our intuition given in section 1. When entities interact in a way that causes data association ambiguities, the Ω machinery of the IMKF enables it to retain enough information to quickly recover. MHT have no such safety mechanism, and thus they do not recover quickly.

VII. SUMMARY

We have presented a new filter called the identity management Kalman filter (IMKF) for tracking large number of objects under uncertain data associations. The IMKF maintains three statistics, a mean and a covariance familiar from the conventional Kalman filter, as well as an identity mapping matrix that associates internal tracks with physical objects in the real world. By augmenting the KF with such an identity mapping matrix, the IMKF can maintain full posteriors over the combinatorial space of associations between the internal




Figure 8. Scaling Results on Challenge Data

filter tracks, and the objects in the world. The form of this matrix makes the update efficient, and it collapses an otherwise exponential probability distribution into quadratic space. While the underlying idea of the information matrix for data association is not new, the IMKF implements this idea for the first time in the context of full continuous tracking. This paper gives the mathematical derivation from first principles, and discusses the key approximations necessary to obtain such a compact filter.

In experimental results we find that the IMKF can solve tracking problems with orders of magnitude more objects than previously possible, even when significant identity uncertainty is present. We believe that the IMKF provides a solution to the outstanding problem of developing multi-object tracking algorithms that scale to very large state spaces.

Obviously, the description of the IMKF in this paper is limited in many ways. For example, it does not address objects entering or leaving the area of concern or the problem of false alarms. Also the derivations in this a paper assume linear system dynamics and linear measurement models. We believe these extensions are relatively straightforward, and will be analogous to the rich literature on nonlinear filtering.

We also observe that as N becomes large, the matrix Ω naturally becomes sparse. This suggests that for large N , the essential update steps can be implemented in $O(N)$ time. Further work investigating this $O(N)$ solution could yield interesting results.

Acknowledgements. The authors would like to thank Chang Liu and Lu Ruan for code to efficiently find simple cycles on digraph [9].

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA) or the U.S. Government. Effort sponsored by the Defense Advanced Research Projects Agency (DARPA) through the Department of the Interior National Business Center under grant number NBCH104009.”

Leonidas Guibas wishes to acknowledge the support of DoD Multidisciplinary University Research Initiative (MURI)

program administered by the Office of Naval Research under Grant N00014-00-1-0637, NSF grant CNS-0435111, and ARO grant W911NF-04-R-0005.

REFERENCES

- [1] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.
- [2] Y. Bar-Shalom and X.-R. Li. *Estimation and Tracking: Principles, Techniques, and Software*. YBS, Danvers, MA, 1998.
- [3] T. Corman, C. Leiserson, R. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2001.
- [4] F. Dellaert, S.M. Seitz, C. Thorpe, and S. Thrun. EM, MCMC, and chain flipping for structure from motion with unknown correspondence. *Machine Learning*, 50(1-2):45–71, 2003.
- [5] A. Doucet, J.F.G. de Freitas, and N.J. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer Verlag, New York, 2001.
- [6] C. Guestrin, R. Thibaux, P. Bodik, M.A. Paskin, and S. Madden. Distributed regression: an efficient framework for modeling sensor network data. In *Proceedings of Information Processing in Sensor Networks*, 2004.
- [7] A.M. Jazwinsky. *Stochastic Processes and Filtering Theory*. Academic, New York, 1970.
- [8] S. Julier and J. Uhlmann. A new extension of the Kalman filter to nonlinear systems. In *International Symposium on Aerospace/Defense Sensing, Simulate and Controls*, Orlando, FL, 1997.
- [9] Chang Liu and Lu Ruan. Finding Good Candidate Cycles for Efficient p-Cycle Network Design. In *ICCCN*, Chicago, 2004.
- [10] Peter S. Maybeck. *Stochastic models, estimation, and control*, volume 141 of *Mathematics in Science and Engineering*. 1979.
- [11] P.S. Maybeck. The Kalman filter: An introduction to concepts. In I.J. Cox and G.T. Wilfong, editors, *Autonomous Robot Vehicles*. Springer Verlag, 1990.
- [12] Songhwai Oh and Shankar Sastry. AA Polynomial-Time Approximation Algorithm for Joint Probabilistic Data Association. In *Proc. American Control Conference (ACC)*, Portland, OR, June 2005.
- [13] H. Pasula, S. Russell, M. Ostland, and Y. Ritov. Tracking many objects with many sensors. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Stockholm, Sweden, 1999. IJCAI.
- [14] M. Pitt and N. Shephard. Filtering via simulation: auxiliary particle filter. *Journal of the American Statistical Association*, 94:590–599, 1999.
- [15] D.B. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Aerospace and Electronic Systems*, AC-24:843–854, 1979.
- [16] D. Schulz, W. Burgard, D. Fox, and A. Cremers. Tracking multiple moving targets with a mobile robot using particle filters and statistical data association. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [17] B. Schumitsch, S. Thrun, G. Bradski, and K. Olukotun. The information-form data association filter. In *Proceedings of Conference on Neural Information Processing Systems (NIPS)*, Cambridge, MA, 2005. MIT Press.
- [18] B. Schumitsch, S. Thrun, L. Guibas, and K. Olukotun. The identity management kalman filter (imkf). Stanford EE Technical Report 2006-Schumitsch01, 2006.
- [19] J. Shin, L. Guibas, and F. Zhao. A distributed algorithm for managing multi-target identities in wireless ad-hoc sensor networks. In *Proc. 2nd International Conference on Information Processing in Sensor Networks (IPSN03)*, pages 223–238, Palo Alto, CA, April 2003. Springer.
- [20] J. Shin, N. Lee, S. Thrun, and L. Guibas. Lazy inference ob object identities in wireless sensor networks. In *Fourth International Conference on Information Processing in Sensor Networks*, Los Angeles, CA, 2005.
- [21] P. Swerling. A proposed stagewise differential correction procedure for satellite tracking and prediction. Technical Report P-1292, RAND Corporation, January 1958.
- [22] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, Cambridge, MA, 2005.