Stochastic Sampling Based Data Association

Matthew Travers[†], Todd Murphey[†], and Lucy Pao[‡]

Abstract— **This paper considers how to determine the origin of a single measurement originating from one of a group of objects moving in close proximity. During the time in which measurements are being received, the dynamics of the various objects are the same except for initial conditions. We present a method that uses techniques from filtering theory to represent a distribution using a finite number of parameters. This method, which we call stochastic sampling based data association (SSBDA), is similar to a particle filter but differs in that we use a modified probabilistic data association filter (PDAF) in the propagation of the distribution associated with the object's location. Using the PDAF it is possible to see the effect that the addition of each measurement has on the covariance of the posterior distribution. We discuss how the covariance of the posterior can be used for making decisions on whether or not a particular measurement originated from a predetermined object of interest.**

I. Introduction

Data association has been used in a wide variety of tracking scenarios, typically focusing on multiple measurements from multiple objects. We consider a specific version of this problem: objects traveling in close proximity. This problem is difficult because the objects' dynamics are essentially the same, the only difference between the objects are initial conditions.

In the problem considered, only a single measurement is received at each time step. The measurement could potentially have originated from any of the objects within the sensor's range. However, we know only the initial distribution associated with a predetermined object of interest.

The main objective of the problem being considered is to either positively associate the measurement at each time step with the object of interest or not. Pre-existing data association methods, such as gating and the Kolmogorov-Smirnov test [13], have attempted to solve this problem. The method, which we refer to as stochastic sampling based data association (SSBDA) is necessary because of the constant close proximity of the various objects in the system combined with the sparseness of measurements.

An intuitive choice of existing methods to solve this problem is gating. In gating, a "gate" or validation region is drawn around the expected location of the object of interest. Measurements that fall within the gate are associated with the object and measurements that fall outside are not. In the systems we consider, it will often be the case that measurements not originating from the object of interest will fall within the validation region and thus would be incorrectly associated with the object of interest. Other than the possibility of incorrectly associating measurements due to the objects close relative proximity, there is another difficulty with the gating algorithm.

In the gating algorithm, as well as data association algorithms where clear associations (i.e., acceptance/rejection) must be made, the user must define some threshold for acceptance. For example, in gating the user must define the size of the validation region. This choice has a tradeoff between incorrect measurements that are positively associated and correct measurements that are not. The primary contribution of this paper is the fact that SSBDA removes the need for a user-defined acceptance region. Like any Bayesian-based filtering algorithm, SSBDA has two steps: prediction and update. In SSBDA, the prediction step is different from most filtering algorithms in the fact that we employ stochastic integration to map forward in time each of the finite parameters that represent the prior distribution. In the update step we use the PDAF to incorporate the addition of the measurement. There is a distribution associated with each of these steps, and thus a covariance. If the covariance of the distribution associated with the update is not smaller than the covariance of the distribution associated with the prediction, the measurement at that time is rejected. The covariances themselves are thus the "metric" upon which the data is associated.

We will consider two example systems in which SSBDA is applied, one linear and one nonlinear. The results will show no discernible difference in SSBDA performance between the two systems in terms of the ability to identify incorrect measurements.

Sample-based data association methods for nonlinear systems have been previously covered in the literature. In [11], a standard particle filter is modified to incorporate probabilities of association. This process is referred to as a joint probabilistic data association filter (JPDAF) with samples. Due to the similarities between this method and SSBDA, we will compare their filtering performances in Section III. In [5],[6],[9], and [12], Markov Chain-Monte Carlo methods (MCMC) for data association are presented. Comparison of SSBDA to MCMC data association is a direction of future work. The major difference between SSBDA and these pre-existing methods is the use of stochastic integration in the prediction step of the filtering scheme.

This paper is organized as follows. In Section II, we introduce stochastic integration, the PDAF, and sample-based methods by giving a brief description of each. In Section III, the SSBDA algorithm is introduced. Analytical results are presented, followed by an explanation of the SSBDA

[†]Department of Mechanical Engineering, Northwestern University, Evanston, Illinois 60208

[‡]Department of Electrical, Computer, and Energy Engineering, University of Colorado at Boulder, Boulder, Colorado 80309

algorithm. Simulation results demonstrating the performance of the SSBDA algorithm are also presented. Finally, Section IV includes conclusions and a discussion of future work.

II. Mathematical Methods

The SSBDA algorithm uses several mathematical methods. We will briefly cover stochastic integration, the PDAF, and using a finite number of samples to represent a distribution. A more thorough treatment of these subjects can be found in [1],[4],[7], and [8].

A. Stochastic Differential Equations

A stochastic dynamical system consists of a base flow on a probability space $(\Omega, F, \mathbf{P})$ and a deterministic flow. The probability space triple $(\Omega, F, \mathbf{P})$ consists of the set Ω which is the sample space, the σ-algebra F which consists of subsets of Ω, and the probability measure $\mathbf{P}$. In this work, the elements $\omega \in \Omega$ are assumed to be derived from a Wiener Process W.

Assume that the stochastic dynamical system evolves in an ℓ dimensional vector space M with vector fields $X_i \in X^\ell(M)$ for $i = 0, ..., m$. The general form of the stochastic differential equation is

$$dx = X_0(z)dt + \sum_{i=1}^{m} X_i(z) \circ dW_i, \;\; x(0) = x_0 \qquad (1)$$

where X_0 is the 'drift vector field' and $X_i, i = 1, ..., m$ are the 'diffusion vector fields'.

B. Probabilistic Data Association Filter (PDAF)

In the SSBDA method, stochastic integration is combined with the PDAF algorithm. The following are the discrete-time equations of the probabilistic data association filter. The results are well known and are thus stated merely for convenience; for a formal derivation, see [1].

$$\begin{aligned}
\hat{x}(k|k-1) &= A(k)\hat{x}(k-1|k-1) \\
P(k|k-1) &= A(k)P(k-1|k-1)A(k)^T + Q(k-1) \\
\hat{z}(k|k-1) &= H\hat{x}(k|k-1) \\
\nu(k) &= z(k) - \hat{z}(k|k-1) \\
S(k) &\triangleq HP(k|k-1)H^T + R(k) \\
W(k) &\triangleq P(k|k-1)H^T S^{-1}(k) \\
\hat{x}(k|k) &= \hat{x}(k|k-1) + \beta_i W(k)\nu(k) \\
P(k|k) &= [I - W(k)H]P(k|k-1)
\end{aligned}$$

where $\hat{x}(k|k-1)$ is the state prediction and $A(k)$ the linear dynamics. $P(k|k-1)$ is the covariance associated with the state predicion, $Q(k-1)$ is the process noise covariance, $\hat{z}(k|k-1)$ is the measurement prediction, and H is the output map (assumed constant). The residual is $\nu(k)$ and its associated covariance is $S(k)$, and $R(k)$ is the measurement noise covariance. The Kalman filter gain is $W(k)$, β_i is the probability of association, $\hat{x}(k|k)$ is the current state estimate and $P(k|k)$ its covariance.

The probability of association is central to creating the effect that we desire in SSBDA when a false measurement is present. The probability of association is calculated by first finding the following parameters: $\lambda = \frac{m_k}{V_k}$, where V_k is the volume of the validation region, m_k is the total number of expected measurements in the validation region. λ is the density of measurements in the validation region. $V_k = c_{n_z}\gamma^{n_z/2}|S(k)|^{1/2}$, and c_{n_z} is the volume of the n_z-dimensional unit hypersphere, n_z being such that $z(k) \in \mathbb{R}^{n_z}$, and γ is the square of the number of standard deviations to allow into the validation region. The $\beta_i(k)$ probabilities, where $i = 1, 2, \ldots, m_k$, are then computed as [1]

$$\beta_i(k) = \frac{\epsilon_i}{b + \sum_{j=1}^{m_k} \epsilon_j} \qquad (2)$$

where

$$\epsilon_i = \exp\left[-\frac{\nu_i'(k)S^{-1}(k)\nu_i(k)}{2}\right],$$

and

$$b \triangleq (2\pi/\gamma)^{n_z/2}\lambda V_k c_{n_z}(1 - P_D P_G)/P_D.$$

In (2) above, b is an "acceptance region" parameter which encodes the geometry of the region where measurements are most likely to fall into the PDAF algorithm. The parameters P_D and P_G are the probability of detection and probability of gating, respectively.

Note that in standard use there may be clutter in the environment that creates erroneous measurements inside of the validation region. In the examples that we consider we assume that there is no clutter, more specifically $m_k = 1$ always. No gating is used.

C. Sampling-Based Methods

The SSBDA method presented in the following section is a sampling-based approach to the data association problem. In sampling-based methods, a set of samples is generated which is used to approximate a distribution. We will show below that in SSBDA there are two distributions of interest, $p(x(k)|x(k-1))$ and $p(x(k)|z_k)$. These two distributions are the distributions associated with the prediction and the update, respectively. The corresponding particle sets will be referred to as $\chi(k|k-1)$ and $\chi(k|k)$, where, for example

$$\chi(k|k) := \{\hat{x}_1(k|k), \hat{x}_2(k|k), \ldots, \hat{x}_N(k|k)\}.$$

Note that each of the $\hat{x}_i(k|k)$ is a realization of the distribution $p(x(k)|z(k))$.

The underlying concept behind using sampling-based approaches has been well covered in the Monte Carlo methods literature ([2],[3],[8],[10]) One key aspect is that the expectations with respect to the distribution $p(x(k)|data)$ are approximated by

$$\int f(x(k))p(x(k)|data)dx(k) \approx \frac{1}{N}\sum_{i=1}^{N} f(x_i(k)), \qquad (3)$$

where N is the number of samples and $f(\cdot)$ is the function associated with various moments. Note that as the number N gets large, the error in this approximation tends to zero.

III. SSBDA

In this section we present the SSBDA algorithm. The main objectives of this algorithm are to use a sampling-based approach combined with stochastic integration and PDAF to represent distributions, and either reject or accept measurements based on a comparison of covariances. At each time step two distributions are found, one associated with the prediction and one associated with the update. The covariances of both of these distributions are found. If the covariance of the distribution associated with the update is not smaller than the covariance associated with the prediction, the measurement is assumed to add no information to the system and is thus rejected as having originated from the object of interest.

In this section we present some analytical results for the SSBDA algorithm. In particular, we will show under what conditions this algorithm will work. Note that the following results assume the state space is $\mathbb{R}^n$.

The condition for rejecting a measurement is that the covariance of the distribution associated with the prediction is smaller than or equal to the covariance of the distribution associated with the update, i.e.,

$$||\sigma(p(x(k)|z(k)))||_F \geq ||\sigma(p(x(k)|x(k-1)))||_F, \qquad (4)$$

where σ is the covariance and $||\cdot||_F$ is the Frobenius norm.

Lemma I For $\hat{x}_i(k|k)$ the i-th update, $\hat{x}_i(k|k-1)$ the i-th prediction, $\bar{x}(k|k)$ the average over all updates at time k, $\bar{x}(k|k-1)$ the average over all predictions and $||\cdot||$ the Euclidean norm, if

$$||\hat{x}_i(k|k)-\hat{x}_j(k|k)|| \geq ||\hat{x}_i(k|k-1)-\hat{x}_j(k|k-1)|| \; \forall i,j \leq N \qquad (5)$$

then

$$||\hat{x}_i(k|k) - \bar{x}(k|k)|| \geq ||\hat{x}_i(k|k-1) - \bar{x}(k|k-1)|| \; \forall i.$$

Proof: Equation (5) is true for all j. As $N \to \infty$, $\hat{x}_j(k|\cdot) = \bar{x}(k|\cdot)$ for some j *almost surely*. ■

Lemma II For $\hat{x}_i(k|k-1)$ the i-th prediction, $W(k)$ the filter gain, $\beta_i(k)$ the i-th probability of association, and $\nu_i(k)$ the i-th innovation, all at time k, if $\theta(k)$ is the angle between the two vectors $\hat{x}_i(k|k-1)-\hat{x}_j(k|k-1)$ and $W(k)(\beta_i(k)\nu_i(k)-\beta_j(k)\nu_j(k))$ defined by the Euclidean norm, if

$$||\hat{x}_i(k|k-1) - \hat{x}_j(k|k-1)|| \geq$$
$$-\frac{1}{2\cos\theta(k)}||W(k)(\beta_i(k)\nu_i(k) - \beta_j(k)\nu_j(k))|| \quad \forall i,j \leq N, \qquad (6)$$

then

$$||\hat{x}_i(k|k)-\hat{x}_j(k|k)|| \geq ||\hat{x}_i(k|k-1)-\hat{x}_j(k|k-1)|| \; \forall i,j \leq N.$$

Proof:

$$\begin{aligned}||\hat{x}_i(k|k) - \hat{x}_j(k|k)|| = ||&\hat{x}_i(k|k-1) - \hat{x}_j(k|k-1)\\ &+ W(k)(\beta_i(k)\nu_i(k) - \beta_j(k)\nu_j(k))||\end{aligned}$$

SSBDA$(x_0, P_0, z_{1\ldots}, N)$
for $k = 1, 2, \ldots$ **do**
  for $i = 1, 2, \ldots N$ **do**
    $\hat{x}_i(k|k-1) = A(k)x_i(k-1|k-1)$
      $+ \int_{k-1}^{k} X(c(s,\omega)) \circ dW(s,\omega)$
    $P(k|k-1) = A(k)P(k-1|k-1)A(k)^T + Q(k-1)$
    $S(k) = HP(k|k-1)H^T + R(k)$
    $W(k) = P(k|k-1)H^T S(k)^{-1}$
    $\nu_i(k) = z(k) - H\hat{x}_i(k|k-1)$
    $\epsilon_i(k) = \exp(-\nu_i(k)^T S(k)\nu_i(k))/2$
    $\beta_i(k) = \epsilon_i(k)/(b + \epsilon_i(k))$
    $\hat{x}_i(k|k) = \hat{x}_i(k|k-1) + \beta_i(k)W(k)\nu_i(k)$
    $P(k|k) = (I - W(k)H)P(k|k-1)$
  $\chi(k|k-1) = \{\hat{x}_1(k|k-1), \hat{x}_2(k|k-1), \ldots, \hat{x}_N(k|k-1)\}$
  $\chi(k|k) = \{\hat{x}_1(k|k), \hat{x}_2(k|k), \ldots, \hat{x}_N(k|k)\}$
  if $covariance(\chi(k|k-1)) \leq covariance(\chi(k|k))$ **then**
    $\chi(k|k) = \chi(k|k-1)$
    Drop z_k
return $\chi(k|k)$

TABLE I: Stochastic sampling based data association algorithm.

$$\begin{aligned}=&[(\hat{x}_i(k|k-1) - \hat{x}_j(k|k-1))^T(\hat{x}_i(k|k-1) - \hat{x}_j(k|k-1))\\ &+ 2(\hat{x}_i(k|k-1) - \hat{x}_j(k|k-1))^T W(k)(\beta_i(k)\nu_i(k)\\ &- \beta_j(k)\nu_j(k)) + (\beta_i(k)\nu_i(k) - \beta_j(k)\nu_j(k))^T W^T(k)\\ &\cdot W(k)(\beta_i(k)\nu_i(k) - \beta_j(k)\nu_j(k))]^{1/2} \qquad (7)\end{aligned}$$

Note that $||\hat{x}_i(k|k-1) - \hat{x}_j(k|k-1)||^2$ is the first term in (7). Note also that $||W(k)(\beta_i(k)\nu_i(k) - \beta_j(k)\nu_j(k))||^2$ is the last term in (7). When (6) is satisfied,

$$||\hat{x}_i(k|k) - \hat{x}_j(k|k)|| \geq ||\hat{x}_i(k|k-1) - \hat{x}_j(k|k-1)|| + c$$

where c is a constant such that $c \geq 0$, and thus

$$||\hat{x}_i(k|k) - \hat{x}_j(k|k)|| \geq ||\hat{x}_i(k|k-1) - \hat{x}_j(k|k-1)||.$$

■

Proposition I If

$$||\hat{x}_i(k|k-1) - \hat{x}_j(k|k-1)|| \geq$$
$$-\frac{1}{2\cos\theta(k)}||W(k)(\beta_i(k)\nu_i(k) - \beta_j(k)\nu_j(k))|| \quad \forall i,j \leq N,$$

then

$$||\sigma(p(x(k)|z(k)))||_F \geq ||\sigma(p(x(k)|x(k-1)))||_F.$$

Proof: If condition (6) in Lemma II is satisfied, using Lemma I, and noting that $||\sigma(p(x|\cdot))||_F$ can be shown to be equal to

$$||\sigma(p(x|\cdot))||_F = ||\hat{x}_i - \bar{x}||^2, \qquad (8)$$

Proposition I comes directly. ■

A. Algorithm

Table I presents the SSBDA algorithm. The inputs into this algorithm are the initial distribution, represented by x_0 and P_0, the measurements, $z_1, z_2, \ldots$, and the number N of samples to use. We assume that the first measurement occurs at time $k = 1$. To "create" the samples at time $k = 1$, we use the initial state x_0 combined with stochastic integration. We start at x_0 and perform N forward stochastic integrations up to the time at which the first measurement occurs. The result of performing these N integrations is a sample-based representation of the distribution associated with the prediction $p(x(1)|x(0))$, of which (3) can be used to approximate the various moments. After performing the N stochastic integrations, the PDAF algorithm is used to "update" each of N predictions. What is meant by "update" is that the measurement at time k is incorporated, through convolution, into the distribution associated with each of the samples. Note that the covariances and gain in the PDAF are the same for each sample. The only difference from sample to sample is the predicted state $\hat{x}_i(k|k-1)_i$, the innovation $\nu_i(k)$, the probability function $\epsilon_i(k)$, the probability of association $\beta_i(k)$, and the updated state $\hat{x}_i(k|k)$.

After running the modified PDAF algorithm for each of the N samples, the two sets $\chi(k|k-1)$ and $\chi(k|k)$ are formed. These two sets are approximations to the distributions $p(x(k)|x(k-1))$ and $p(x(k)|z(k))$, or the distributions associated with the prediction and update, respectively. Note that the process of approximating these two distributions is parallel to the particle filter. The difference is that in the particle filter, the measurement is incorporated into the filtered state by resampling. In SSBDA there is no resampling, the measurement is incorporated into the filtered estimate by updating the state of each of the N samples with a PDAF.

Forming the sets $\chi(k|k-1)$ and $\chi(k|k)$ allows us to approximate the distributions associated with the prediction and update of our current state at time k. These two distributions are used to perform the main objective of SSBDA. If the measurement does not improve the estimate of the current state, then the measurement originated from another nearby object. To improve the estimate of the current state, the addition of the measurement results in a smaller covariance of the distribution associated with the update than that associated with the prediction. This means that under SSBDA, if (4) is satisfied the measurement at time k is rejected. Note that (4) being satisfied is equivalent to (6) holding, which is the test condition for the SSBDA algorithm.

B. Simulated Results

In this section we present simulated results of applying the SSBDA method to dynamic systems examples. The two systems that we consider both contain three cars constrained to travel on a road in a "convoy." They remain in close proximity throughout the entire period in which measurements are received. There is no relative maneuvering between the cars. The dynamics for the cars are exactly the same in simulation except for initial conditions. The measurement is of the car's position and occurs every $0.1s$. The car of interest is the car in the middle, Car 2, as shown in Figure 1. A sensor is pointed at Car 2, but can occasionally accidently measure one of the other cars instead.

We consider two types of roads, one flat and straight, the other flat and curved. In the results presented below we have assumed that the measurements in the x and y directions are uncorrelated. We are also assuming that the various

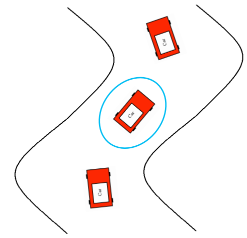

Fig. 1: Convoy of three cars. The car in the middle, Car 2, is the car of interest. Note that in the examples, the vertical direction is the x axis.

noises due to the system are uncorrelated (these assumptions are not necessary to the success of the method, but reduce the number of computations needed for the examples). We compare the relative sizes of two covariance matrices with the Frobenius norm.

The results of applying SSBDA to the linear convoy of cars problem can be seen in Figure 2. The plotted positions (the points which make up the black lines) in this figure correspond to update positions. In this example the number of samples is 75 (which implies that there are 75 black lines), i.e., $N = 75$. The green points are the measurements that originated from Car 2. The red points are the measurements that originated from one of the other two cars in the convoy, in this case Car 3. The horizontal axis is time and the vertical axis is the longitudinal x position along the road.

As measurements are received, each of the N-samples traces out a trajectory. Figure 3 shows the variance in the $x-$direction over time for both the prediction (solid blue line) and the update (dashed orange line). In Figures 2 and 3, it can be seen that the covariance associated with the update goes up due to the addition of the incorrect measurements. Figure 3 shows that the covariance of the update distribution is in fact higher than the covariance of the prediction's distribution at these time steps where incorrect measurements are received.

Figures 4 and 5 show the same results as those seen in Figures 2 and 3, correspondingly, but in this case the road to which the cars are constrained is nonlinear (note that in Figure 4 the vertical axis is now the $y-$component of Car 2's trajectory). We are again only looking at a single dimension because we are making the same assumptions about the noise as we did in the linear case.

In Figure 4 we can see the same "spreading out" of the stochastic trajectories that was present in Figure 2. Figure 5 confirms that for this nonlinear example, the variance of the distribution associated with the update is in fact higher than the variance associated with the prediction at the times erroneous measurements are received.

Figure 6 shows the reason that a PDAF is used in the SSBDA algorithm as opposed to a Kalman filter. In Figure 6, the orange dot-dash line represents the mean of all 75 sample

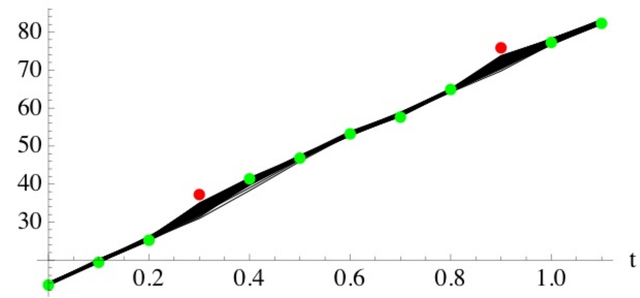

Fig. 2: SSBDA using a PDAF for the linear system. The red points are erroneous measurements, the green points are correct measurements, and the black lines are the stochastic trajectories.

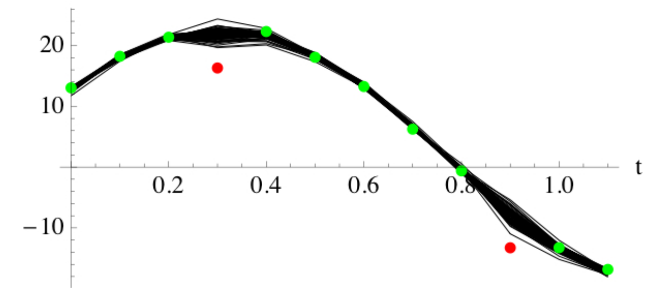

Fig. 4: SSBDA using a PDAF for the nonlinear system. The red points are erroneous measurements, the green points are correct measurements, and the black lines are the stochastic trajectories.

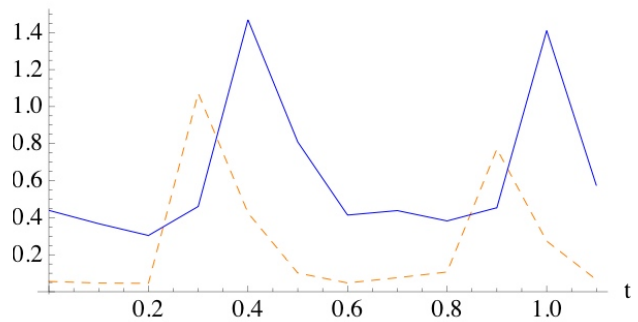

Fig. 3: Variances of the distributions associated with the update (dashed orange) and the prediction (solid blue) for the linear system. The horizontal axis is time and the vertical axis represents the variance in the $x-$direction.

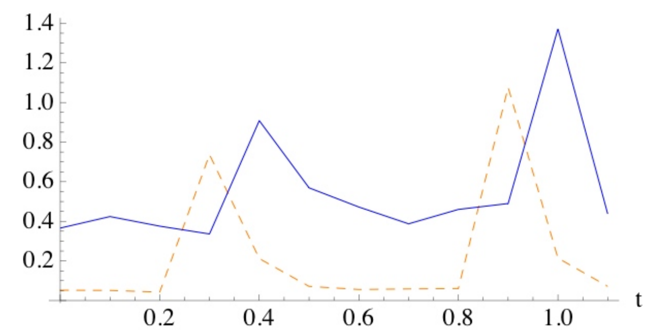

Fig. 5: Variances of the distributions associated with the update (dashed orange) and the prediction (solid blue) for the nonlinear system. The horizontal axis is time and the vertical axis represents the variance in the $y-$direction.

trajectories at each time step for the linear example (i.e., this orange line is the mean at each time step of the trajectories in Figure 2). The blue dotted line is again a mean trajectory over 75 individual trajectories, but in this case the SSBDA algorithm was run using a Kalman filter in the update step as opposed to a PDAF. Using the Kalman filter, each sample trajectory experiences a larger effect in its filtered path due to the erroneous measurement. The result is that the overall average position of the filtered path is drawn away from the true trajectory by the occurrence of an erroneous measurement.

Figure 6 shows that using a Kalman filter in the SSBDA algorithm skews the filtered trajectory in an undesirable way around the erroneous measurements. Due to this skewing, it is impossible to make correct associations. The standard PDAF is a filtering method that takes measurement proximity into account when forming an updated state estimate (which is unlike the Kalman filter). Although the PDAF does not include an engine for clear acceptance/rejection of measurements, a comparison in filtering performance between a PDAF and SSBDA is of interest.

A comparison between a JPDAF with samples [11] and SSBDA is presented in Figures 7 and 8. In these results, filtering performance is designated by RMS error between the filtered trajectory and the simulated true trajectory of Car 2. In Figure 7, the three cars are close together, where close is designated as within two standard deviations of each other (based on the measurement noise covariance). In this figure, the solid blue line represents the JPDAF with samples and the dotted red line SSBDA. It can seen by inspection that the JDAF with samples does slightly better than SSBDA, but the performance is comparable. In Figure 8, everything is the same as Figure 7, except that for this case the cars are far apart, where far is designated as outside of three standard deviations. In this case, SSBDA maintains the same performance characteristics as in the close proximity case, but the JPDAF with samples has error that gets larger over time. The degradation of the filtering performance of the JPDAF with samples is due to the structure of this problem. Due to the large relative proximity of the single incorrect measurement and the predicted state estimate at each time step, the variance of the resulting update distribution becomes artificially skewed through the resampling process. This skewing leads to the increasingly bad performance of the JPDAF with samples as more measurements are received, which is shown in Figure 8.

In these two examples comparing performance between SSBDA and the JPDAF with samples, the three cars were either all in close proximity or relative far proximity. In practice, there may be more than three cars with varying

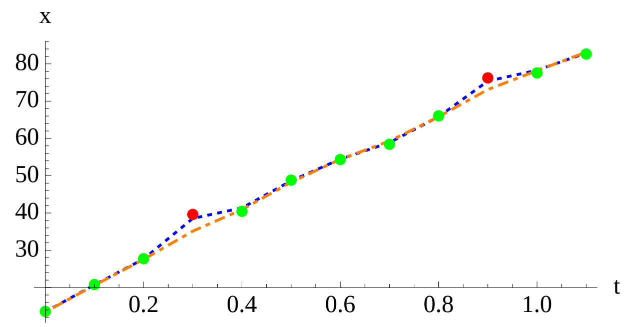

Fig. 6: The dot-dashed orange line represents the mean trajectory of the SSBDA algorithm using a PDAF in the update step. The dotted blue line represents the mean trajectory of the SSBDA algorithm using a Kalman filter in the update step, i.e., does not include any probability of association.

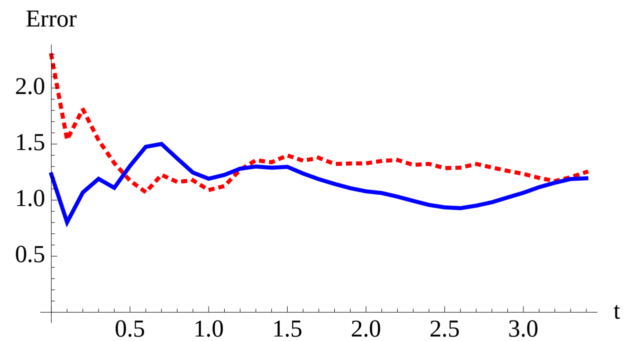

Fig. 7: RMS errors for a JPDAF with samples (solid blue) and SSBDA (dotted red) for the nonlinear system when the cars are close together (inside two standard deviations).

degrees of proximity and it may not be clear what the threshold between close and far is. For this reason we wish to have a single algorithm that works regardless of the which object produces the measurement.

IV. Conclusions & Future Work

By taking a sample-based approach and using stochastic integration along with a PDAF, we showed that when Proposition I holds, SSBDA produces a distribution whose covariance does go up. We also showed that the size of this covariance can be used to reject measurements that did not originate from the object of interest.

At each time step, the SSBDA algorithm uses a sample-based approach to represent two distributions, one associated with the prediction step and the other with the update. The covariance of these two distributions are compared as a basis for acceptance/rejection of measurements. The SSBDA algorithm and this simple acceptance/rejection calculation are the main contribution of this work.

Analytical as well as simulation results for two separate systems were given. One of these systems was linear and the other nonlinear. The results showed that SSBDA performance was robust to the addition of the nonlinear dynamics.

In another set of results, the SSBDA algorithm of Table I is modified to use a Kalman filter instead of a PDAF. These

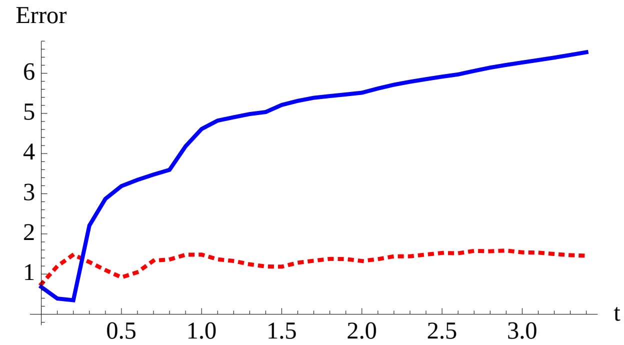

Fig. 8: RMS errors for a JPDAF with samples (solid blue) and SSBDA (dotted red) for the nonlinear system when the cars are far apart (outside 3 standard deviations).

results show that the Kalman filter based version sees more of a negative effect due to erroneous measurements. We also compared the JPDAF with samples and SSBDA. We showed that, in terms of nonlinear filtering, the SSBDA algorithm performed well regardless of object proximity, while the JPDAF with samples had errors that grow with time when the object relative proximity was large.

The extension of the SSBDA algorithm to a wider variety of nonlinear systems is a focus of future work, which includes nonlinear dynamics as well as non-Gaussian stochastic forcing. We are also exploring the use of MCMC data association and how it compares to SSBDA.

References

[1] Y. Bar-Shalom and T.E. Fortmann. *Tracking and Data Association*. Academic Press, Inc., 1988.

[2] A. Doucet, N. De Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer, 2001.

[3] A. Doucet, S. Godsill, and C. Andrieu. Sequential Monte Carlo Methods for Bayesian Filtering. *Statistics and computing*, 10(3):197–208, 2000.

[4] N. Ikeda and S. Wantanabe. *Stochastic Differential Equations and Diffusion Processes*. North-Holland, 1989.

[5] Z. Khan, T. Balch, and F. Dellaert. *An MCMC-Based Particle Filter for Tracking Multiple Interacting Targets*. Springer, Berlin / Heidelberg, Germany, 2004.

[6] Z. Khan, T. Balch, and F. Dellaert. MCMC Data Association and Sparse Factorization Updating for Real Time Multi-target Tracking with Merged and Multiple Measurements. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(12):1960–72, 2006.

[7] H. Kunita. *Stochastic Flows and Stochastic Differential Equations*. Cambridge University Press, 1990.

[8] D. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.

[9] S. Oh, S. Russell, and S. Sastry. Markov Chain Monte Carlo Data Association for General Multiple-Target Tracking Problems. *43rd IEEE Conference on Decision and Control*, 1:735–742, 2004.

[10] B. Ristic, S. Arulampalam, and N. Gordon. *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, 2004.

[11] 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 *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2001.

[12] D. Schulz, W. Burgard, D. Fox, and A.B. Cremers. People Tracking with Mobile Robots Using Sample-Based Joint Probabilistic Data Association Filters. *The International Journal of Robotics Research*, 22:99, 2003.

[13] M. Travers, T. Murphey, and L. Pao. Data Association for Systems with Measurement Origin Uncertainty. In *American Control Conference*, Seattle, WA, 2008.