2016 IEEE Intelligent Vehicles Symposium (IV)
Gothenburg, Sweden, June 19-22, 2016

Multi-Sensor Multi-Object Tracking of Vehicles Using High-Resolution Radars

Alexander Scheel[1], Christina Knill[2], Stephan Reuter[1], and Klaus Dietmayer[1]

Abstract— **Recent advances in automotive radar technology have led to increasing sensor resolution and hence a more detailed image of the environment with multiple measurements per object. This poses several challenges for tracking systems: new algorithms are necessary to fully exploit the additional information and algorithms need to resolve measurement-to-object association ambiguities in cluttered multi-object scenarios. Also, the information has to be fused if multi-sensor setups are used to obtain redundancy and increased fields of view. In this paper, a Labeled Multi-Bernoulli filter for tracking multiple vehicles using multiple high-resolution radars is presented. This finite-set-statistics-based filter tackles all three challenges in a fully probabilistic fashion and is the first Monte Carlo implementation of its kind. The filter performance is evaluated using radar data from an experimental vehicle.**

I. INTRODUCTION

Radar sensors are a crucial element in modern vehicle environment perception systems which is mostly due to their ability to supply precise range and range rate (Doppler) measurements, the robustness in adverse weather and lighting conditions, and the comparatively low cost. Yet, radar sensors have long provided limited information as they often yielded a single detection per object. In recent years, however, progress in automotive radar technology has led to increasing sensor resolution which enables sensors to resolve multiple detections per object. Thus, a more precise image of the environment is available and the additional information allows to deduce object dimensions or even contours.

However, this development poses two major challenges for tracking algorithms. First, so-called extended objects, i.e. objects with an extent that is not negligible in comparison to sensor resolution and which give rise to several detections, violate the point object assumption which is commonly made in classical tracking literature, e.g. for the standard Kalman filter. Under this assumption, objects are assumed to be small in relation to sensor resolution and to produce at most one measurement. Secondly, the increased amount of measurements aggravates the measurement-to-object association problem which arises if multiple closely-spaced objects or clutter measurements are present. Vehicle tracking constitutes such a scenario as there are oftentimes several vehicles and high-resolution radar sensors also provide a multitude of measurements from other objects in the environment. Also, high-resolution sensors are able to resolve spurious detections from spinning wheels which deteriorate tracking results if falsely used for estimating the vehicle speed. A third challenge emerges if multiple radar sensors are used to increase the field of view (FOV) or to achieve redundancy: the available information has to be fused to obtain a single and consistent model of the environment.

There are two approaches to solve the extended object problem. The first and probably most common method is to add a preprocessing routine which reduces the measurements to single meta-measurements by first clustering the detections which presumably originate from one object and then inferring properties such as the centroid position or velocity. An exemplary and recent approach has been presented in [1] and [2], where the velocity and yaw rate of an object is computed from high-resolution radar data in a single scan by extracting velocity profiles using one and two sensors, respectively. In general, preprocessing routines tend to fail if there is ambiguity in the data of a single scan and the object cannot be extracted clearly. The second approach is to design extended object filter algorithms that work on the raw measurements directly and are able to process multiple measurements per object. Prominent examples have been presented in [3], [4], and [5]. Yet, these methods are not tailored to vehicle tracking with radar sensors and their application is impeded by complexity, the lack of support for range rate measurements, or inept modelling assumptions. In contrast, a direct scattering approach as well as an extended object model based on typical reflection centers have been specifically designed for vehicle tracking with multiple radar measurements in [6] and [7], respectively.

Random Finite Set (RFS) filters which are based in finite set statistics (FISST) ([8], [9], and [10]) have recently become popular tools for solving the second challenge, the multi-object tracking problem. The key idea is to represent the multi-object state as an RFS, i.e. a set which contains all object state vectors but where not only the states but also the number of objects are unknown and random. The posterior density over this multi-object state is then estimated recursively in a similar fashion to the standard Bayes filter. Since both the object states and the measurements are modeled as RFS, association ambiguities and clutter can be treated probabilistically and are filtered over time. Different extended object RFS filters have been proposed, including Probability Hypothesis Density (PHD) [11] and Cardinalized Probability Hypothesis Density (CPHD) [12] filters which propagate the first moment of the multi-object density. A recent development is the Generalized Labeled Multi-Bernoulli (GLMB) filter for extended objects [13] which

[1]A. Scheel, S. Reuter, and K. Dietmayer are with the Institute for Measurement, Control, and Microtechnology at Ulm University, 89081 Ulm, Germany {alexander.scheel, stephan.reuter, klaus.dietmayer}@uni-ulm.de

[2]C. Knill is with the Institue of Microwave Engineering at Ulm University, 89081 Ulm, Germany christina.knill@uni-ulm.de

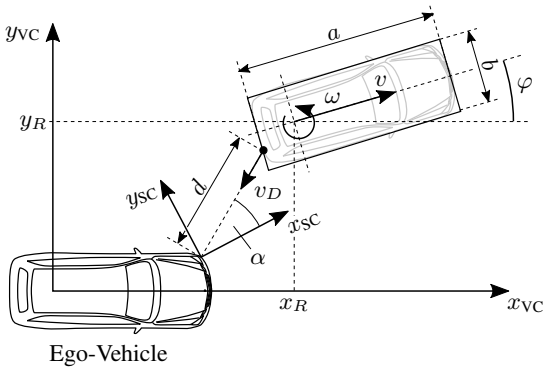

Fig. 1: Schematic overview of the state vector and measurement quantities in the vehicle (VC) and sensor (SC) coordinate systems.

yields improved tracking performance in comparison to the former filters and its approximate as well as computationally faster Labeled Multi-Bernoulli (LMB) variant [13].

In this paper, a vehicle tracking approach is presented that tackles all mentioned challenges. An extended object LMB filter core is used in conjunction with the extended object model from [6] which was selected because of its simplicity and ability to handle complicated and ambiguous situations. With this setup, the entire filtering process from the raw target-level measurements to the tracks is dealt with probabilistically. Furthermore, it is demonstrated that the method can be used to fuse data from two or more sensors.

In the following, the tracking problem is first formulated in Section II. Then, the measurement model and the multi-object filter are presented in Section III and Section IV, respectively. Section V demonstrates the accuracy and performance of the proposed approach on real sensor data and Section VI concludes the paper.

II. PROBLEM FORMULATION

The goal is to estimate the state and dimensions of all vehicles in the surrounding of the ego-vehicle using available radar measurements. For this purpose, each vehicle is described by a composed state vector $x_k = [\xi_k^T, \zeta_k^T]^T \in \mathbb{X}$, where $\mathbb{X}$ denotes the state space and k the time step index. The first portion $\xi_k = [x_{R,k}, y_{R,k}, \varphi_k, v_k, \omega_k]^T$ denotes the kinematic state and consists of the Cartesian positions of the center of the rear axle x_R and y_R, the orientation φ, the speed v, and the yaw rate ω. These states are due to the constant turn rate and velocity model which is used for describing the vehicle motion. Additionally, the extent portion $\zeta_k = [a_k, b_k]^T$ describes the width a and the length b of the vehicle. All quantities are estimated in the vehicle coordinate system which is located at the ego-vehicle's center of the rear axle with the x-axis pointing to the front and the y-axis to the left. See Fig. 1 for a schematic illustration. Since the location of the rear axle is very similar in most passenger cars, a fixed length ratio is used to position the center rear axle in the rectangle.

Since there may be multiple vehicles present, each object is additionally assigned a label ℓ from the discrete label space $\mathbb{L}$ which uniquely identifies it. The labels are augmented to the state vector which yields the labeled single-object state $\mathbf{x}_k = [x_k^T, \ell]^T$. The multi-object state which describes the entire scene, i.e. the state of each vehicle and the number of present vehicles, is given by the RFS $\mathbf{X} \subset \mathbb{X} \times \mathbb{L}$ which constitutes a finite set of labeled state vectors with random cardinality. The notation $|\mathbf{X}|$ is used to determine the cardinality for a specific realization.

Note that this paper follows the notation from [14] which uses lower case letters (x) for single-object quantities, upper case letters (X) for finite sets, and bold letters ($\mathbf{x}$, $\mathbf{X}$) for labeled states. The function $\mathcal{L}(\mathbf{X}) = \{\ell \mid [x, \ell]^T \in \mathbf{X}\}$ is used to project the labeled multi-object state to the label space and to retrieve the contained labels.

The radar data is provided by several sensors which deliver radar detections on target level. That is, the i-th measurement $z_k^i = [d_k^i, \alpha_k^i, v_{D,k}^i]^T$ provides its radial distance d^i, its azimuth angle α^i, and its Doppler velocity v_D^i in sensor coordinates, see Fig. 1. Since there are multiple measurements in each cycle and the number of measurements m varies, all measurements of one scan are described as the RFS $Z_k = \{z_k^1, ..., z_k^m\}$ which is a subset of the measurement space $\mathbb{Z}$.

Using this formulation, the multi-object Bayes filter [10], which is a rigorous extension of the classical Bayes filter to RFSs, can be used to provide an estimate of the posterior distribution of the labeled multi-object state given the available measurements up to the current time step $\pi_{k|k}(\mathbf{X}_k|Z_{1:k})$. Hence, it provides a probabilistic framework to obtain an estimate of the number of vehicles and their respective state vectors directly from the raw measurements.

The multi-object Bayes filter iteratively propagates the the multi-object distribution using a prediction and update step. During prediction, the posterior multi-object density from time step $k-1$ evolves according to the Chapman-Kolmogorov equation

$$\pi_{k|k-1}(\mathbf{X}_k|Z_{1:k-1}) = \int f_{k|k-1}(\mathbf{X}_k|\mathbf{X}_{k-1})\pi_{k-1|k-1}(\mathbf{X}_{k-1}|Z_{1:k-1})\delta\mathbf{X}_{k-1} \tag{1}$$

and yields the prior multi-object density $\pi_{k|k-1}(\mathbf{X}_k|Z_{1:k-1})$ at time k. The dynamics are captured in the multi-object transition density $f_{k|k-1}(\mathbf{X}_k|\mathbf{X}_{k-1})$. Note that the integral is a set integral as defined in [10].

In the update step, the new measurement set Z_k is used to compute the posterior multi-object density

$$\pi_{k|k}(\mathbf{X}_k|Z_{1:k}) = \frac{g_k(Z_k|\mathbf{X}_k)\pi_{k|k-1}(\mathbf{X}_k|Z_{1:k-1})}{\int g_k(Z_k|\mathbf{X}_k)\pi_{k|k-1}(\mathbf{X}_k|Z_{1:k-1})\delta\mathbf{X}_k} \tag{2}$$

using Bayes' theorem. Here, $g_k(Z_k|\mathbf{X}_k)$ is the multi-object likelihood function which captures the measurement process. As the following explanations focus on a single filter iteration, the time index k is dropped in the remainder of the paper to improve readability. Prior components are distinguished with a $+$ as subscript.

III. MEASUREMENT MODEL

Before introducing the tracking filter equations, the centerpiece of the update step, the employed multi-object measurement model, is presented. This is done in two steps: First, the general multi-object likelihood is discussed. Afterwards, the single-object likelihood function which is required to compute the multi-object likelihood and which includes the radar-specific portion is outlined.

A. Multi-Object Measurement Likelihood

Multi-object likelihood functions describe the relationship between a multi-object state and the expected measurements. Many RFS filters for extended targets use an observation model which is based on the following assumptions:

1) Each object in the multi-object state is detected by the sensor with the detection probability $p_D(x,\ell)$ or misdetected with probability $q_D(x,\ell) = 1 - p_D(x,\ell)$.
2) If an object is detected, it generates an independent set of measurements W which is distributed according to the single-object likelihood function $\tilde{g}(W|\mathbf{x})$.
3) Additionally, the sensor produces clutter measurements which are independent of the object states and follow the Poisson RFS g_C [10] with intensity function $\kappa(z) = \lambda_C p_C(z)$. That is, the number of clutter measurements is Poisson distributed with parameter λ_C and the elements follow the clutter distribution $p_C(z)$. In this work, a uniform distribution over the entire measurement space is used for the clutter density.

In [13], it was demonstrated that the multi-object model under these assumptions can be written as

$$g(Z|\mathbf{X}) = g_C(Z) \sum_{i=1}^{|\mathbf{X}|+1} \sum_{\substack{\mathcal{U}(Z)\in\mathcal{P}_i(Z)\\ \theta\in\Theta(\mathcal{U}(Z))}} \left[\psi_{\mathcal{U}(Z)}(\cdot;\theta)\right]^{\mathbf{X}} \tag{3}$$

with

$$\psi_{\mathcal{U}(Z)}(x,\ell;\theta) = \begin{cases} \frac{p_D(x,\ell)\tilde{g}(\mathcal{U}_{\theta(\ell)}(Z)|x,\ell)}{[\kappa]^{\mathcal{U}_{\theta(\ell)}(Z)}}, & \theta(\ell) > 0 \\ q_D(x,\ell), & \theta(\ell) = 0 \end{cases} \tag{4}$$

and

$$g_C(Z) = e^{\lambda_C}[\kappa(\cdot)]^Z. \tag{5}$$

In the above equations, the multi-object exponential notation

$$h^X \triangleq \prod_{x\in X} h(x) \tag{6}$$

is used for brevity. It denotes products of real-valued functions $h(x)$ for all elements in a set X. For an empty set $\emptyset$, $h^\emptyset = 1$. Also, $\mathcal{P}_i(Z)$ denotes the set of all partitions which separate the measurement set in exactly i mutually exclusive clusters and $\mathcal{U}(Z)$ is one particular partition. For a given partition, association mappings $\theta : \mathcal{L}(\mathbf{X}) \to \{0,1,\ldots,|\mathcal{U}(Z)|\}$ assign the objects to the measurement clusters in the partition such that $\theta(\ell) = \theta(\ell') > 0$ implies $\ell = \ell'$ and an association to the index 0 stands for a misdetection. The space of all association mappings is denoted by $\Theta(\mathcal{U}(Z))$ and $\mathcal{U}_{\theta(\ell)}(Z)$ identifies the cluster which has been assigned to the object with label ℓ.

Intuitively, (3) computes the probability of obtaining the measurement set Z by testing all possibilities of how the measurements could have been created and summing the respective probabilities. The increment of one in the superscript of the first sum makes sure that there is always one additional cluster which contains unassigned measurements, i.e. clutter. The factor $g_C(Z)$ denotes the probability that all detections are clutter measurements and individual elements are canceled by the term in the denominator of (4) if the corresponding measurement has been assigned to an object.

B. Single-Object Measurement Likelihood

For the single-object likelihood in (4), the direct scattering model from [6] with adaptions to fit to the presented multi-object measurement model is used. It models vehicles as rectangles and provides a simple but effective approach to represent the relationship between vehicle state and the received measurements. The model is able to use the entire Doppler information even in complex constellations such as cross traffic or highly dynamic situations with high yaw rates and works in situations with strong ambiguities. In the following, the adapted model is presented briefly. For a detailed discussion of the original measurement model and evaluation of accuracy in single-object single-sensor problems, refer to [6].

The single-object likelihood for a set of measurements W is modeled as the Poisson RFS

$$\tilde{g}(W|x,\ell) = e^{-\lambda_T}\left[\lambda_T p_z(\cdot|x)\right]^W \tag{7}$$

where the multi-object exponential from (6) was again used for brevity. This model is based on two assumptions: First, the number of expected measurements is Poisson distributed with expected value λ_T and secondly, all measurements are assumed to be conditionally independent and distributed according to the measurement likelihood for a single measurement $p_z(z|x)$. In practice, the number of received measurements is difficult to model and does not necessarily follow a Poisson distribution as it depends on the the sensor-to-object constellation. By using an identical distribution for both the measurement and clutter rate and setting $\lambda_T = \lambda_C$, however, it is achieved, that the portions for the number of target and clutter measurements in (4) cancel and no preference is given to a specific number of measurements.

The single-object measurement likelihood $p_z(z|x)$ describes the likelihood of obtaining the three measurement quantities for a given vehicle state and can be factored to

$$\begin{aligned} p_z(z|x) &= p_z(d,\alpha,v_D|x) \\ &= p_{v_D}(v_D|d,\alpha,x)p_d(d|\alpha,x)p_\alpha(\alpha|x). \end{aligned} \tag{8}$$

A trapezoidal probability density is used to model the likelihood for the azimuth angle $p_\alpha(\alpha|x)$. That is, the azimuth angle portion of the sensor FOV which is covered by the rectangular vehicle model is determined and every angle in this interval receives an identical density value. The density values linearly decrease to zero to both sides of the vehicle which introduces transition region. Measurements outside these regions receive a likelihood of zero.

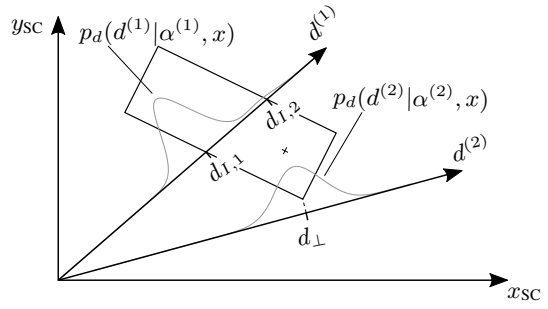

Fig. 2: Range measurement likelihood model with qualitative range densities: Ray 1 targets the rectangle, ray 2 passes it.

For a given state and azimuth angle, the expected range measurement can be computed easily by intersecting an imaginary ray with the rectangular vehicle model as illustrated in Fig. 2. Rays that target the rectangle directly result in the two intersection points $d_{I,1}$ and $d_{I,2}$ where the ray enters and leaves the rectangle, respectively. The range density is then modeled as a mixture of two Gaussian distributions which are centered at the intersection points. The weights of both Gaussians c_i sum to one and are set to $c_1 = 0.95$ and $c_2 = 0.05$. This models the fact that measurements are typically received from the surface facing the sensor but ensures that rarely occurring measurements that lie in the rectangle or on the averted side due to reflections from the underbody can be handled. The standard deviation of both Gaussians is given by σ_r and models the range accuracy of the sensor as well as deviations of the actual vehicle shape from the rectangular model. If a ray does not target the rectangle but passes it, only a single Gaussian, centered at the point on the ray that is closest to the rectangle, is used. The standard deviation of this Gaussian increases with the distance between ray and rectangle.

The expected Doppler measurement for a given vehicle state is determined by

$$\hat{v}_D(x) = v \cdot \cos(\alpha - \varphi) + \omega \left(y_R \cdot \cos(\alpha) - x_R \cdot \sin(\theta) \right). \quad (9)$$

As the Doppler velocity does not depend on the measured range, $p_{v_D}(v_D|d,\alpha,x) = p_{v_D}(v_D|\alpha,x)$. It is again modeled as a Gaussian distribution with the expected Doppler velocity $\hat{v}_D$ as mean and standard deviation σ_{v_D} which depends on the sensor specifications.

IV. MULTI-OBJECT TRACKING

The Labeled Multi-Bernoulli filter for extended objects, which was introduced in [13], is used to estimate the posterior multi-object density with the presented measurement model. It provides a solution to the general multi-object Bayes filter from (1) and (2) which is based on specific types of multi-object distributions, the class of LMB distributions.

A. Labeled Multi-Bernoulli Distributions

Labeled RFS were introduced in [14] and are one form of modeling the multi-object density $\pi(X)$. As presented in Section II, a label is attached to all components of the finite set thus allowing to keep track of the object identity and its trajectory over several time steps. To ensure that track labels are distinct and refer to one object only, the distinct label indicator $\Delta(\mathbf{X}) = \delta(|\mathcal{L}(\mathbf{X})| - |X|)$ is introduced where $\delta(\cdot)$ denotes the Kronecker-delta function. The following paragraphs provide an introduction to the two labeled multi-object distributions that are used in this paper, the LMB RFS and the GLMB RFS.

1) Labeled Multi-Bernoulli RFS: An LMB RFS consists of multiple independent and labeled components. The labels ℓ of all components constitute the discrete label space $\mathbb{L}$. Every component is described by a distribution of its state $p^{(\ell)}(x) = p(x,\ell)$ and an existence probability $r^{(\ell)}$. Hence, the multi-object density of an LMB RFS is given by

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) w(\mathcal{L}(\mathbf{X})) p^{\mathbf{X}} \quad (10)$$

where

$$w(L) = \prod_{i \in \mathbb{L}} \left(1 - r^{(i)}\right) \prod_{\ell \in L} \frac{1_{\mathbb{L}}(\ell) r^{(\ell)}}{1 - r^{(\ell)}} \quad (11)$$

defines the probability that only the objects with labels $L = \mathcal{L}(\mathbf{X})$ exist. Here, $1_{\mathbb{L}}(\ell)$ is the inclusion function which is equal to 1 if and only if $\ell \in \mathbb{L}$. Since LMB RFSs assume independent tracks, they are limited in representing complex situations where a multimodal cardinality distribution is required due to strong interdependencies.

2) Generalized Labeled Multi-Bernoulli RFS: The GLMB distribution overcomes the limitations of LMB RFSs by allowing for multiple possible realizations from an index set $\mathbb{C}$ for a set of track labels. It is given by

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{c \in \mathbb{C}} w^{(c)}(\mathcal{L}(\mathbf{X})) \left[p^{(c)}\right]^{\mathbf{X}} \quad (12)$$

where the weights have to satisfy

$$\sum_{L \subseteq \mathbb{L}} \sum_{c \in \mathbb{C}} w^{(c)}(L) = 1. \quad (13)$$

Note that in contrast to (10), there is no explicit definition of the weights $w^{(c)}(\mathcal{L}(\mathbf{X}))$ and various definitions which satisfy (13) are possible. Hence, the GLMB RFS allows multimodal cardinality distributions and is able to model interdependencies among objects.

B. Labeled Multi-Bernoulli Filter for Extended Objects

The LMB filter for extended objects proceeds in three steps: prediction, update, and approximation.

1) Prediction: The posterior multi-object density in LMB RFS form is predicted using the standard multi-object transition density [10]. This transition assumes that each object survives to the next time step with probability $p_S(x,\ell)$ or disappears with probability $q_S(x,\ell) = 1 - p_S(x,\ell)$. If an object survives, its states evolve according to the single-object transition density $f_+(x|x',\ell)$. Additionally, new components may be created and are combined in a birth density of LMB form with label space $\mathbb{B}$ ($\mathbb{L} \cap \mathbb{B} = \emptyset$), component weights $r_B^{(\ell)}$, and densities $p_B^{(\ell)}(x)$. This density is then added to the predicted density of existing objects. As proven in [15], the prior multi-object density which results from (1) under these

assumptions is again an LMB RFS with augmented label space $\mathbb{L}_+ = \mathbb{L} \cup \mathbb{B}$ and parameter set

$$\boldsymbol{\pi}_+ = \left\{ \left(r_{+,S}^{(\ell)}, p_{+,S}^{(\ell)} \right) \right\}_{\ell \in \mathbb{L}} \cup \left\{ \left(r_B^{(\ell)}, p_B^{(\ell)} \right) \right\}_{\ell \in \mathbb{B}} \tag{14}$$

with

$$r_{+,S}^{(\ell)} = \eta_S(\ell) r^{(\ell)} \tag{15}$$

$$p_{+,S}(x, \ell) = \frac{\int p_S(x, \ell) f_+(x|x', \ell) p(x', \ell) dx'}{\eta_S(\ell)} \tag{16}$$

$$\eta_S(\ell) = \iint p_S(x, \ell) f_+(x|x', \ell) p(x', \ell) dx' dx. \tag{17}$$

2) Update: GLMB RFSs are a conjugate prior to the multi-object likelihood from (3) [13]. Since an LMB distribution is a special case of a GLMB RFS, the update step of the multi object Bayes filter from (2) yields

$$\begin{aligned} \boldsymbol{\pi}(\mathbf{X}|Z) =& \Delta(\mathbf{X}) \sum_{i=1}^{|\mathbf{X}|+1} \sum_{\substack{\mathcal{U}(Z) \in \mathcal{P}_i(Z) \\ \theta \in \Theta(\mathcal{U}(Z))}} w_{\mathcal{U}(Z)}^{(\theta)}(\mathcal{L}(\mathbf{X})) \\ & \times \left[p^{(\theta)}(\cdot|\mathcal{U}(Z)) \right]^{\mathbf{X}} \end{aligned} \tag{18}$$

with

$$w_{\mathcal{U}(Z)}^{(\theta)}(L) = \frac{w_+(L) \left[\eta_{\mathcal{U}(Z)}^{(\theta)}(\cdot) \right]^L}{\sum_{J \subseteq \mathbb{L}} \sum_{i=1}^{|J|+1} \sum_{\substack{\mathcal{U}(Z) \in \mathcal{P}_i(Z) \\ \theta \in \Theta(\mathcal{U}(Z))}} w_+(J) \left[\eta_{\mathcal{U}(Z)}^{(\theta)}(\cdot) \right]^J}, \tag{19}$$

$$p^{(\theta)}(x, \ell|\mathcal{U}(Z)) = \frac{p_+(x, \ell) \psi_{\mathcal{U}(Z)}(x, \ell; \theta)}{\eta_{\mathcal{U}(Z)}^{(\theta)}(\ell)}, \tag{20}$$

$$\eta_{\mathcal{U}(Z)}^{(\theta)}(\ell) = \int p_+(x, \ell) \psi_{\mathcal{U}(Z)}(x, \ell; \theta) dx, \tag{21}$$

and the prior weight $w_+(L)$ which is obtained from the prior multi-object density through (11). The posterior distribution captures all different hypotheses of measurement-to-object associations and is now in the form of a GLMB, i.e. the objects are not independent anymore. This is due to the fact that the objects share the same measurement set and measurement clusters can only be assigned to one object.

3) Approximation: In comparison to the prior LMB distribution, the number of components in the posterior distribution has drastically increased. Further propagating the GLMB would result in a more complex filtering procedure with increased computation time, which is especially detrimental if a Sequential Monte Carlo (SMC) implementation is pursued. The LMB filter hence approximates the posterior GLMB density by an LMB density which matches its first moment. The parameters for the approximate posterior LMB density are [15]

$$r^{(\ell)} = \sum_{L \subseteq \mathbb{L}_+} \sum_{i=1}^{|L|+1} \sum_{\substack{\mathcal{U}(Z) \in \mathcal{P}_i(Z) \\ \theta \in \Theta(\mathcal{U}(Z))}} w_{\mathcal{U}(Z)}^{(\theta)}(L) 1_L(\ell), \tag{22}$$

$$\begin{aligned} p^{(\ell)}(x) =& \frac{1}{r^{(\ell)}} \sum_{L \subseteq \mathbb{L}_+} \sum_{i=1}^{|L|+1} \sum_{\substack{\mathcal{U}(Z) \in \mathcal{P}_i(Z) \\ \theta \in \Theta(\mathcal{U}(Z))}} w_{\mathcal{U}(Z)}^{(\theta)}(L) 1_L(\ell) \\ & \times p^{(\theta)}(x, \ell). \end{aligned} \tag{23}$$

During this process, the state densities of the objects remain identical and the information loss mostly concerns the cardinality estimate as well as interdependencies among objects.

C. Sequential Monte Carlo Implementation

The single-object measurement model presented in Section III-B is highly nonlinear due to the the rectangular shape and the intersection process between rays and rectangle. Therefore, particle distributions are used to represent the state distributions $p^{(\ell)}(x)$. In particular, a Rao-Blackwellized particle filter (RBPF) scheme [16] is used in the underlying estimation process for each object which consists of the single-object prediction (16) and single-object update (20). This procedure is identical to the single-object case from [6]. The state density

$$\begin{aligned} p^{(\ell)}(x|Z) =& p^{(\ell)}(\xi, \zeta|Z) = p^{(\ell)}(\xi|Z) p^{(\ell)}(\zeta|\xi, Z) \\ =& \sum_{i=1}^{n} w_p^{(i,\ell)} \delta(\xi^{(i)} - \xi) p^{(\ell)}(\zeta^{(i)}|\xi^{(i)}, Z) \end{aligned} \tag{24}$$

is split into a portion for the kinematic state which is represented by n particles and a geometric portion which is modeled as a discrete distribution. Due to the conditioning on ξ of the i-th particle, each kinematic state particle carries its own discrete distribution for the object extent. The reader is referred to [16] for the update equations of a RBPF filter with discrete distributions. After the particles have been updated, a resampling step is included to avoid sample impoverishment.

D. Fusion Approach

For the fusion of data from multiple sensors, a centralized fusion approach is chosen since it is makes the best use of all available information and yields optimal results [17]. The fusion system consists of a central processing unit and one instance of the measurement model for each sensor which is parameterized with corresponding sensor-specific values such as the sensor position. As new measurements become available, the multi-object density is predicted to the current time and the prior density is then updated using the respective measurement model instance.

E. Implementation Details

The following paragraphs address implementation issues that have to be considered to apply the filter in practice.

1) Variation of the extent: The main goal of estimating the vehicle dimensions is to adapt the detailed measurement model to the required size and to get a coarse estimate with decimeter accuracy. This is achieved by representing the extent state as discrete distribution which allows a simple and approximate estimation procedure. During prediction, pseudo noise is added to the current estimate. That is, additional size hypotheses which vary both the width and length with

fixed step size in positive and negative direction are added. To avoid a drastic increase in size hypotheses over multiple time steps, the extent estimate is approximated by its mean value after each update step. Due to the strong correlation of object extent and the position of the center rear axle, changes in position may be wrongly explained by a change in dimension. To avoid this issue, the corner of the rectangular model that is closest to the sensor is kept fixed and size changes are only allowed relative to this anchor point. The variation of the dimensions is restricted to only include reasonable vehicle hypotheses. The bounds are $a_{\text{min}} = 2.5$ m and $a_{\text{max}} = 7$ m for the length as well as $b_{\text{min}} = 1.4$ m and $b_{\text{max}} = 2.5$ m for the width. Also, the length-to-width ratio is constrained to the interval $[1.7, 3.5]$.

2) Partitioning and association: Theoretically, the measurement model requires the likelihood to be evaluated for all possible partitions and associations, see (3) and (18). As this is computationally intractable even for a moderate amount of measurements, only relevant and viable clusters as well as associations are evaluated in practice. Hence, the computation time is reduced drastically while the possibility to maintain several plausible and ambiguous association hypotheses is still given. In this work, measurement clusters are obtained by repeatedly applying distance-based clustering algorithms with varying thresholds between 0.4 and 5 m. Moreover, additional cluster hypotheses are added by further splitting measurements in a cluster if the range rate measurements indicate multiple objects that move differently. This approach allows to obtain clusters which exclude spurious measurements, e.g. from rotating wheels.

3) Initialization and pruning of objects: The birth density in (14) introduces new objects to the filter procedure. A new particle density is initialized as soon as a cluster of measurements that has not been successfully used to update an existing track is found and which also includes at least one measurement with a Doppler velocity of more than 2.5 m/s. Hence, the filter only initializes moving objects and the false alarm rate due to measurements from static objects is kept low. Once an object is initialized, it may stop and will still be tracked. New objects are initialized with an existence probability of $r_B^{(\ell)} = 0.1$. Objects are pruned if their probability of existence drops below 1%.

4) Ego-motion compensation: Since the ego-vehicle moves and the other vehicles are tracked with respect to the vehicle coordinate system, ego-motion compensation steps are included to correct the objects' states before prediction as well as obtained Doppler measurements.

V. EVALUATION

The proposed tracking algorithm was implemented in MATLAB and tested on real world data obtained from two short range radar sensors in a test vehicle. The two sensors are mounted in the right and left corners of the front bumper, have a range of approximately 45 m, and an opening angle of $\pm 85°$. The FOVs of both sensors overlap in front of the vehicle, whereas the side areas are only covered by one sensor. Additionally, the test vehicle is equipped with a DGPS and inertial measurement unit (IMU) which is able provide accurate measurements of the ego-vehicle position, orientation and motion. Two scenarios have been chosen to evaluate the filter performance: a single-object scenario for evaluating the estimation accuracy when using two sensors and a scenario with two vehicles that enter and leave the FOV to demonstrate the multi-object capability. All results were obtained by using 300 particles per object.

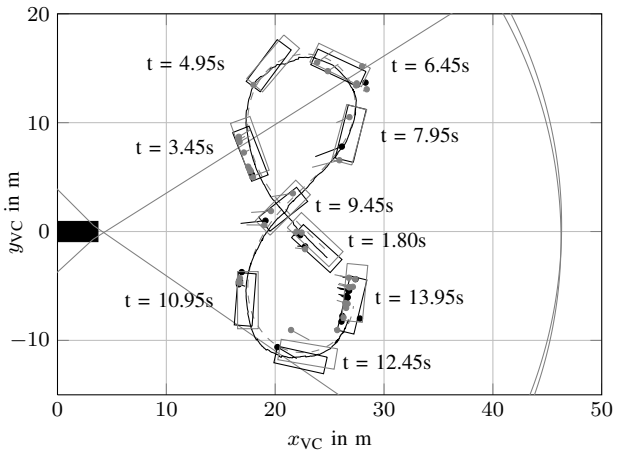

Fig. 3: Extract of single-object scenario: estimated trajectory (black) and vehicle (black rectangles), ground truth trajectory (dashed gray) and vehicle (gray rectangles), measurements of the left (gray) and right (black) sensor in the vicinity of the vehicle, FOVs (solid gray).

A. Single-Object Accuracy with two Sensors

In the single-object scenario, the ego-vehicle is at standstill and a second vehicle is driving multiple horizontal eights in front of the ego-vehicle. This scenario is highly dynamic as the vehicle makes many turns, and changes its orientation, speed, and yaw rate constantly. This requires an accurate measurement model for correctly processing the Doppler measurements. The target vehicle is also equipped with a DGPS / IMU system that allows to compute ground truth values for the kinematic states. As observable in Fig. 3, which depicts an extract of estimation results to visualize the scenario, the filter is able to track the vehicle well, even if the vehicle leaves the FOV of the right sensor intermittently. This is mostly due to the extended object model which precisely represents the relationship between object state and measurements. Also, the filter is able to process a strongly varying amount of measurements which originate from all four sides. To eliminate non-deterministic effects of the SMC implementation, the filter was run 100 times. Average estimation results along with the corresponding errors are shown for each estimated state in Fig. 4. Note that the position error is strongly correlated with the estimated object dimensions. As the vehicle extent is not easily observable from the measurements, the filter tends to underestimate width and length which in turn increases the position error.

B. Multi-Object Scenario

To demonstrate the algorithm in a more complex setting, a parallel traffic scenario with two target vehicles which

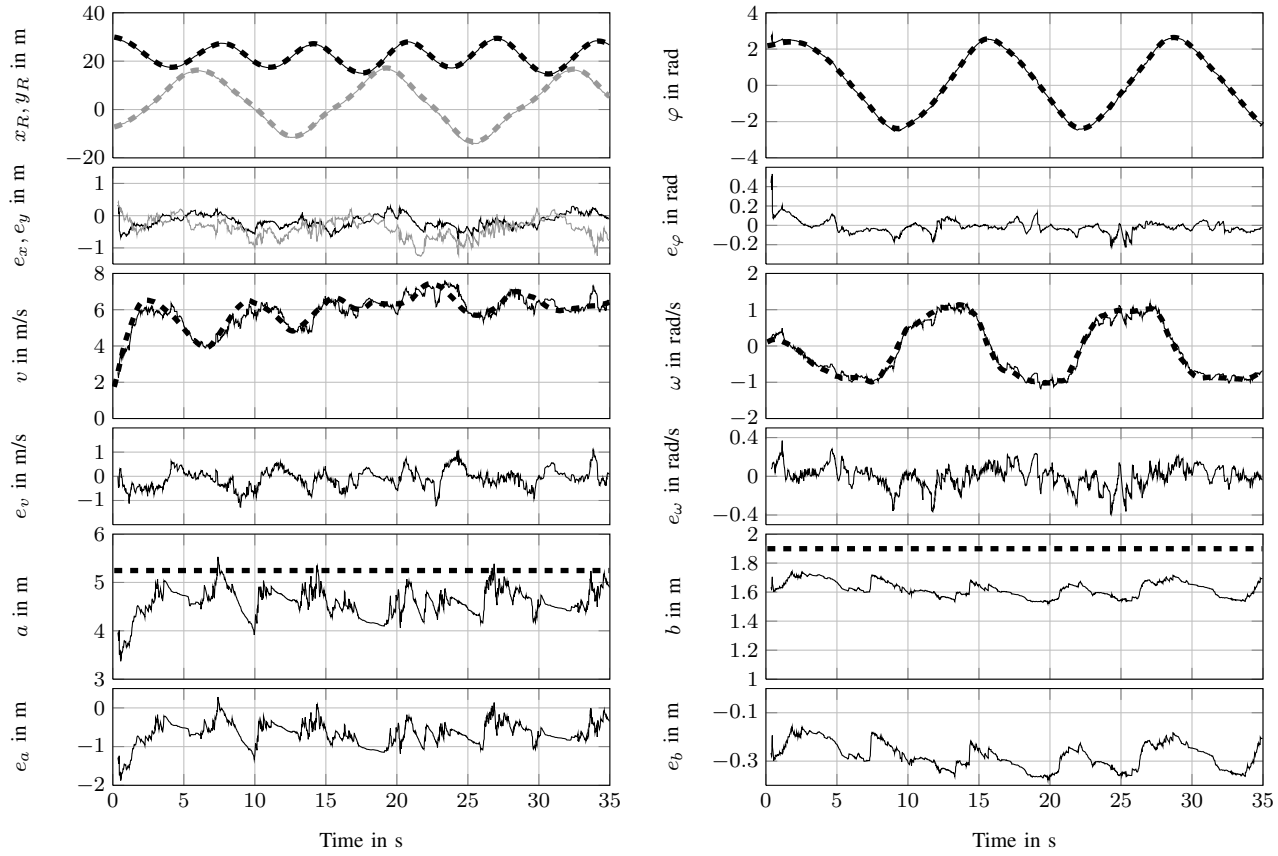

Fig. 4: Single-object scenario results: Mean value of the estimated states (solid) over 100 simulation runs, ground truth values (dashed), and respective errors e. In the position plot, the x-position is depicted in black and the y-position in gray.

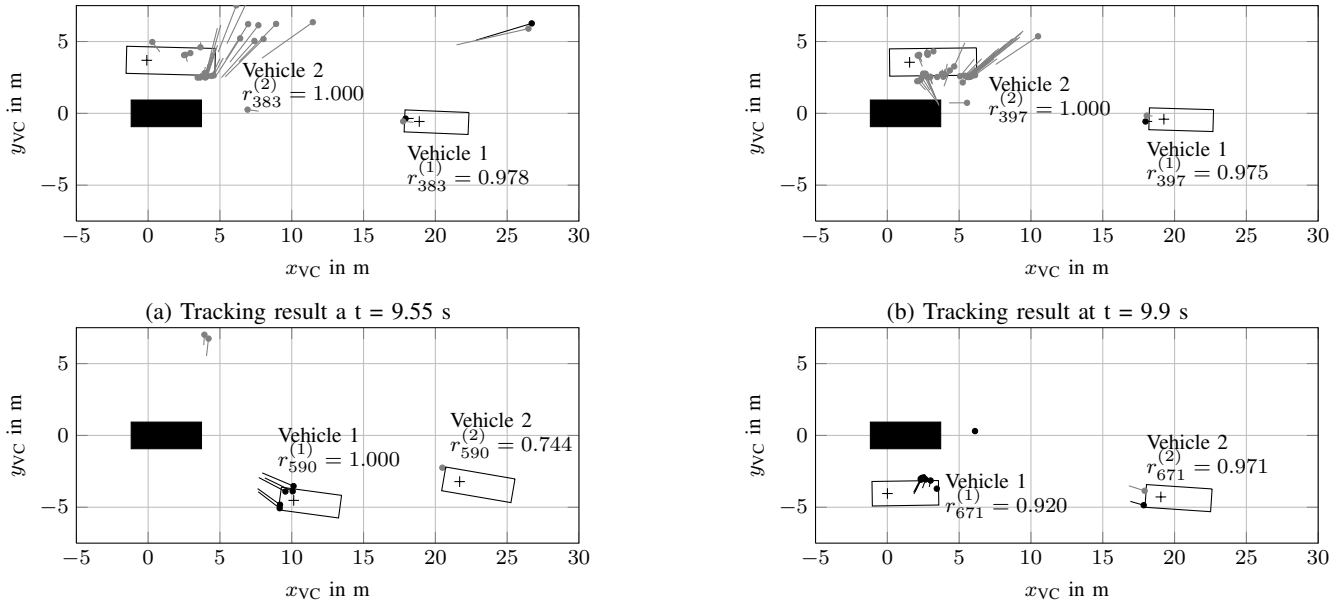

(a) Tracking result a t = 9.55 s

(b) Tracking result at t = 9.9 s

(c) Tracking result at t = 14.72 s

(d) Tracking result at t = 16.76 s

Fig. 5: Exemplary time steps of the multi-object scenario: ego-vehicle (filled rectangle) and estimated objects with measurements form the left (gray) and right (black) sensor, the magnitude of Doppler measurements is indicated with a line

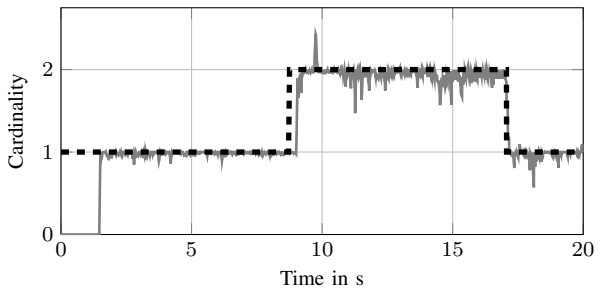

Fig. 6: Cardinality estimate (solid gray) for the multi-object scenario compared to the true value (dashed black)

comprises appearance and disappearance of objects was recorded. Particular time steps which illustrate the scenario and the estimation results are shown in Fig. 5. At the beginning, the ego-vehicle follows the first target vehicle. Then the second target vehicle enters the FOV to the left (see Figs. 5a and 5b) and overtakes the ego-vehicle as well as the first target vehicle (Fig. 5c). The first target vehicle leaves the FOV as the ego-vehicle also overtakes it (Fig. 5d). Clutter measurements are a major difficulty in this scenario which is observable in Fig. 5a, where target vehicle two is very close to clutter measurements from static objects next to the road, and Fig. 5b, where the clutter is caused by spinning wheels. In both cases, the association of measurements to the vehicle is not trivial and a wrong association may cause the filter to diverge. By probabilistically testing multiple different association hypotheses, the filter is able to find the correct association and to maintain the tracks.

Apart from the object states, estimating the number of objects in the FOV, i.e. the cardinality, is the second important aspect in multi-object scenarios. Figure 6 depicts the cardinality estimate for this scenario which is given by the sum of the existence probabilities for all elements in the posterior LMB distribution. The cardinality is underestimated if few or no measurements are available to support the vehicle hypotheses (e.g. at 14.5 s). It is overestimated if the filter initializes false tracks due to clutter measurements with significant Doppler velocity. In this scenario, the cardinality estimate shows only small and short deviations from the true value, which indicates that the filter successfully tracks both objects and corrects false alarms quickly. The deviation at the beginning exists because the filter is only allowed to initialize the track as the first vehicle begins to move.

VI. CONCLUSION

In this paper, a multi-object tracking approach for tracking vehicles in high-resolution radar data was presented. This first SMC implementation of the LMB filter for extended objects avoids preprocessing steps, fuses data from multiple sensors and uses the available information in a fully probabilistic fashion. This includes the entire processing routine from raw sensor measurements to multiple association and clustering hypotheses, fusion, track initialization as well as pruning, and the final estimates. Experiments on radar data demonstrated the filter's ability to track vehicles in complex driving maneuvers and to handle clutter measurements form static objects or spinning wheels. To reach the goal of a fully probabilistic and robust environment perception system, important future steps include porting the filter to real-time capable languages and to fuse the radar data with other complementary sensor measurements.

ACKNOWLEDGMENT

This work was supported by the German Research Foundation (DFG) within the Transregional Collaborative Research Center SFB/TRR 62 Companion-Technology for Cognitive Technical Systems.

REFERENCES

[1] D. Kellner, M. Barjenbruch, K. Dietmayer, and J. Klappstein, "Instantaneous Lateral Velocity Estimation of a Vehicle using Doppler Radar," in *Proceedings of the 16th International Conference on Information Fusion*, 2013, pp. 877–884.

[2] D. Kellner, M. Barjenbruch, J. Klappstein, J. Dickmann, and K. Dietmayer, "Instantaneous Full-Motion Estimation of Arbitrary Objects using Doppler Radar," in *Proveedings of the 2014 IEEE Intelligent Vehicles Symposium*, 2014, pp. 324–329.

[3] K. Gilholm and D. Salmond, "Spatial distribution model for tracking extended objects," *IEE Proceedings - Radar, Sonar and Navigation*, vol. 152, no. 5, pp. 364–371, 2005.

[4] M. Baum and U. D. Hanebeck, "Random Hypersurface Models for Extended Object Tracking," in *Proceedings of the IEEE International Symposium on Signal Processing and Information Technology*, 2009, pp. 178–183.

[5] J. W. Koch, "Bayesian Approach to Extended Object and Cluster Tracking using Random Matrices," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 44, no. 3, 2008.

[6] C. Knill and A. Scheel, "A Direct Scattering Model for Tracking Vehicles with High-Resolution Radars," in *Proceedings of the 2016 IEEE Intellignet Vehicles Symposium*, 2016.

[7] L. Hammarstrand, L. Svensson, F. Sandblom, and J. Sörstedt, "Extended Object Tracking using a Radar Resolution Model," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 3, pp. 2371–2386, 2012.

[8] Mahler, Ronald P. S, ""Statistics 101" for Multisensor, Multitarget Data Fusion," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 19, no. 1, pp. 53–64, 2004.

[9] Mahler, Ronald P. S, ""Statistics 102" for Multisource-Multitarget Detection and Tracking," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 376–389, 2013.

[10] Mahler, Ronald P. S, *Statistical multisource-multitarget information fusion*, ser. Artech House information warfare library. Boston: Artech House, 2007.

[11] Mahler, Ronald P. S, "PHD filters for nonstandard targets, I: Extended targets," in *Proceedings of the 12th International Conference on Information Fusion*, 2009, pp. 915–921.

[12] C. Lundquist, K. Granström, and O. Orguner, "An Extended Target CPHD Filter and a Gamma Gaussian Inverse Wishart Implementation," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 3, pp. 472–483, 2013.

[13] M. Beard, S. Reuter, K. Granström, B.-T. Vo, B.-N. Vo, and A. Scheel, "Multiple Extended Target Tracking with Labelled Random Finite Sets," *IEEE Transactions on Signal Processing*, vol. 64, no. 7, pp. 1638–1653, 2016.

[14] B.-T. Vo and B.-N. Vo, "Labeled Random Finite Sets and Multi-Object Conjugate Priors," *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3460–3474, 2013.

[15] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, "The Labeled Multi-Bernoulli Filter," *IEEE Transactions on Signal Processing*, vol. 62, no. 12, pp. 3246–3260, 2014.

[16] K. Murphy and S. Russel, "Rao-Blackwellized particle filtering for dynamic Bayesian networks," in *Sequential Monte Carlo methods in practice*, A. Doucet, N. d. Freitas, and N. Gordon, Eds. Berlin [etc.]: Springer, 2001, pp. 499–515.

[17] Y. Bar-Shalom, X. Tian, and P. K. Willett, *Tracking and data fusion: A handbook of algorithms*. Storrs and Conn: YBS Publishing, 2011.