

Map-aided Dead-reckoning using only Measurements of Speed

Johan Wahlström, Isaac Skog, *Member, IEEE*, João G. P. Rodrigues, *Student Member, IEEE*, Peter Händel, *Senior Member, IEEE*, and Ana Aguiar, *Member, IEEE*

Abstract—We present a particle-based framework for estimating the position of a vehicle using map information and measurements of speed. The filter propagates the particles' position estimates by means of dead-reckoning, and then updates the particle weights using two measurement functions. The first measurement function is based on the assumption that the lateral force on the vehicle does not exceed critical limits derived from physical constraints. The second is based on the assumption that the driver approaches a target speed derived from the speed limits along the upcoming trajectory. Assuming some prior knowledge of the initial position, performance evaluations of the proposed method indicate that end destinations often can be estimated with an accuracy in the order of 100 [m]. These results expose the sensitivity and commercial value of speed data collected in many of today's insurance telematics programs, where the data is used to adjust premiums and provide driver feedback. We end by discussing the strengths and weaknesses of different methods for anonymization and privacy preservation in telematics programs.

Index Terms—Vehicle positioning, dead-reckoning, insurance telematics, map-matching, usage-based-insurance, privacy.

I. INTRODUCTION

The emergence of connected vehicles has opened up possibilities for services which combine the power of cloud computing with the expanding capabilities of modern vehicle technology. Connected vehicles not only enable remote analysis of vehicle condition and driving behavior, but also allow drivers to benefit from a multitude of convenience-related applications such as theft tracking and accident detection. However, as the amounts of generated data and the means of connectivity increase, so do the dangers related to privacy and security [1]. Recently, connected vehicles and associated aftermarket devices have been shown to be vulnerable to wireless hacking via a broad range of attack vectors [2]–[4]. In the worst case, the hacker will be able to remotely manipulate any of the vehicle's electronic control units (ECUs). Moreover, privacy concerns have been raised about the large databases of sensitive data collected from vehicle-installed telematics units [5]–[8]. As illustrated in Fig. 1, a telematics insurer uses data on driving behavior to set the premium offered to each policyholder [9]–[11]. Typically, the risk profile of each driver will be based on e.g., the number of speeding or harsh braking

J. Wahlström, I. Skog, and P. Händel are with the ACCESS Linnaeus Center, Dept. of Signal Processing, KTH Royal Institute of Technology, Stockholm, Sweden. (e-mail: {jwahlst, skog, ph}@kth.se).

J. G. P. Rodrigues and A. Aguiar are with the Instituto de Telecomunicações, Departamento de Engenharia Eletrotécnica e de Computadores, Faculdade de Engenharia da Universidade do Porto, 4200-465 Porto, Portugal. (e-mail: {joao.g.p.rodrigues, ana}@fe.up.pt)

This work was supported by a grant from the Future Cities summer school, which was part of the Future Cities project at the University of Porto.




Fig. 1. Several insurance telematics providers are today collecting measurements of speed, here denoted by s_k , from vehicles' wheel speed sensors to adjust premiums and provide feedback to drivers.

events, as detected using measurements from vehicle-installed black boxes or from the vehicle's on-board-diagnostics (OBD) system [12], [13]. This study examines the locational privacy of such insurance telematics programs.

A. Locational Privacy

The digital revolution has made violations of locational privacy both cheaper and harder to detect [14]. Temporally isolated digital position inferences were first made possible by the widespread usage of credit cards and access cards. Today, the abundance of e.g., global navigation satellite system (GNSS) receivers in smartphones and connected vehicles enables continuous streams of position measurements to be sent to central servers [15], [16]. In addition, it is often possible to perform smartphone-based WiFi-localization by exploiting system permissions required by many popular apps [17]. Although some might argue that consumers are willingly giving up personal information in exchange for services and products, the actual privacy implications of sharing specific data are often obscured. This is especially true when the data itself does not include position measurements, but rather, information which in combination with additional databases can be used for positional inference.

Initially, privacy and security were not major concerns in the insurance telematics industry, and the amount of data collected was primarily determined by restrictions on transmission and storage. However, these topics have received more attention as the industry has matured, and privacy concerns have even been cited as a contributing reason to suspend telematics trials [18]–[21]. Today, several telematics providers only collect measurements of speed. When responding to privacy concerns, these companies will refer to the fact that no position data is recorded. However, as shown in [7], location-based information can still be extracted from many trips.

B. Contributions and Outline

In this article, we use map-aided dead-reckoning (DR) to geographically locate a vehicle purely based on its speed profile and digital map information. Since the study utilizes the same type of information that is available for many of the current insurance telematics providers, the results have privacy implications for millions of policyholders. The estimation is carried out using a particle filter, and thereby parallels previous filtering-based approaches where measurements of the vehicle's yaw rate have been available [22]–[26]. Though the idea of estimating a vehicle's position using only speed measurements and map information is not new [7], [8], this is the first time that it has been formulated as a filtering problem. In addition to making it possible to motivate the implementation using optimality arguments, our hope is that this also will make the method more accessible, and facilitate extensions based on established methods for particle filtering. Previous algorithms have been very reliant on the matching of vehicle stops in the speed data with intersections or traffic signals in the map. For example, in [8], data where the vehicle made a large number of stops due to traffic congestion or similar was removed altogether since this type of driving was not accounted for in the algorithm. In this study, we avoid these issues by processing the speed measurements one-by-one, rather than in segments delimited by vehicle stops. Moreover, the performance evaluations are made on data sets with significantly longer trip lengths than in previous studies.

Section II reviews the state-of-the-art in map-aided dead-reckoning and discusses how the problem is altered depending on what sensor information is available. The employed system model and the filter implementation are detailed in Section III. Section IV illustrates the performance characteristics of the proposed method by studying the accuracy of estimated end positions and the dependence on a priori knowledge of the initial position. The results indicate that GNSS or OBD measurements of speed often are sufficient to extract a substantial amount of personal information on daily activities and location-based behavior. Section V discusses the implications of the presented results and possible ways forward for privacy-aware telematics providers. Finally, the article is concluded in Section VI.

II. PRELIMINARIES ON MAP-AIDED DEAD-RECKONING

DR refers to the process of recursively estimating the position of an object by numerically integrating measurements or estimates of velocity. The general concept has been widely used to navigate e.g., aircrafts, ships, submarines, robots, automobiles, and pedestrians [27]. Obviously, perfect integration requires perfect knowledge of the velocity at all time instances. In practice, this means that stand-alone dead-reckoning systems always will be subject to position errors that accumulate with time. To prevent the position error from growing without bound, additional information is required. This type of information can be provided by, e.g., satellite-based positioning systems [12], WiFi fingerprinting systems [28], magnetic fingerprinting systems [29], terrain measurements [30], or road maps.

A. Map-aided DR with Sensor-based Yaw Information

DR aided only by information from road maps (and some a priori knowledge of the initial position) have been the topic of several studies within land-vehicle navigation [31]. Historically, the primary motivation of these studies have been to increase the accuracy and integrity of existing low-cost navigation systems during failures or performance deteriorations of the employed positioning systems (e.g., during GNSS outages). Assuming that measurements of both left and right wheel speeds are available from the controller area network (CAN) bus, estimates of the vehicle's longitudinal velocity and yaw rate can be obtained by means of differential geometry [22]. These estimates are then used for DR. If only a scalar speed measurement is available at each sampling instance, measurements of the yaw angle or yaw rate can instead be obtained from the steering wheel angle sensor, from a vehicle-mounted gyroscope, or from a magnetometer [32]. The error growth of the navigation solution is typically mitigated by comparing the current position and yaw estimates with the corresponding quantities at the nearest map segment [23], [33]. To avoid the logistics of accessing measurements of wheel speed, DR can instead be performed by means of inertial navigation. However, this will increase the number of integration steps, and thereby also the rate of the position error growth. As a result, the position estimates will be very sensitive to e.g., mounting misalignments and sensor biases. In general, the performance of inertial measurement unit (IMU)-based methods can be expected to be inferior to that of odometric methods [24].

B. Map-aided DR without Sensor-based Yaw Information

Since the standard parameters available from the vehicle's OBD system do not include any information on the vehicle's yaw rate or yaw angle, it is interesting to see what level of positioning accuracy that can be obtained by using only OBD measurements of speed. (As previously mentioned, the industry of insurance telematics today includes several actors, e.g., Progressive, Geico, and Allstate, that collect speed measurements from their policyholders [34], [35].) Related studies were conducted in [7] and [8], which estimated the position of a vehicle by matching intersections in the map with detected vehicle stops in the speed measurements. A list of possible map traces was continuously updated by adding new traces after each vehicle stop, and deleting traces deemed infeasible after 1) comparing their lengths with the corresponding distance implied by the speed measurements; 2) comparing the measured speeds with speed limits; and 3) comparing the measured speeds with the maximum speeds as implied by the curvature of the road. The estimation was not statistically motivated and was not formulated within any established estimation framework. In this study, we aim to fill this gap.

The DR problem arising when only speed measurements are available is fundamentally different from the DR problem when also sensor-based yaw information is available. In the latter case, the position estimate can be propagated in two dimensions using only sensor measurements. Spatial information

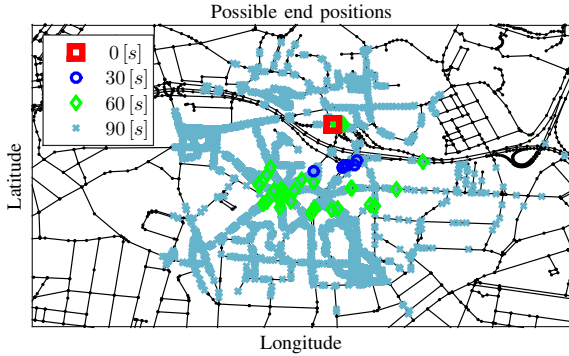


Fig. 2. Possible positions 0 [s], 30 [s], 60 [s], and 90 [s] after initializing DR (using only measurements of speed and map information) from a known starting point. The example is taken from a trip in Porto, Portugal, and the figure shows an area of about $3.5 [(km)^2]$ ($1.4 [km] \times 2.5 [km]$).

from the road map is then conveniently incorporated into the estimation as soft constraints. Usually, this information will be sufficient to avoid filter divergence, and additional information, such as speed limits, will be of secondary importance [22], [23]. When only measurements of speed are available, the position estimate cannot be propagated solely based on sensor measurements. Instead, the vehicle's yaw angle has to be estimated from the road map. In this case, sensor measurements and spatial information from the road map will never be sufficient as the posterior distribution will continue to "spread out" at intersections (see Fig. 2). Additional models which relate the expected vehicle speed to the vehicle's position on the map are required, and speed limits will become of utmost importance.

III. MODEL AND ESTIMATION FRAMEWORK

We present a filtering-based approach to the problem of estimating a vehicle's position using map information and measurements of speed. This section describes both the employed system model and the particle-based filter implementation.

A. Overview of System Model

The map data from OpenStreetMap (OSM) can be described as a mathematical graph where each vertex (node) defines a latitude-longitude pair. The connectivity of the graph is defined by a set of ways, i.e., ordered lists of nodes indicating that there is an edge (link) between the first and the second node, between the second and the third node, etc. Each way can be characterized by attributes specifying road type, speed limits, restrictions on travel direction, etc. The problem at hand is to use information from a digital map such as OSM, together with measurements of speed, to estimate a vehicle's position. The considered estimation framework is based on the model

$$\mathbf{x}_{k+1} \sim p(\mathbf{x}_{k+1} | \mathbf{x}_k, s_k), \quad (1a)$$

$$\mathbf{y}(\mathbb{S}) \sim p(\mathbf{y}(\mathbb{S}) | \mathbb{X}). \quad (1b)$$

Here, \mathbf{x} is a state vector comprising the vehicle's two-dimensional position \mathbf{r} , and the vehicle's direction of travel. Assuming that the vehicle always travels along links in the




Fig. 3. Illustration of the process model when the vehicle passes a road junction. If the probability of the initial position \mathbf{r}_k is p , the probability of each of the updated positions \mathbf{r}_{k+1} are p/N where N denotes the number of links along which \mathbf{r}_{k+1} can be expected. The distance between \mathbf{r}_k and (any) \mathbf{r}_{k+1} , measured along the links, is equal to $\Delta t_k s_k$.

graph, the state vector can be given as an ordered list of two connected nodes and a constant indicating the ratio of the vehicle's distance to the first node and the distance between the two nodes. The vehicle's speed is denoted by s (no notational distinction is made between the true and measured speed). As can be seen, the speed measurements act both as input in the state equation (1a) and as measurements in the measurement equation (1b). Moreover, $p(\cdot | \cdot)$ is used to denote a conditional probability density function (pdf). All pdfs are implicitly conditioned on the available map information. Throughout the paper, we will follow the convention of using a subindex k to denote quantities at sampling instance k . The complete set of state vectors and speed measurements are then given by $\mathbb{X} \triangleq \{\dots, \mathbf{x}_{k-1}, \mathbf{x}_k, \mathbf{x}_{k+1}, \dots\}$ and $\mathbb{S} \triangleq \{\dots, s_{k-1}, s_k, s_{k+1}, \dots\}$, respectively. The employed measurements are collected in a generic function \mathbf{y} , and will be discussed more thoroughly in Section III-C. For clarity, we will in the text differentiate between s and \mathbf{y} by referring to speed measurements and measurements, respectively.

B. Process Model

The process model (1a) describes the time development of the vehicle's position, using speed as input. The update is performed by distributing the travelled distance $\Delta t_k s_k$ along the upcoming links. We have here used Δt_k to denote the sampling interval between sampling instances k and $k+1$. When the vehicle passes a road junction, it is modeled to continue along any of the connected links (excluding the link it has just traversed) with equal probability. This is illustrated in Fig. 3. As may be realized, less probable events such as unexpected u-turns have been disregarded. If needed, the model can easily be modified to account for errors in the estimated travelled distance. Discrepancies between the estimated travelled distance $\Delta t_k s_k$ and the true travelled distance can arise due to map errors and off-road driving [25]. In addition, GNSS speed measurements are subject to multipath propagation [36], whereas OBD measurements are affected by wheel slips, quantization errors in the measured speed [37], and scale factor errors dependent on tire pressure [22]. Previous attempts to include the scale factor in the state vector did not result in any improvement in navigation performance, but rather, increased the risk of filter divergence [23], [25]. We further note that the

speed measured by the OBD system is the absolute value of the vehicle's three-dimensional velocity, whereas the studied model only considers two spatial dimensions (most nodes in OSM do not include any altitude information). As a result, the travelled distance in the two-dimensional plane of the map will tend to be overestimated when travelling along steep inclines [23]. For example, if the vehicle is traveling with a speed of $25 [m/s]$ along a road where the angle of inclination is $5 [^\circ]$, the speed in the two-dimensional plane of the map will be $25 \cdot \cos(5 [^\circ]) [m/s] \approx 24.62 [m/s]$. This error contribution can be compensated for by using a commercial digital map with altitude information and extending the dimension of the state vector accordingly.

Obviously, the estimated travelled distance $\Delta t_k s_k$ is subject to several modeling and measurement errors. Despite this, the resulting accumulated position error will in many cases be negligible in comparison to the position error resulting from the uncertainty of the vehicle path, i.e., the list of nodes or links that the vehicle has traversed.

C. Measurement Model

When propagating the state distribution as described in the preceding subsection, the number of feasible vehicle paths grows exponentially as a function of the traveled distance [7]. To bound the number of paths, measurement updates must be employed. We will give two examples of possible measurements. First, we study the lateral forces exerted on the vehicle. When cornering, the lateral g-forces on the vehicle at sampling instance k can be approximated by [38]

$$y^{(1)}(s_k, s_{k+1}) \triangleq \sqrt{a_k^2 + s_k^2 \omega_k^2} \quad (2)$$

where a denotes the vehicle's longitudinal acceleration and ω denotes the vehicle's yaw rate. The acceleration is approximated by the difference quotient

$$a_k \triangleq \frac{s_{k+1} - s_k}{\Delta t_k}. \quad (3)$$

The yaw rate ω must be estimated from map data. We choose to estimate the yaw rate in a separate Kalman filter based on the state-space model

$$\theta_{k+1} = \theta_k + \Delta t_k \omega_k, \quad (4a)$$

$$\omega_{k+1} = \omega_k + \eta_k^\omega, \quad (4b)$$

$$\tilde{\theta}_k = \theta_k + \epsilon_k^\theta. \quad (4c)$$

Here, the noises η_k^ω and ϵ_k^θ are assumed to be white with variances $\Delta t_k \sigma_\omega^2$ and σ_θ^2 , respectively. Further, θ is the vehicle's yaw angle, and $\tilde{\theta}$ is the direction of the current link. Due to the periodicity of the yaw angle, the measurement residuals must be chosen within the interval $[-180, 180] [^\circ]$. An alternative to using a Kalman filter is to estimate the yaw rate directly from the difference quotients $(\tilde{\theta}_{k+1} - \tilde{\theta}_k) / \Delta t_k$. However, due to the limited resolution of the nodes in OSM (many sharp corners will have links meeting at a $90 [^\circ]$ angle) these difference quotients will be fairly unstable.

Typically, passenger cars are not designed for g-forces higher than $0.8 g$, and will seldom exceed $0.6 g$ during normal driving. (A sufficiently high lateral force will cause the vehicle

to rollover [38].) Here, g denotes the gravitational acceleration at the surface of the earth. Motivated by this, the distribution of $y^{(1)}$ is heuristically modeled as

$$p(y^{(1)}(s_k, s_{k+1}) | \mathbf{x}_1, \dots, \mathbf{x}_k) \propto \begin{cases} 1, & y^{(1)} < g^{(1)} \\ \frac{y^{(1)} - g^{(2)}}{g^{(1)} - g^{(2)}}, & g^{(1)} \leq y^{(1)} < g^{(2)} \\ 0, & g^{(2)} \leq y^{(1)} \end{cases}. \quad (5)$$

The design parameters $g^{(1)}$ and $g^{(2)}$ are chosen so that the lateral forces fall below $g^{(1)}$ during the larger part of the driving while allowing for outliers giving estimated lateral forces up to $g^{(2)}$.

To derive the second type of measurements, we note that drivers tend to maintain a speed close to the speed limit for the larger part of the route. Hence, we propose the use of the measurements

$$y^{(2)}(s_k, s_{k+1}) \triangleq a_k + c s_k \quad (6)$$

distributed according to

$$p(y^{(2)}(s_k, s_{k+1}) | \mathbf{x}_1, \dots, \mathbf{x}_{k+K}) = p_C(y^{(2)}, c \bar{s}_{k+K}, \sigma) \quad (7)$$

where $K \in \mathbb{N}$ is a design parameter describing the driver's tendency to adjust his speed to the road ahead, and p_C denotes the pdf of the Cauchy distribution with location and scale parameters given by the second and third arguments, respectively. The probabilistic assumption in (7) is derived from the model

$$a_k = c(\bar{s}_{k+K} - s_k) + \epsilon_k^s \quad (8)$$

which says that the vehicle's acceleration will be approximately proportional to the difference between the current speed and the target speed \bar{s}_{k+K} . As should be obvious, we have assumed that ϵ_k^s is Cauchy distributed with location 0 and scale σ . The design parameters c and σ describe the characteristics of each individual driver and car. Modeling the error ϵ as distributed according to the heavy-tailed Cauchy distribution will make the estimation robust to measurement errors caused by unexpected driving behavior.

The target speed is, just as the yaw rate, estimated in a separate Kalman filter. In this case, the state-space model is

$$\bar{s}_{k+1} = \bar{s}_k + \eta_k^s. \quad (9a)$$

$$\tilde{s}_k = \bar{s}_k + \epsilon_k^s. \quad (9b)$$

Here, the noises η_k^s and ϵ_k^s are white with variances $\Delta t_k \sigma_{s1}^2$ and σ_{s2}^2 , respectively, while \tilde{s} is the speed limit of the current link. By letting the target speed be dependent on the speed limits on the road ahead, we are able to capture the forward-looking behavior of typical driving. At each iteration, the estimated target speed is projected to the closest speed in the space

$$\{s_k : \sqrt{a_k^2 + s_k^2 \omega_k^2} < g^{(1)}\}. \quad (10)$$

In other words, the estimated target speed is adjusted so as to never imply an excessively large lateral force. This means

Algorithm 1 : Particle-filter algorithm.

- 1: **for** $k = 0, 1, \dots$ **do**
 - 2: *Time update*: Update each particle state \mathbf{x}_k and its associated particle weight w_k by distributing the travelled distance $\Delta t_k s_k$ along the upcoming links.
 - 3: *Parameter update*: Perform a time update in the Kalman filters estimating the yaw angle, the yaw rate, and the target speed for each particle. Use the travel direction and the speed limit at each particle to perform the associated measurement updates.
 - 4: *Measurement update*: Update the particle weights according to equation (11). The estimated yaw rate is used in the updates based on both $y^{(1)}$ and $y^{(2)}$, whereas the estimated target speed only is needed in the update based on $y^{(2)}$. Set the weights of all particles which violate constraints on travel directions to zero.
 - 5: *Particle weight redistribution*: Redistribute the prior weight of particles which have had their weight set to zero in the measurement update to particle "siblings" (if such siblings exists).
 - 6: *Particle elimination*: Eliminate all particles with zero weight. If the total number of particles exceeds a given threshold, eliminate all particles with insufficient weight.
 - 7: *Particle resampling*: If $k = n \cdot N_0$ for some $n \in \mathbb{N}$ and some integer parameter N_0 , and if the number of particles is sufficiently low, resample particles using systematic resampling. Displace each sampled particle along its current link by sampling from a uniformly distributed distance. Re-initialize the particles so that all particle weights are equal.
 - 8: *Particle merging*: Merge particles that are sufficiently similar.
 - 9: **end for**
-

The algorithm is detailed in Sections III-B, III-C, and III-D.

that particles which pass sharp corners while the vehicle slows down often will have their weights increased, and hence, we can detect the vehicle turning even though all nearby links have the same speed limit.

As an alternative to using the model (8), it is possible to derive a measurement function based on the assumption that the current speed is close to the target speed. This model is obtained as a special case of (8) in the limit of $\bar{s}_{k+K} = s_k$. However, this model tends to be too restrictive on discrepancies between the vehicle speed and the target speed, and often causes the measurement errors to have a significant temporal correlation (e.g., if the vehicle's speed exceeds the target speed at one sampling instance it is very likely to do so also at the next sampling instance). By using (8), we allow for temporally correlated discrepancies between the vehicle speed and the target speed, but still expect the speed to be at least approaching the target speed.

Both types of measurements $y^{(1)}$ and $y^{(2)}$ are based on ideas similar to those in [7]. Additional measurements can be derived from the fact that vehicle stops tend to coincide with traffic lights or stop signs in the map. This information was not utilized here as the number of unpredictable stops (e.g.,

stops due to congestion) was very large, and information on traffic lights and stop signs was lacking over the larger part of the map area.

Since all speed measurements are used both in the process model (1a) and in the measurements (2) and (6) one might argue that the system model includes several error correlations that must be considered in the estimation. However, the larger part of the variance in $y^{(1)}(s_k, s_{k+1})|\mathbf{x}_1, \dots, \mathbf{x}_k$ and $y^{(2)}(s_k, s_{k+1})|\mathbf{x}_1, \dots, \mathbf{x}_{k+K}$ stem from variations in driving behavior, rather than from errors in the measured speed (for example, the absolute OBD measurement error is typically smaller than 1 [km/h], while the speed at a given position often can differ by 20 [km/h] depending on the traffic conditions). As a result, the correlations caused by errors in the measured speed are generally negligible.

D. Particle Filter Implementation

Due to the multimodal characteristics of the conditional pdfs in the system model (1), assuming unimodal posterior distributions will typically lead to irrecoverable localization errors in a very short period of time. To avoid this, we propose a solution based on the particle filter framework [39]. Each particle must, in addition to the current state vector \mathbf{x}_k , also store its estimates of the yaw angle, the yaw rate, and the target speed. However, as the state-space models for these parameters are time-invariant (up to variations in the sampling interval), the covariance matrices only need to be stored once (i.e., while each particle needs to run its own filter, the filter parameters will be the same for all particles). The particle propagation is performed according to the model described in Section III-B. Instead of sampling a single continuing path when a particle passes a road junction (i.e., ignoring all but one of the possible links), the particle is split into several new particles with equal weights (the sum of which is equal to the weight of the original particle). In this way, we increase the particle diversity and decrease the risk of filter divergence due to unlucky sampling. Obviously, this also means that the number of particles in many cases will increase during the propagation step.

Aside from the resulting increase in computational cost, the choice of "splitting" particles as they pass a road junction has an additional downside. This becomes clear when the true path of the vehicle passes a large number of road junctions. In this case, a particle following this path will have its weight diminished with each passing road junction. If the weight of the particle is not to fall short of resampling thresholds, it must then increase substantially during measurement updates. In practice, this means that particles sometimes will be eliminated mainly because they passed many road junctions. To mitigate this effect, we once again utilize the fact that the sum of the weights of particles emerging from a particle passing a road junction should be equal to the weight of the original particle. Hence, if a particle which has emerged from a "particle split" at a road junction is immediately eliminated due to an estimated lateral force exceeding $g^{(2)}$, its weight, prior to the elimination, is distributed to its "siblings", i.e., other particles which originated from the same road junction and particle. The

same rules are applied to particles which violate constraints on travel directions, i.e., which travel in the wrong direction on a one-way link. Implementing this modification will tend to reduce the risk of filter divergence when the vehicle travels along a main road with a large number of intersections.

As motivated in the last paragraph in Section III-B, the particle propagation is performed without any simulated process noise in the travelled distance. Similar modeling choices have previously been discussed in [40] and references therein.

In the measurement update, each particle weight w_k is updated according to

$$w_{k+1} \propto p(\mathbf{y}_{k+1}(\mathbb{S})|\mathbb{X})w_k \quad (11)$$

where the total set of measurements are

$$\mathbf{y}_{k+1}(\mathbb{S}) \triangleq [y^{(1)}(s_k, s_{k+1}) \quad y^{(2)}(s_{k-K}, s_{k+1-K})]^\top. \quad (12)$$

Note that the measurements $y^{(2)}$ use earlier speed measurements than $y^{(1)}$ as the conditional distribution of $y^{(2)}(s_{k-K}, s_{k+1-K})$ is modeled as dependent on the target speed at sampling instance k .

The filter is implemented with systematic resampling [41]. New particles are resampled at equidistant points in time given that the number of particles is sufficiently low. Each new particle is displaced with a uniformly distributed distance along its trajectory. This will make it possible to detect any eventual accumulated errors in the total travelled distance. Particles with insufficient weight are eliminated whenever the total number of particles exceed a given threshold. However, particles which are given a zero weight are eliminated immediately. In addition, whenever two particles are sufficiently similar (i.e., are on the same link, have the same parameter estimates, and have sufficiently similar position estimates) they are merged into a single particle which will have its weight equal to the sum of the original weights and its position equal to the weighted average of the original positions. This will ease the computational burden somewhat as well as mitigate problems with particle clustering [22]. The particle filter is summarized in Algorithm 1.

IV. FIELD STUDY

GNSS and OBD data were collected from five drivers performing a total of eighteen different trips in both rural and urban areas in and around Porto, Portugal. No specific trip was conducted for the purpose of this study, and hence, the data represent a sample of typical routes in the area. However, since we were interested in the algorithm's ability to avoid divergence during longer trips, the studied trips were somewhat longer than what would have been expected from a completely random sample. All data were collected using the SenseMyCity app, which was developed as part of the Future Cities project [42]. The median and mean trip lengths were 28.3 [km] and 49.5 [km], respectively. The update rate of the GNSS data was 1 [Hz], while the update rate of the OBD data varied between 1 [Hz] and 3 [Hz] depending on the car (the majority of the trips had an update rate close to 1 [Hz]). The GNSS position measurements were used as ground truth. Assuming knowledge of the vehicle's initial position (this will

TABLE I
FILTER PARAMETERS.

	Parameter	Value
<i>Process noise</i>	σ_ω	5 [°/s ² /√Hz]
	σ_{s1}	0.5 [m/s ² /√Hz]
<i>Measurement noise</i>	σ_θ	15 [°]
	σ_{s2}	10 [m/s]
$y^{(1)}$ update	$g^{(1)}$	0.55 g
	$g^{(2)}$	0.65 g
$y^{(2)}$ update	K	12
	c	0.05 [1/s]
	σ	1.5 [m/s ²]
Particle threshold (resampling trigger) [†]		100
Weight threshold (elimination)		1/200
Distance threshold (particle merge)		20 [m]
Resampling displacement interval		(−100, 100) [m]

[†] Checked every hundredth second.

often be the driver's home address, the driver's work address, etc.), the filter was initialized by sampling particles on all links in a circle of radius 50 [m] centered at the true initial position. Due to the stochastic nature of the filter, all displayed performance measures have been averaged over 10 runs. The parameters used in the field study are shown in Table I.

Occasionally, data was lost from either of the two sensors. In these cases, the missing data was replaced by data from the other sensor (GNSS replacing OBD or OBD replacing GNSS)¹. The replacement data was scaled to take the OBD scale factor into account. The scale factor can be estimated using data from sampling instances where measurements from both sensors are available. While some GNSS outages always can be expected, we stress that the larger part of the data collection disruptions occurred in the OBD data and was related to the chosen communications setup. In other words, data losses of this kind should not be expected in e.g., a commercial OBD-based insurance telematics program. In two (five) of the trips, data losses caused the measurements of the trip to start (end) at nonzero speeds. If the data starts or ends close to the motorway, this will usually simplify the estimation since filter divergence is more prone to occur along roads with lower speed limits.

A. Results

The performance of the filter when using GNSS and OBD measurements of speed are shown in Fig. 4 and Fig. 5, respectively. The figures display the empirical distribution functions (edfs) of the horizontal $\|\hat{\mathbf{r}} - \mathbf{r}\|$ and relative horizontal position errors $\|\hat{\mathbf{r}} - \mathbf{r}\|/L$ of the estimated end position. Here, we have used L to denote the driving length of the trip. If all particles were eliminated while running the filter, the horizontal position error was considered to be equal to the driving length of the trip. The performance of two estimators are shown. These are

¹Since the errors of the GNSS and OBD measurements are in the same order of magnitude, we do not expect this to distort the results in any substantial way.




Fig. 4. The empirical distribution function of the horizontal and relative horizontal position errors of the end position. The filter used GNSS measurements of speed.

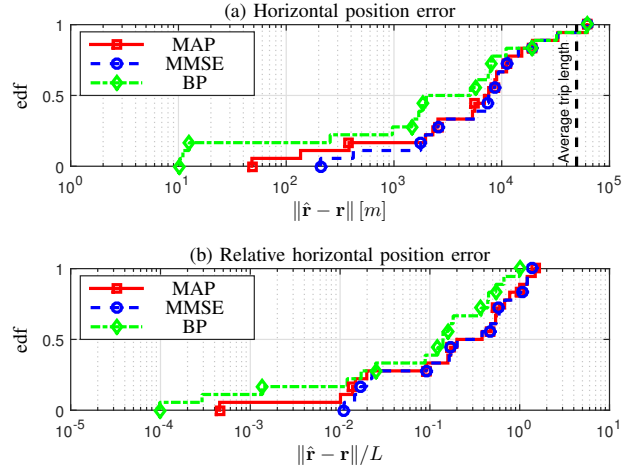


Fig. 5. The empirical distribution function of the horizontal and relative horizontal position errors of the end position. The filter used OBD measurements of speed.

the maximum a posteriori (MAP) estimator, i.e., the position of the particle with the largest weight, and the minimum mean square error (MMSE) estimator, i.e., the weighted average of all particle's end points. In addition, we also display the accuracy of the particle with the position estimate that is closest to the true end position. This is called the best-particle (BP). Obviously, this is not a realizable estimator as we can only find the BP by assuming knowledge of the true end position. However, we still believe the accuracy of the BP is relevant as it gives an indication of the performance that can be expected when additional information on likely end positions are available (whereas the MAP and MMSE estimators give an indication of the performance that can be expected when processing trips independently without access to any additional personal information). Many drivers tend to start and end most of their trips at a very limited number of locations. Hence, when a vehicle is detected to have stopped for an extended period of time, the current particles can be evaluated based on personal information, previously extracted driving habits, and map information (for example, a driver will seldom stop in the middle of the motorway). Under ideal circumstances, only one particle will be considered to be a realistic end point.

As can be seen, the BP lies less than 200 [m] from the end position in about a fourth of the trips, and lies less than 2 [km] from the end position in about half of the trips. If the true end position is e.g., the driver's home, this level of accuracy will often be sufficient to determine that the driver went home during the trip. Similarly, the MAP estimator is accurate to within approximately 200 [m] in a tenth of the trips. Moreover, the errors of the MAP and MMSE estimators are about a fifth of the total trip length in about half of the trips. This means that even though we are not able to accurately estimate the end position of a trip, we might still be able to estimate the overall direction of travel (e.g., whether the driver drove towards or away from the city center).

The dependence on the initial uncertainty is illustrated in Fig. 6 which shows the quantiles $\text{edf}^{-1}(0.1)$ and $\text{edf}^{-1}(0.25)$ of the BP estimate as dependent on the uncertainty radius, i.e., the radius of the circle in which the initial particles were




Fig. 6. The empirical 0.1 and 0.25 quantiles of the horizontal position error of the end position of the BP, as dependent on the radius of the circle in which the initial particles were sampled.

sampled. All other filter parameters were kept constant. Even with an uncertainty radius of 1 [km], a fourth of the trips have a horizontal position error below 600 [m]. This demonstrates the capability to extract positional information from speed measurements even when there is significant uncertainty in the initial position.

The estimated OBD scale factor had a sample mean and sample standard deviation (over the studied trips) of 1.012 and 0.024, respectively. Despite this, the choice of using GNSS or OBD measurements of speed seem to have a limited effect on the filter performance.

B. Filter Divergence

When no particle is closer than approximately 200 [m] to the true position, the filter should be considered to have diverged. In other words, the filter ultimately diverged in a majority of the trips studied here. However, in many cases the filter did not diverge until late into the trip, and hence, the error of the end position could still be small as compared to the total driving length. The filter will usually diverge in urban environments where the road density is high and most roads have the same speed limit. In these cases, the number of possible paths grows very fast and the measurement updates provide little information. While the rate of divergence may seem rather high, the results (e.g., an MMSE estimator with a median horizontal error of about a fifth of the driving length) are comparable to those obtained in [7] where the median trip

TABLE II
CAPABILITY OF PROPOSED METHOD FOR NAVIGATION.

Inference of end destination	Share of trips
Exact location ($\sim 200 [m]$)	1/4
Approximate region ($\sim 2 [km]$)	1/2

length was $7.5 [km]$ (about four times shorter than in the data used here), the maximum trip length was $16.0 [km]$ (shorter than both the median and mean trip length in the data used here), and no uncertainty in the initial position was mentioned. It can also be noted that filter divergence occasionally can occur even if measurements of yaw rate are available [22]–[24].

V. DISCUSSION

In the preceding section, digital map information and measurements of speed were used to geographically locate a car during everyday trips. The estimates were shown to be accurate enough to enable violations of the driver’s locational privacy in a substantial amount of the trips. (A summary of the indicated capability is provided in Table II.) As a result, this also means that the data collected in many of the current insurance telematics programs is far more sensitive than stated by the insurance companies themselves. While many of these companies will have privacy policies that regulate their data practices, the data can easily be misused if it falls into the wrong hands. Moreover, since the privacy risks will not always be obvious for someone without experience in navigation, the data might not be treated as sensitive, thereby increasing the risk of data theft. Normally, the data is saved indefinitely without consideration of any “right to be forgotten”.

Insurance companies offering telematics services will often have years’ worth of data attributed to individual policyholders, and can therefore make use of previously extracted personal information (e.g., commonly visited locations) to improve the position estimation over time. One way to do this could be to apply more sophisticated weight distributions in the propagation step presented in Section III-B. Put differently, the weights of different particles emerging from the crossing of a road junction could be adjusted based on previous inferences on driving habits.

The prospect of being able to extract positional information from speed measurements does not only pose privacy risks, but also means that insurance companies can use a larger number of features to distinguish and characterize policyholders. For example, while speed measurements alone only enable studies of absolute speeding, i.e., whether the driver exceeds a given speed threshold, accurate knowledge of both speed and position makes it possible for the insurer to directly identify eventual speed limit violations. In addition, the insurer can penalize driving in areas where accidents are prone to happen.

There are several ways to implement a privacy-preserving insurance telematics program. For example, the data assembled at the central server does not need to reveal from which vehicle a specific set of data was gathered. This approach is taken in [43], which associates each collected data tuple with a

time stamp and a random time-varying vehicle identifier. Any information extraction (e.g., the computation of a driver score) requiring data from a single vehicle must then be performed as a secure multi-party computation involving both the server and the vehicle-fixed client application. In this case, ideal privacy preservation means that the server data by itself cannot be used to extract more information about an individual vehicle’s data than can be done from the same set of data but without any vehicle identifiers. However, if position data is collected, it may still be possible to identify data from policyholders who live in secluded areas (since they might be the only people driving in this area) [44], [45]. The corresponding location traces may then be found by using that location updates close in time and space are likely to originate from the same vehicle [46]. Further, location traces from different drivers can be clustered by studying individual driving characteristics, so called driver fingerprinting [47]. Obviously, it is possible to reduce the risk of inference attacks by simply reducing the data quality, assuming that the primary privacy-preserving information extraction can still be performed with tolerable errors. This can be done by e.g., omitting samples or deliberately perturbing measurements [48].

Many insurance telematics providers are not interested in speed measurements or location updates per se, but rather, in the driver score or risk profile that can be extracted from them. Hence, it will often be sufficient to compute the measures of interest directly on the telematics unit, send these to a central server, and discard the original measurements [49]. This will both reduce the required communications and the risks of privacy intrusion. On the downside, this may increase the computational demands on the telematics units.

VI. CONCLUSIONS

We have considered the problem of estimating the position of a vehicle using only map information and measurements of speed. As opposed to previous studies utilizing the same set of sensors, the problem was here formulated within the well-known particle filter framework. Particle-based estimation makes it straightforward to 1) study the estimation accuracy using estimators suited for multimodal posteriors; 2) draw use of experience from previous filter-based approaches where sensor-based yaw information have been available; and 3) extend the presented framework with additional measurement functions. Performance evaluations of the proposed method indicated that location-based information such as end destinations often can be estimated with an accuracy in the order of $100 [m]$. This level of accuracy will in many cases be sufficient to extract information on e.g., visited locations or areas. As a result, the collection of speed measurements from policyholders in current insurance telematics programs both poses privacy risks and makes it possible for the insurer to discriminate among drivers based on visited areas. Telematics providers who wishes to decrease the risk of privacy infringement for their policyholders can for example make use of time-varying vehicle identifiers, or perform the computations of relevant driver measures directly in the vehicle-fixed telematics units.

REFERENCES

- [1] P. Lawson, "The connected car: Who's in the driver's seat?" Mar. 2015, FIPA.
- [2] L. Pike, J. Sharp, M. Tullsen, P. C. Hickey, and J. Bielman, "Securing the automobile: A comprehensive approach," in *Embedded Security in Cars*, Detroit, MI, May 2015.
- [3] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno, "Comprehensive experimental analyses of automotive attack surfaces," in *Proc. 20th USENIX Conf. Security*, Berkeley, CA, Aug. 2011.
- [4] I. Foster, A. Prudhomme, K. Koscher, and S. Savage, "Fast and vulnerable: A story of telematic failures," in *Proc. 9th USENIX Workshop on Offensive Technol.*, Washington, D.C., Aug. 2015.
- [5] S. Duri, J. Elliott, M. Gruteser, X. Liu, P. Moskowitz, R. Perez, M. Singh, and J.-M. Tang, "Data protection and data sharing in telematics," *Mob. New. Appl.*, vol. 9, no. 6, pp. 693–701, Dec. 2004.
- [6] N. Rizzo, E. Sprissler, Y. Hong, and S. Goel, "Privacy-preserving driving style recognition," in *Proc. IEEE Int. Conf. Connected Veh. Expo*, Shenzhen, China, Oct. 2015.
- [7] X. Gao, B. Firner, S. Sugrim, V. Kaiser-Pendergrast, Y. Yang, and J. Lindqvist, "Elastic pathing: Your speed is enough to track you," in *Proc. ACM Int. Joint Conf. Pervasive and Ubiquitous Comput.*, Seattle, WA, 2014.
- [8] R. Dewri, P. Annadata, W. Eltarjaman, and R. Thurimella, "Inferring trip destinations from driving habits data," in *Proc. 12th ACM Workshop on Privacy in the Electron. Soc.*, Berlin, Germany, Nov. 2013, pp. 267–272.
- [9] P. Händel, I. Skog, J. Wahlström, F. Bonawiede, R. Welch, J. Ohlsson, and M. Ohlsson, "Insurance telematics: Opportunities and challenges with the smartphone solution," *IEEE Intell. Transport. Syst. Mag.*, vol. 6, no. 4, pp. 57–70, Oct. 2014.
- [10] J. Wahlström, I. Skog, and P. Händel, "Driving behavior analysis for smartphone-based insurance telematics," in *Proc. 2nd Workshop on Physical Analytics*, Florence, Italy, May 2015, pp. 19–24.
- [11] J. Engelbrecht, M. J. Booyens, G.-J. van Rooyen, and F. J. Bruwer, "Survey of smartphone-based sensing in vehicles for intelligent transportation system applications," *IET Intell. Transport. Syst.*, vol. 9, no. 10, pp. 924–935, Dec. 2015.
- [12] J. Wahlström, I. Skog, and P. Händel, "IMU alignment for smartphone-based automotive navigation," in *Proc. 18th IEEE Int. Conf. Inf. Fusion*, Washington, DC, Jul. 2015, pp. 1437–1443.
- [13] J. Engelbrecht, M. J. Booyens, and G.-J. van Rooyen, "Recognition of driving manoeuvres using smartphone-based inertial and GPS measurement," in *Proc. 1st Int. Conf. Use of Mobile ICT*, Stellenbosch, South Africa, Dec. 2014.
- [14] A. J. Blumberg and P. Eckersley, "On locational privacy, and how to avoid losing it forever," Aug. 2009, Electron. Frontier Foundation.
- [15] M. Iqbal and S. Lim, "Privacy implications of automated GPS tracking and profiling," *IEEE Technol. Soc. Mag.*, vol. 29, no. 2, pp. 39–46, Jun. 2010.
- [16] K. Michael, A. McNamee, M. G. Michael, and H. Tootell, "Location-based intelligence - Modeling behavior in humans using GPS," in *Proc. IEEE Int. Symp. Technol. Soc.*, Queens, NY, Jun. 2006, pp. 1–8.
- [17] L. Nguyen, Y. Tian, S. Cho, W. Kwak, S. Parab, Y. Kim, P. Tague, and J. Zhang, "UnLocIn: Unauthorized location inference on smartphones without being caught," in *Proc. Int. Conf. Privacy and Security in Mobile Syst.*, Atlantic City, NJ, Jun. 2013, pp. 1–8.
- [18] C. Troncoso, G. Danezis, E. Kosta, J. Balasch, and B. Preneel, "PriPAYD: Privacy-friendly pay-as-you-drive insurance," *IEEE Trans. Dependable and Secure Comput.*, vol. 8, no. 5, pp. 742–755, Sep. 2011.
- [19] J. Paefgen, T. Staake, and F. Thiesse, "Evaluation and aggregation of pay-as-you-drive insurance rate factors: A classification analysis approach," *Decision Support Syst.*, vol. 56, pp. 192–201, Dec. 2013.
- [20] M. Courtney, "Premium binds," *Eng. Technol.*, vol. 8, no. 6, pp. 68–73, Jul. 2013.
- [21] J. Ohlsson, P. Händel, S. Han, and R. Welch, *BPM - Driving Innovation in a Digital World*. Springer, 2015, ch. Process Innovation with Disruptive Technology in Auto Insurance: Lessons Learned from a Smartphone-Based Insurance Telematics Initiative, pp. 85–101.
- [22] P. Hall, "A Bayesian approach to map-aided vehicle positioning," Master's thesis, Linköping University, Jan. 2001.
- [23] N. Svenzén, "Real time implementation of map aided positioning using a Bayesian approach," Master's thesis, Linköping University, Dec. 2002.
- [24] G. Hedlund, "Map aided positioning using an inertial measurement unit," Master's thesis, Linköping University, Nov. 2008.
- [25] J. Kronander, "Robust automotive positioning: Integration of GPS and relative motion sensors," Master's thesis, Linköping University, Dec. 2004.
- [26] A. Selloum, D. Bétaille, E. L. Carpentier, and F. Peyret, "Robustification of a map aided location process using road direction," in *Proc. 13th Int. Conf. Intell. Transport. Syst.*, Funchal, Portugal, Sep. 2010, pp. 1504–1510.
- [27] P. D. Groves, *Principles of GNSS, inertial, and multisensor integrated navigation systems*, 1st ed. Artech House, 2008.
- [28] Y. Chen and H. Kobayashi, "Signal strength based indoor geolocation," in *Proc. IEEE Int. Conf. Commun.*, vol. 1, Apr., New York, NY 2002, pp. 436–439.
- [29] J. Haverinen and A. Kemppainen, "Global indoor self-localization based on the ambient magnetic field," *Robot. Auton. Syst.*, vol. 57, no. 10, pp. 1028–1035, Oct. 2009.
- [30] N. Bergman, "A Bayesian approach to terrain-aided navigation," in *IEEE Proc. 11th Symp. Syst. Identification*, Fukuoka, Japan, Sep. 1997, pp. 1531–1536.
- [31] F. Gustafsson, U. Orguner, T. Schön, P. Skoglar, and R. Karlsson, *Handbook of Intelligent Vehicles*. Springer, 2012, ch. Navigation and tracking of road-bound vehicles, pp. 397–434.
- [32] P. Davidson, J. Collin, J. Raquet, and J. Takala, "Application of particle filters for vehicle positioning using road maps," in *Proc. 23rd Int. Tech. Meeting of the Satellite Division of the Institute of Navigation*, Portland, OR, Sep. 2010, pp. 1653–1661.
- [33] I. Skog and P. Händel, "In-car positioning and navigation technologies - A survey," *IEEE Trans. Intell. Transport. Syst.*, vol. 10, no. 1, pp. 4–21, Mar. 2009.
- [34] "The state of usage based insurance today," Apr. 2014, Ptolemus - Consulting Group.
- [35] S. Derikx, M. D. Reuver, M. Kroesen, and H. Bouwman, "Buying-off privacy concerns for mobility services in the internet-of-things era: A discrete choice experiment on the case of mobile insurance," in *28th Bled eConference*, Bled, Slovenia, Jul. 2015.
- [36] S. N. Sadrieh, A. Broumandan, and G. Lachapelle, "Doppler characterization of a mobile GNSS receiver in multipath fading channels," *The J. Navigation*, vol. 65, no. 7, pp. 477–494, Jul. 2012.
- [37] F. Gustafsson, "Rotational speed sensors: limitations, pre-processing and automotive applications," *IEEE Mag. Instrum. Meas.*, vol. 13, no. 2, pp. 16–23, Apr. 2010.
- [38] J. Wahlström, I. Skog, and P. Händel, "Detection of dangerous cornering in GNSS-data-driven insurance telematics," *IEEE Trans. Intell. Transport. Syst.*, vol. 16, no. 6, pp. 3073–3083, Dec. 2015.
- [39] N. Gordon, D. Salmond, and A. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEEE Proc. Radar and Signal Process.*, vol. 140, no. 2, pp. 107–113, Apr. 1993.
- [40] J.-O. Nilsson and P. Händel, "Recursive Bayesian initialization of localization based on ranging and dead reckoning," in *IEEE, Intell. Robots and Syst. Int. Conf.*, Tokyo, Japan, Nov. 2013, pp. 1399–1404.
- [41] J. D. Hol, "Resampling in particle filters," Linköping University, Tech. Rep., May 2004.
- [42] J. G. Rodrigues, A. Aguiar, and J. Barros, "SenseMyCity: Crowdsourcing an urban sensor," *CoRR*, vol. arXiv:1412.2070, Dec. 2014.
- [43] R. A. Popa, H. Balakrishnan, and A. J. Blumberg, "VPriv: Protecting privacy in location-based vehicular services," in *Proc. 18th Int. Conf. USENIX Security Symp.*, Montreal, QC, Aug. 2009, pp. 335–350.
- [44] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady, "Enhancing security and privacy in traffic-monitoring systems," *IEEE Pervasive Comput.*, vol. 5, no. 4, pp. 38–46, Oct. 2006.
- [45] J. Krumm, "Inference attacks on location tracks," in *Proc. 5th Int. Conf. Pervasive Comput.*, Berlin, Germany, May 2007, pp. 127–143.
- [46] M. Gruteser and B. Hoh, "On the anonymity of periodic location samples," in *Proc. 2nd Int. Conf. Security in Pervasive Comput.*, Berlin, Germany, Apr. 2005, pp. 179–192.
- [47] K. K. Miro Enev, Alex Takakuwa and T. Kohno, "Automobile driver fingerprinting," in *Proc. Privacy Enhancing Technol.*, vol. 1, Jan. 2016, pp. 34–50.
- [48] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady, "Preserving privacy in GPS traces via uncertainty-aware path cloaking," in *Proc. 14th ACM Int. Conf. Comput. Commun. Security*, Alexandria, VA, Oct. 2007, pp. 161–171.
- [49] P. Händel, J. Ohlsson, M. Ohlsson, I. Skog, and E. Nygren, "Smartphone-based measurement systems for road vehicle traffic monitoring and usage-based insurance," *IEEE Syst. J.*, vol. 8, no. 4, pp. 1238–1248, Dec. 2014.



Johan Wahlström received his MSc degree in Engineering Physics from the KTH Royal Institute of Technology, Stockholm, Sweden, in 2014. He subsequently joined the Signal Processing Department at KTH, working towards his PhD. His main research topic is insurance telematics. In 2015, he received a scholarship from the Sweden-America foundation and spent six months at Washington University, St. Louis, USA.



Isaac Skog (S'09-M'10) received the BSc and MSc degrees in Electrical Engineering from the KTH Royal Institute of Technology, Stockholm, Sweden, in 2003 and 2005, respectively. In 2010, he received the Ph.D. degree in Signal Processing with a thesis on low-cost navigation systems. In 2009, he spent 5 months at the Mobile Multi-Sensor System research team, University of Calgary, Canada, as a visiting scholar and in 2011 he spent 4 months at the Indian Institute of Science (IISc), Bangalore, India, as a visiting scholar. He is currently a Researcher at KTH

coordinating the KTH Insurance Telematics Lab. He was a recipient of a Best Survey Paper Award by the IEEE Intelligent Transportation Systems Society in 2013.



João G. P. Rodrigues (S'11) received the MSc degree in electrical and computer engineering from the University of Porto, Porto, Portugal, in 2009. He is currently working toward the PhD degree with the University of Porto. He develops his work at the Institute for Telecommunications, and the main topics of his thesis are data gathering and mining in intelligent transportation systems. His main research interests include sensor networks and intelligent transportation systems. He received a Doctoral Scholarship from the Portuguese Foundation for

Science and Technology in 2009.



Peter Händel (S'88-M'94-SM'98) received a PhD degree from Uppsala University, Uppsala, Sweden, in 1993. From 1987 to 1993, he was with Uppsala University. From 1993 to 1997, he was with Ericsson AB, Kista, Sweden. From 1996 to 1997, he was a Visiting Scholar with the Tampere University of Technology, Tampere, Finland. Since 1997, he has been with the KTH Royal Institute of Technology, Stockholm, Sweden, where he is currently a Professor of Signal Processing and the Head of the Department of Signal Processing. From 2000 to

2006, he held an adjunct position at the Swedish Defence Research Agency. He has been a Guest Professor at the Indian Institute of Science (IISc), Bangalore, India, and at the University of Gävle, Sweden. He is a co-founder of Movelo AB. Dr. Händel has served as an associate editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING. He was a recipient of a Best Survey Paper Award by the IEEE Intelligent Transportation Systems Society in 2013.



Ana Aguiar (S'94-M'98-S'02-M'09) received the Electrical and Computer Engineering degree from the University of Porto, Porto, Portugal, in 1998, and the PhD in telecommunication networks from the Technical University of Berlin, Berlin, Germany, in 2008. Since 2009, she has been an Assistant Professor with the Faculty of Engineering, University of Porto. She began her career as an RF Engineer working for cellular operators, and she worked at Fraunhofer Portugal AICOS on service-oriented architectures and wireless technologies applied to

ambient assisted living. She is the author of several papers published and presented in IEEE and ACM journals and conferences, respectively. She contributes to several interdisciplinary projects in the fields of intelligent transportation systems and well-being (stress). Her research interests include wireless networking and mobile sensing systems, specifically vehicular networks, crowd sensing, and machine-to-machine communications. She is a Reviewer for several IEEE and ACM conferences and journals.