

Research Paper

Cite this article: Lopez Fernandez S, Samarasekera ACJ, Feger R, Stelzer A (2025) Velocity vector estimation in automotive radar networks. *International Journal of Microwave and Wireless Technologies*, 1–10. <https://doi.org/10.1017/S175907872400117X>

Received: 15 February 2024

Revised: 28 October 2024

Accepted: 9 November 2024

Keywords:


radar; system applications; standards

Corresponding author:

Sergio Lopez Fernandez;

Email: sergio.lopez_fernandez@jku.at

Velocity vector estimation in automotive radar networks

Sergio Lopez Fernandez , A. Chaminda J. Samarasekera, Reinhard Feger and Andreas Stelzer

Institute for Communications Engineering and RF-Systems, Johannes Kepler University, Linz, Upper Austria 4040, Austria

Abstract

In this work, we develop a method for robust single-cycle measurement velocity vector estimation for automotive radar. Building upon our previous work, we introduce a methodology that leverages spatial diversity for accurate estimation of the velocity vector of targets in the medium to close ranges. We extend our initial conceptual framework, addressing limitations from our first approach and proposing necessary enhancements for real-world applicability. Our improved process excels in target separation, identification, and velocity vector estimation, proving effective across various scenarios and minimizing errors. The system, tested on pedestrians and metal targets, presents a promising avenue for exploring its performance with varying target sizes. Simultaneously, our in-depth study on Doppler-multiplex modulation reveals new relevant constraints, prompting a modulation change for improved response separation. Despite the necessity of increasing module numbers for enhanced performance, our structured approach to target itemization and classification positions our methodology as a valuable framework for future systems, offering a comprehensive solution to diverse challenges in target estimation and classification within the automotive landscape.

Introduction

This work constitutes an extension to an earlier paper presented at the 20th European Radar Conference (EuRAD) and published in its Proceedings [1]. In the quickly advancing world of automotive technology, radar systems have become essential components, transforming the manner in which vehicles perceive and traverse their environments. Currently and next to cameras and Lidar, radar sits at the forefront of advanced driver-assistance systems (ADAS) and autonomous driving [2, 3]. Radar differentiates itself from the rest of sensors due to its better performance in darkness and adverse weather conditions [4], and they currently provide crucial sensing for collision avoidance, adaptive cruise control, and multiple safety features in the close ranges.

Currently, there is a notable trend toward the advancement of high-end single radar modules situated at the front of vehicles for sensing purposes [5, 6]. Nevertheless, we advocate for a shift toward employing simpler radar modules that form a coherent radar network within automotive applications. Radar networks have demonstrated enhancements in detection rates, signal-to-noise ratio (SNR) estimations [7], improved angular resolution [8], and various applications such as the focal point of this work: estimating the velocity vector of targets.

Estimating the velocity vector of targets has conventionally relied on tracking methods [9, 10], involving the maintenance of a target list over successive measurements to infer the trajectory and speed of the targets. In our earlier research [1], we proposed an innovative methodology that leverages the spatial diversity inherent in coherent radar networks to achieve this in a single measurement cycle with independence of module position.

This approach capitalizes on the distinct Doppler velocities observed by individual sensors within the network during each measurement cycle, providing a means to estimate the velocity vector of targets. Various methods for estimating tangential velocity are discussed in [11]. Some studies, such as [12], focus on a single target and address the limitations of Doppler ambiguity. In [13], the authors tackle velocity estimation by using multiple detections of targets with a single module, leveraging clustering techniques. Additionally, distributed system approaches are explored, as seen in [14], where the authors estimate the velocity of a target using bi-static signals. This concept has also been applied to ego-estimation, demonstrating the effective use of data from multiple radar sensors [15–17]. In contrast, our approach utilizes the complete response from a radar network, incorporating both quasi-monostatic and bi-static signals. We introduce a robust method designed for multi-target scenarios, with a particular focus on enhancing safety measures for vulnerable

© The Author(s), 2025. Published by Cambridge University Press in association with The European Microwave Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

road users. In our previous work, the foundational idea was introduced, and we presented preliminary radar measurements as a proof of concept, specifically tailored for close-range automotive applications.

In this subsequent phase of our research, our objective is to further elaborate on and expand the conceptual framework introduced in our previous work. Our aim is to conduct a more comprehensive investigation, delving into the inherent limitations of the proposed method and identifying crucial extensions required for its practical application in real-world scenarios within the automotive domain. We have modified the system by replacing the fixed threshold used in detection with a dynamic constant false alarm rate (CFAR) approach which makes the process more robust in unpredictable scenarios. We have also extended the applicability of the method to multi-target scenarios, improving multipath and side lobe rejection. Additionally, we identified the limitations with the current modulation and proposed a viable alternative. All these changes have been validated through measurement campaigns. By undertaking this extended analysis, we strive to bridge the gap between theoretical advancements and practical feasibility, paving the way for the meaningful integration of our proposed approach into existing systems in the landscape of automotive applications.

The Method section introduces the methodology, placing particular emphasis on redefining bistatic velocities. In the following section, which constitutes the core of this research, we delve into essential extensions required to adapt the method for real-world applications, along with a discussion on Doppler multiplex modulations. The experimental results with the newly implemented modifications are presented in the section after that. Lastly, the last section encapsulates the conclusions drawn from our work.

The method

In this method, we rely on a radar network, and a key feature we make use of is its spatial diversity. This network has various modules placed in a line able to capture Doppler velocities from the targets in front of it. Now, the crucial characteristic is that each module provides different Doppler measurements. These differences are not random – they are influenced by factors like the distance between modules, where they are positioned in the network, and how the targets are moving. The modules in the network function as a group of observers, each offering a unique perspective of the environment and the targets.

Each module's viewpoint gives us a nuanced understanding of the Doppler velocities of our targets. With these varied readings and our knowledge of the network's layout, we get down to estimating the velocity vector of the targets. It is this combination of spatial diversity and network knowledge that helps us unravel the patterns of velocity vectors and make sense of the dynamic movements within the radar network.

The easiest way of visualizing the different Doppler velocities and how they relate to each other on the system is by downsizing the network to just two modules and analyzing the scenario for one moving target. Figure 1 illustrates this case, where \mathbf{v}_T represents the actual velocity vector of the target and \mathbf{v}_1 and \mathbf{v}_2 represent the perceived Doppler components by each module. These perceived velocities are the result after Doppler processing various frequency-modulated continuous-wave (FMCW) chirps of the quasi-monostatic responses. Additionally, Figure 1 includes $\mathbf{v}_{b_{2-1}}$, which is the perceived velocity through the bistatic response when we consider a signal being transmitted by module 2, reflected off the target and received and processed at module 1.




Figure 1. Target moving with velocity vector \mathbf{v}_T . The perceived velocity from the quasi-monostatic response of two stations \mathbf{v}_1 and \mathbf{v}_2 as well as the bistatic response $\mathbf{v}_{b_{2-1}}$ is shown as well as the angle relations between the four.

It can be observed that \mathbf{v}_1 and \mathbf{v}_2 are the projections of \mathbf{v}_T onto the look direction of each module, while $\mathbf{v}_{b_{2-1}}$ is the projection onto the bisector of the bistatic angle ($\beta/2$). Through Doppler processing and estimation we have access to \mathbf{v}_1 , \mathbf{v}_2 and, $\mathbf{v}_{b_{2-1}}$. Also, through digital beamforming (DBF), we can estimate both θ_1 and θ_2 – the direction of arrival for each module.

This was the assumption for [1], however, we have undergone a deeper study of the bistatic signals that reveal some extra details on the perceived velocities [18]. Let us focus for a moment on the bistatic signal depicted in Fig 1. In the diagram module 2 is the transmitter and θ_2 the transmitting angle, module 1 is the receiver and θ_1 the receiver angle. β is the bistatic angle that characterizes the bistatic triangle and is

$$\beta = \theta_2 - \theta_1. \quad (1)$$

To define the relation between the velocity of the target V_T and the actual bistatic component $\mathbf{v}_{b_{2-1}}$, we can start by defining the variation of the range from the transmitter to the target R_2 in the bistatic response as

$$\frac{dR_2}{dt} = V_T \cos(\delta - \beta/2), \quad (2)$$

where δ is the angle between the true velocity direction of the target and the bisector of the bistatic angle β . Analogously, the term $\frac{dR_1}{dt}$ is the projection of the target velocity vector onto the receiver-to-target line of sight (LOS) defined as

$$\frac{dR_1}{dt} = V_T \cos(\delta + \beta/2). \quad (3)$$

The bistatic Doppler is the time rate of change of the total path length normalized by the wavelength λ , being the total path $R_2 + R_1$ the Doppler shift is

$$f_b = \frac{1}{\lambda} \left(\frac{dR_2}{dt} + \frac{dR_1}{dt} \right). \quad (4)$$

We can combine (2), (3), and (4) to express the bistatic shift as

$$f_b = (2V_T/\lambda) \cos \delta \cos(\beta/2) \quad (5)$$

which can also be interpreted as a projection of the velocity vector of the target over the bisector of the bistatic angle. However, it is important to note that the projection appears scaled by the cosine of the bisector of the bistatic angle. The importance of this formulation is due to the fact that the method is based on the scalar projection of a vector (also called the scalar component or the magnitude component). It is the length of the projection of the vector

onto a given line or direction, it is a scalar value. The scalar projection of \mathbf{v}_T for the response n over the LOS of its corresponding module is

$$\|\mathbf{p}_{\mathbf{v}_n}^{\mathbf{v}_T}\| = \frac{|\mathbf{v}_T \cdot \mathbf{v}_n|}{\|\mathbf{v}_n\|} \tag{6}$$

In accordance to Figure 1, we can understand $\|\mathbf{p}_{\mathbf{v}_n}^{\mathbf{v}_T}\|$ as the norm of any of the perceived velocities ($\|\mathbf{v}_n\|$), if we decompose and simplify the scalar product we can express the relation between \mathbf{v}_T and \mathbf{v}_n as

$$V_n^2 = V_{T_x} V_{n_x} + V_{T_y} V_{n_y} \tag{7}$$

where V_{T_x} and V_{T_y} are the Cartesian components of \mathbf{v}_T taking the target as origin as shown in Figure 1. V_{n_x} and V_{n_y} are the components of \mathbf{v}_n and V_n its norm. With this description of the system, we have access to equations that relate the information gathered by each module and response in the network to the actual velocity vector of the target when we understand them as scalar projections of \mathbf{v}_T . This emphasizes the necessity to present the perceived bistatic velocities as a projection of the velocity vector of the target as we derived in (5). With access to at least two of these responses we can form a system of equations following (7) (both quasi-monostatic and bistatic) and solve it for the Cartesian components of \mathbf{v}_T . Most of the time, however, we will be presented with an over-determined system of equations that combines the information gathered at each module:

$$\begin{bmatrix} V_1^2 \\ V_2^2 \\ V_{b_{2-1}}^2 \\ \vdots \\ V_n^2 \end{bmatrix} = V_{T_x} \begin{bmatrix} V_{1_x} \\ V_{2_x} \\ V_{b_{2-1}_x} \\ \vdots \\ V_{n_x} \end{bmatrix} + V_{T_y} \begin{bmatrix} V_{1_y} \\ V_{2_y} \\ V_{b_{2-1}_y} \\ \vdots \\ V_{n_y} \end{bmatrix} \tag{8}$$

Simplified (8) can be written as:

$$\mathbf{b} = \mathbf{A}\mathbf{v}_T \tag{9}$$

where $\mathbf{A} \in \mathbb{R}^{n \times 2}$ and $n > 2$. In those cases, we resort to a least-squares approach that minimizes the sum of the squared errors between the collected data and the expected data based on the model. We look for vector $\tilde{\mathbf{v}}_T$ such that:

$$\tilde{\mathbf{v}}_T = \arg \min_{\mathbf{v}_T} \|\mathbf{A}\mathbf{v}_T - \mathbf{b}\|_2^2 \tag{10}$$

The original work [1] operated with a network of four different modules. Delving further into the intricacies of the parameter extraction step, we used to start with a set of 16 (4 quasi-monostatic responses and 12 bistatic responses) distinct 3D matrices capturing range, Doppler, and angle information. To streamline computational efficiency, the initial move involved collapsing the angle dimension by identifying the maximum value in each matrix. Subsequently, a fixed range-dependent threshold was applied to the matrix, with a more stringent criterion for closer ranges. This threshold was designed taking into account the measurement chamber where the experiments were conducted. The identified detections initially resided solely in the range-Doppler domain, prompting an extension into the angle domain. In this extension, the threshold was determined by selecting the stricter value between the range-dependent threshold and the expected level for secondary lobes resulting from windowing.

This process culminated in a 3D cloud of detections, where we applied clustering to identify the moving target in the scene, and for this, a modified iteration of DBSCAN – specifically EDBSCAN

[19] – was employed. Through this clustering approach, we not only identified the target of interest present in the scene but also extracted essential parameters, including range, radial Doppler, and receiving angle per module. As it will be explained in Section III, bi-static responses will require an additional layer of processing to resolve the geometry of the system. This multi-step methodology not only enhanced computational efficiency but also ensured a robust extraction of pertinent information.

Required extensions for practical applicability

The earlier-discussed method was initially designed as a proof of concept within controlled scenarios. In essence, our goal is to identify its limitations and create the necessary tools to maintain its performance when dealing with real-world targets and environments.

Thresholding

The initial challenge we tackle is the range-dependent threshold. This approach thrives when armed with a full understanding of our operating ranges, the specific target(s) in focus, and the surrounding environment. By fixing the threshold values, we secure the detection of targets within our radar's range, enabling accurate estimation. Our ultimate objective is to estimate the velocity vector of vulnerable road users in close to medium ranges. However, since different targets exhibit diverse radar cross-sections (RCSs), relying on pre-established threshold values is not feasible. Moreover, the environment is inherently unpredictable, necessitating adaptability in order to capture all potential targets without any oversight.

We have resorted to CFAR detection, a technique specifically used to keep false alarms at a suitably low rate in priori unknown, time varying and spatially non-homogeneous environments [20]. The basic function of CFAR is to determine the threshold above which any peak can be considered a target. To determine this threshold dynamically so it adapts to the background against which the targets are to be detected we are using a cell-averaging (CA) approach. For this work, we employ a two-dimensional version of the algorithm to enable its application to the angle-flattened range-Doppler matrices generated after the DBF step outlined in Section II. A less computationally intensive alternative could be a one-dimensional application along the range for each Doppler bin. However, when using Doppler-division multiple access (DDMA), the probability of false alarm calculation is different from the traditional one; to exploit the best use of DDMA and increase the probability of detection we use 2D-CFAR [21]. The second axis improves the detection of small targets by considering both dimensions when setting up the threshold [22, 23]. The result is a list of detections that can be subsequently be clustered to later identify each target.

Target association

The method in Section II is straightforward under the assumption of a singular target scenario, simplifying the process by consolidating all relevant information into a cohesive whole. This approach is advantageous in its simplicity, requiring minimal additional processing, we know all this information from all different modules and responses belongs to this one target. However, as we transition to addressing the complexities of the real world, where multiple targets may coexist, a shift in strategy becomes imperative. The dynamic nature of multiple targets necessitates a more intricate

layer of processing to disentangle and appropriately account for the diverse information streams.

Let us consider targets after detection to be characterized by three distinct parameters: range, receiving angle, and radial velocity. In the optimal scenario, each target is detected by every module and observed in every response (bistatic and quasi-monostatic), this introduces two additional identifiers for each target (in which module and in which response it was detected). Therefore, following this initial processing step (a more detailed explanation will be given in Section III.D), we possess a set of detections, F_n , for each module and response in the system where $1 \leq n \leq N$. Here N is the total amount of responses in the system (quasi-monostatic and bistatic combined for every module) that will be processed and formulated as in (6)

$$F_n = \{d_1, d_2, \dots, d_K\} \quad (11)$$

where

$$d_k = [R_k, \nu_k, \theta_k] \quad (12)$$

and $k = 1, 2, \dots, K$ the total number of detections within a response. The objective is to combine all N sets to establish the associations between each detection present and the various targets within the scene. This association is pivotal for the subsequent velocity vector estimation procedure.

Before that, the first step involves aligning each detection onto a common plane, facilitating comparison and association. To achieve this, we rely on information regarding the spatial arrangement of each module in the network. In our context, tailored for automotive applications, we make the assumption that all modules align linearly along the front bumper of the car. Utilizing available data on module distances and their relative order, we are able to position each detection onto the range/cross-range plane. This is achieved by transforming the polar pair of range (R_k) and angle (θ_k) values into Cartesian coordinates.

The quasi-monostatic responses require adjustment to a common reference, and the correction depends on the position of the receiving module on the network. This entails a lateral shift of all points in Cartesian coordinates. Deeper processing is applied to bistatic responses, where, in addition to potential lateral shifts due to the module positions, each detected range undergoes correction by solving the bistatic triangle

$$R_R = \frac{(R_T + R_R)^2 - L^2}{2(R_T + R_R + L \sin \theta_R)} \quad (13)$$

where R_R is the actual receiver range we want to estimate. L is the baseline or distance between the two modules and θ_R is the angle of arrival at the receiver. After corrections, we can fuse all N sets together resulting in a point cloud of the detections on the Cartesian plane. Subsequently, a second layer of clustering using traditional DBSCAN is employed over it. One notable advantage is that we are not required to pre-define the number of clusters in the dataset which makes the process more versatile and robust.

It is worth mentioning that in this process, we deliberately omit the use of velocity information. This decision is intentional. We don't anticipate detections landing in the same position on the Cartesian plane to necessarily share the same radial velocity. It is precisely this divergence that the method relies on. Introducing that extra dimension would only complicate the clustering procedure, requiring a more lenient metric for velocity and potentially grouping detections with a broader range of velocity values.

This is in contrast to the modified DBSCAN (EDBSCAN [19]) used previously in the process to create the different F_n sets. Our

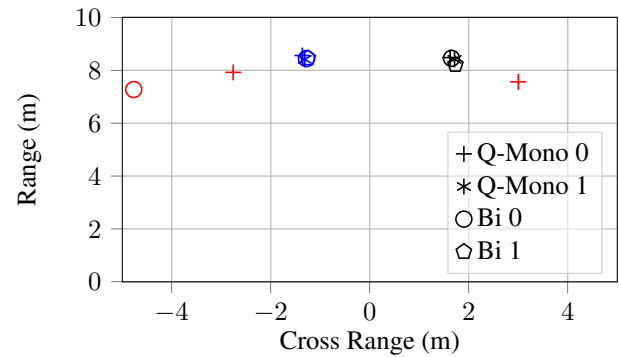


Figure 2. Second layer of clustering for detections of two distinct targets at 8 m. DBSCAN is able to group the detections from two different modules (0 and 1) and their two different responses (quasi-monostatic and bistatic) into two clusters (blue and black). Three different detections are classified as noise (in red).

main focus initially is to avoid the influences that changes in the densities of the search area can have over DBSCAN. We are looking for targets and trying to avoid them getting clustered with clutter and noise. That version specifically uses the three different metrics (range, angle, and Doppler information) and power to group all the points that surpass the CFAR threshold. It adds an extended density-based spatial clustering adding a new size metric to the process making it work well in both dense and sparse search areas and specifically in high dynamic range environments. We want to make sure that multiple detections of the same target (for example both arms and torso of a pedestrian) are interpreted as just one in order to add them as a single entry in its set. Now we are looking to associate each detection in the fused F_n sets and thanks to the initial clustering process we work already with a sparse point cloud that has already been “filtered” and so, the original DBSCAN algorithm proves to be sufficient for our needs.

This additional layer of clustering serves a dual purpose in the process, aiding in the discrimination of false positives within the detections. The clustering process will identify detections that do not belong to actual targets in the scene and label them as noise. Upon retracing the data, these detections typically align with excessively high secondary lobes that manage to surpass the CFAR threshold. When DBSCAN assesses groupings, these detections become isolated from the rest, enabling the estimation process to proceed seamlessly. This behavior has also been noted with multi-path components, which eventually undergo filtration through the clustering procedure as well. In Figure 2, the depicted example illustrates the outcome of the second layer of clustering, grouping detections from two distinct modules and encompassing all responses within the network (quasi-monostatic and bistatic). Utilizing DBSCAN, this process accurately discerns two real targets, denoted by distinct colors (blue and black), within the scene. Notably, DBSCAN adeptly discriminates three detections of ghost targets surpassing the CFAR threshold, which have been highlighted in red.

Limitations of Doppler modulation in widely separated MIMO radar networks

The fundamental principle behind the estimation process lies in accessing the various measurements of radial velocity for a target available in a radar network. We work with a four-module network whose detailed specifications will be outlined in the next section.




Figure 3. Doppler multiplex scheme for four modules. The Doppler spectrum is divided into 16 distinct slots, with a strategic allocation of 12 slots dedicated to the three individual transmitters within each module.




Figure 4. To demultiplex the MIMO channels, we shift the scheme in Figure 3 along the Doppler, aligning each transmitter at the spectrum’s center. This ensures complete transmitter overlap only in the targeted region – column 0 (if a target with velocity zero is assumed). The image represents a single range-Doppler map adjusted 12 different times, one per transmitter in the system. The bottom line shows the expected overlap “weights” for each slot.

Each module features three transmitters and four receivers, operating in a multiple-input multiple-output (MIMO) configuration [24, 25], totaling 192 channels. To separate each channel, we have utilized a Doppler-multiplex modulation, capitalizing on the Doppler dimension of signals by shifting each transmitter in the Doppler domain. Unlike techniques such as code division multiplex, this modulation type eliminates the need for a dedicated multiplexer on the receiver side. Essentially, we work with virtual transmitters along the Doppler dimension and only need to account for these shifts.

This modulation scheme works well in controlled measurement scenarios and predominantly for longer ranges. However, there are limitations in the modulation that require adjustments to our setup and the modulation usage.

So, initially, we utilized a 16-phase shifting key (PSK) scheme, of which we employed 12 different phase increments from the constellation to separate the 12 transmitters. The shifts along the Doppler dimension of the different transmitters are illustrated in Figure 3. The figure illustrates a cross-section of a range-Doppler map where the Doppler domain is divided up in 16 different slots, 12 of which get assigned to each transmitter in the system. The distribution is structured in a way that, upon shifting each transmitter to the center of the spectrum (rolling the range-Doppler map along its Doppler dimension) and essentially “demodulating” by eliminating the phase shift we introduced, a peak is achieved in the center of the spectrum, where there is complete overlap of transmitters. In every other slot of the 16-PSK, there is consistently a gap. This design ensures that we can seek the maximum in the superposition of all rolled versions of the range-Doppler maps to identify the actual target in the whole Doppler domain without detecting ambiguities. Upon completing the requisite shifts, the final distribution for a zero-velocity target is visually represented in Figure 4. Notably, the most pronounced overlap is discernible only within slot 0 of the Doppler domain. This effect is highlighted in Figure 4 by showing the “weights” of each slot in the Doppler domain.

However, a significant limitation arises with this modulation for widely separated MIMO radar modules, and it directly aligns with the fundamental principle of our method. To achieve complete overlap of every transmitter across all modules, we must assume uniform Doppler shift measurements due to target movement across them all. However, reality does not conform to this assumption. When the estimated radial velocity varies from module to module, the spectrum peak for a target shifts to a different Doppler bin in each module, introducing an additional shift beyond the intended PSK that we cannot compensate for. Consequently, the assumption of finding the maximum peak after demultiplexing to identify the target does not hold exactly. Moreover, this behavior is more prominent in our specific use case when the distance to the interest targets approaches the size of our radar network. This can introduce errors in the first layer of clustering, where we estimate parameters for each target (radial Doppler, receiving angle, and range). These errors in estimated velocity can then propagate through the subsequent steps of the process.

Furthermore, a second limitation becomes pronounced, particularly due to the imperative need to process and employ the distinct responses within a module separately. As detailed in Section II, it is essential to treat each perceived velocity separately, given its correlation with the target’s velocity vector. Moreover, to achieve successful target association, distinguishing between different responses is crucial for implementing the appropriate corrections in the second layer of clustering.

Consider the scenario where we aim to process only the monostatic information from module 0. Referring to Figure 4, this entails working with the data received at module 0 that was transmitted from module 0. This information corresponds to the first three rows, aligning with the shift of the three transmitters within module 0. While we anticipate the spectrum’s maximum to be centered after demultiplexing, the presence of additional overlaps from other stations like in slots 12, 13, 4, and 6 (highlighted in dashed lines) introduces a level of inaccuracy to this search. Notably, this effect is not consistently significant, as it occasionally coincides with the first limitation. In some instances, the fact that different modules estimate different Doppler shifts prevents overlaps from non-interest modules. However, these secondary overlaps gain significance in close ranges or when targets exhibit substantial angle dependencies in their RCSs. The positioning of a target in space can lead to a much stronger reflection for one module compared to others, causing peaks on the spectrum to be more pronounced at unexpected Doppler frequencies. This, ultimately, compromises the assumption of maximum overlap. Similar to the first limitation, this primarily impacts the initial layer of clustering, where we set the parameters for each target and response. These parameters are crucial for subsequent tasks, including the association between modules and later velocity estimation.

One way to address these two limitations is to increase the unambiguous velocity and the separations between transmitters in the Doppler domain by modifying the distribution. However, due to hardware constraints, we have opted for a strategic approach: downsizing the network. Considering the current landscape of automotive systems [26–29], it is evident that the limited space within automobiles poses a challenge for accommodating multiple radar sensors. Employing two radars serves as a solution, mitigating the effects mentioned and optimizing space utilization at the front of the car. Additionally, to rigorously test the robustness of our method, we curtailed the number of transmitters to two per module, creating a more relaxed Doppler spectrum. While retaining the 16-PSK constellation, we assigned only four slots. This decision




Figure 5. The first row presents the Doppler multiplex scheme for two modules. It divides the spectrum into 16 distinct slots, strategically allocating 4 slots for each of the 2 individual transmitters within every module. We have highlighted the specific areas of overlap in individual responses, emphasizing that these overlaps exclusively occur in the center of the spectrum. Additionally, this central area is surrounded and protected by guard cells.

serves a dual purpose: it allows for the existence of guard cells around the center of the spectrum, and it ensures that when examining quasi-monostatic and bistatic responses separately, the only overlap of transmitters in the Doppler domain occurs in the region of interest. Figure 5 illustrates the final distribution. Here, we can discern that by focusing on either the initial two rows or the concluding two, the overlap of transmitter slots is limited to the central region of the Doppler spectrum. Furthermore, we incorporate a single slot on each side, serving as a guard cell to enhance signal integrity.

The process

Hence, the process for estimating the velocity vector in a single measurement cycle for multiple targets within a radar network unfolds as follows: initially, each module down-converts the received signals, followed by range-Doppler processing. Given the Doppler modulation used, this step allows us to separate information from different transmitters. Once all received signals are identified, we employ DBF to generate a 3D matrix for each response and module. These matrices encompass range, Doppler, and angle of arrival dimensions. In our specific case, with the network down-scaled to address the limitations of the Doppler multiplex modulation, we have four such matrices – two for quasi-monostatic information and two for bistatic.

Next, we simplify the angle dimension in each matrix by taking the maximum value across it, resulting in 2D range-Doppler matrices. We then apply 2D CA-CFAR, yielding a set of detections in the range-Doppler domain. To extend these detections to the angular domain, we compare the threshold obtained by CA-CFAR with the expected level of secondary lobes after beamforming windowing, using the stricter of the two. For each detection in the range-Doppler domain, we select every point along its angular dimension that exceeds the chosen threshold. At this stage, EDBSCAN [19] is implemented to group detections, ascertain the number of targets, and identify the center of each. These detections are grouped per module and response in sets before association. Here begins the preprocessing before grouping, as we position every detection in the Cartesian plane. To ensure the integrity of the data for the final estimation, as explained in Section II, we retain the original parameters untouched, while also applying necessary corrections for successful association, storing them separately. Adjustment of bistatic signals is crucial to account for the module disposition, ensuring correct receiving angles and ranges. Furthermore, every detection undergoes a lateral shift to align with the baseline and its module's position.

Finally, after fusing each corrected set, a second round of DBSCAN is executed resulting in a number of sets equal to the number of targets in the scene grouping all relevant detections.

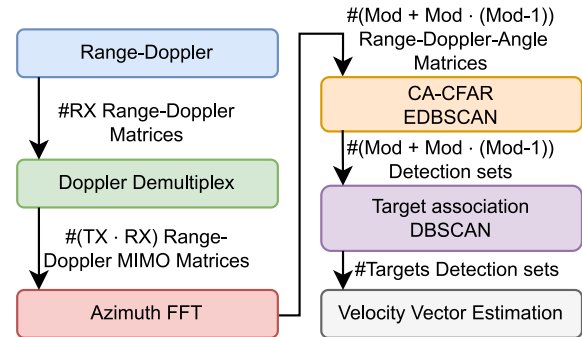


Figure 6. Streamlined flowchart depicting the entire process leading to the estimation of velocity vectors for each target within the scene. The diagram delineates the format and quantity of data available between each step.

Once we establish which detections belong to each target, we proceed with the final estimation, employing unmodified parameters for each separately through a least-squares approach. The process described ensures that the estimation can be indeed performed separately for each measurement cycle and is achieved independently of the module positions. Figure 6 shows a simplified diagram of the whole process culminating in the velocity vector estimation for each target in the scene.

To end, it is interesting to note how the process will work even when modules of the network miss detections of the targets: As long as we can cluster enough detections (no matter the origin) in the second layer of DBSCAN for association, a velocity vector will be calculated. Furthermore, by concentrating our computational power on actual targets, this second layer provides an avenue for introducing specific layers of filtering, enabling a more tailored treatment of the detections as needed. This implies that upon completion of the process, we will possess a meticulously curated list of targets, along with their estimated radial velocities and the system's estimation of their actual velocity vectors. This comprehensive information will empower future target management efforts.

Experimental results

Several measurement campaigns were conducted in order to test the proposed process. The primary goal was to ascertain the accuracy and effectiveness of the association process, assess the performance of the newly adapted Doppler modulation, and evaluate the network's capabilities at extended ranges. It was important to check to the influence of downscaling the network to just two modules in comparison to the four used in our previous work. Numerous measurements were conducted across diverse targets and scenarios, and this paper will delve into the findings from three specific instances:

- (i) Indoor measurement with two moving targets: a metal pole and a person.
- (ii) Outdoor measurement with two moving targets: two people.
- (iii) Outdoor measurement with one person walking at higher ranges.

Throughout all the measurement campaigns, the radar network employed three distinct modules, with one designated as




Figure 7. This module's hardware snapshot reveals on the left the backside of a radar node with the FMCW chirp-receiving waveguide (a), data-transmitting Ethernet connector (b), synchronization signal reception (c) beneath the PCB, and its power connector (d). On the right side, the transmitting and receiving arrays can be seen.

the network's master, connected to the other two in a star configuration. The central node does not engage in signal transmission or reception but plays a crucial role in ensuring coherent network operation. It generates FMCW chirps within the 38–40.5 GHz frequency range, transmitting them through coaxial cables to the remaining modules. An additional connection facilitates the transmission of three signals: a reference clock, synchronization signal for the analog-to-digital converter (ADC), and ramp synchronization signal. This setup guarantees simultaneous transmission, synchronized clocks, and aligned ADCs for both modules.

Prior to transmission, the modules up convert the FMCW chirps, the start frequency is 76.5 GHz, with a bandwidth of 900 MHz (range resolution of 17 cm) and a chirp duration of 32 μ s. Processing-wise, we utilized a total of 512 intermediate-frequency (IF) samples, analyzing 256 chirps for Doppler estimation. As detailed in Section III, each module employs two transmitters for transmission while employing the full four-element uniform linear array (ULA) for reception. [Figure 7](#) shows the hardware of a single module within the network.

For this application, the different responses are processed individually at each module in a non-coherent manner, without treating the network as a whole. While this approach means we cannot leverage certain radar network benefits, such as increased angular resolution, it allows us to relax setup requirements and reduce the need for additional calibration efforts. The signal processing method follows the approach detailed in [30]. The results presented here are likely optimal due to perfect synchronization but could be achieved with a non-coherent setup.

During indoor measurements, the network was arranged on a cart and the measurements performed inside an anechoic chamber. This setup served to validate the hardware's proper functioning and the modulation's accuracy before transitioning to outdoor measurements. For the latter, the network was mounted on a car, with radar nodes strategically placed atop the front bumper, spanning the vehicle's width. Refer to [Figure 8](#) for the arrangement of all three modules, noting that Module 0, for future reference, is positioned on the right from the driver's perspective. In both cases, the positioning limits the setup to azimuth estimation only for the direction of arrival. We chose this limitation due to the increased complexity of the installation, although all the procedures in Section II could easily be extended to a third dimension. We aim to maximize




Figure 8. Radar network installed on the car for outdoor measurements. The distance between the two modules is 1.01 m.




Figure 9. In our initial measurement scenario, we introduced two targets: a pole affixed to a rail, facilitating perpendicular movement, and a person traversing parallel to the network. The image provides the perspective from Module 1.

the distance between the modules to increase spatial diversity and the total aperture of the network. A larger aperture correlates with greater differences in radial velocities, thereby enabling successful velocity vector estimation at longer ranges. For the standard 1.5 m width of a car, we have achieved acceptable results up to 25 m.

Indoor measurement

The first measurement campaign followed the reduction in size of the radar network. Having previously assessed our method's performance with the full capabilities of four transmitting and receiving modules, we aimed to draw a comparison with the efficacy of using only two. Additionally, our objective was to scrutinize the association process within a more controlled setting. To achieve this, we conducted measurements featuring both static and moving targets within an anechoic chamber. The showcased measurement scenario, depicted in [Figure 9](#), involved affixing a 150 cm metal pole to a rail, enabling perpendicular movement to the network at 1 m/s. Concurrently, an individual traversed parallel to the network, intersecting the trajectory of the metal pole. The image shows the targets before starting the data collection. The distance between the modules is 1.01 m.

In conducting the measurements, we adhered to the process outlined in Section III, delving into the analysis of the two quasi-monostatic responses within the network and the two corresponding bistatic responses. The information extracted from the bistatic responses may appear somewhat redundant, given the near-equivalence of the signal paths in both directions. Nevertheless, this redundancy proves advantageous, fortifying the estimation process for each measurement cycle.

In [Figure 10](#), we present the culmination of our estimation efforts. By processing and fusing information from the four distinct responses, we successfully detected the targets with an SNR of

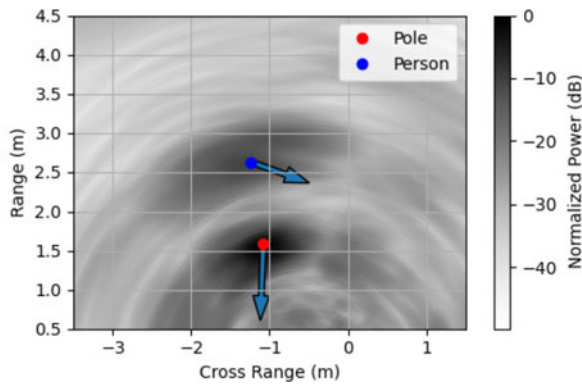


Figure 10. Fused information from all four responses [30], the image showcases detections for both targets alongside the corresponding estimated velocity vectors. Notably, the vector lengths align with the respective estimated modules. The power values are normalized to the scene's maximum for enhanced clarity.




Figure 11. In the first outdoor scenario, two pedestrians move in opposing directions, walking perpendicular to the radars.

30 dB and derived their velocity vectors. The SNR was determined by comparing the power level of the targets to the average noise level in a target-free selected area. Notably, in this scenario, precise ground truth data are available solely for the pole, obtained from its predetermined movement along the rail at 1 m/s, post a brief acceleration period. The estimated velocity for the pole exhibits a root-mean-square error of 0.032 m/s, a commendable outcome aligning with our past estimations and deemed suitable for the given scenario. The person's speed, though slightly slower than average walking speeds estimated at nearly half that of the pole, aptly mirrors both the observed circumstances during the measurements and the direction and module of their movement. We specifically selected this frame due to the proximity of the targets, aiming to verify the accurate association of the estimated parameters.

Outdoor measurements

In the subsequent two sets of measurements, the network was affixed to the front bumper of a car, as illustrated in Figure 8. Following the approach employed in indoor measurements, we endeavored to uphold a consistent setup by maintaining the inter-module distance at 1.01 m. The primary focus in this phase was to assess the network's performance with two pedestrians at extended distances. The latter part of the outdoor measurements is specifically dedicated to higher ranges. Throughout these measurements, the radar parameters remain constant, and we will present selected frames from the acquired data to illuminate our findings.

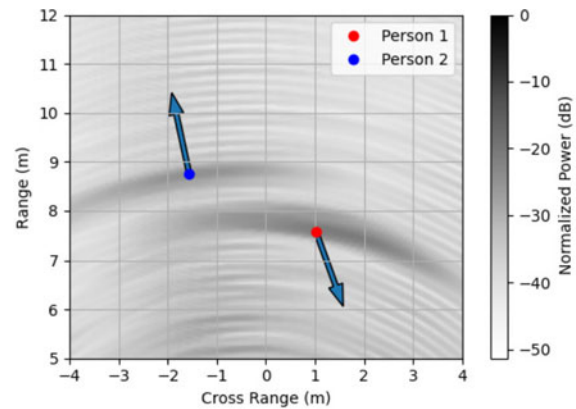


Figure 12. The image consolidates data from all four responses, displaying detections for both targets with their corresponding velocity vectors. A slight tilt is noticeable due to network positioning, natural walking motions, and minor performance degradation at higher ranges. Power values are normalized for improved clarity, relative to the scene's maximum.

In the initial scenario, two pedestrians moved anti-parallel to each other, quasi-perpendicular to the network, in opposite directions. Refer to Figure 11 for their positions relative to the measurement car. In the presented frame, the left pedestrian is slightly farther, moving away from the network, while the right pedestrian is approaching it. We processed the data akin to the initial measurements – handling and fusing the four responses, detecting targets, extracting parameters, and estimating velocity vectors. The outcomes are depicted in Figure 12 which result in an estimation with an SNR of 20 dB.

The estimated vectors align closely with the anticipated outcomes for both norm and direction. The estimated velocities for the targets – 1.694 m/s and 1.745 m/s, respectively – approach typical walking speeds [31, 32]. However, there is a slight tilt observed in the estimated directions, influenced partly by the actual positioning of the network and the natural tilting motions of the human body during walking. The process encounters some challenges as distances increase, leading to a reduction in the disparities between the multiple estimated radial velocities.

We aimed to assess the algorithm's proficiency in determining accurate velocity vectors over extended distances. Since our primary focus remains on the 5–15 m range, typical for short range automotive scenarios, we sought to explore the algorithm's capabilities in those longer ranges. To conduct this evaluation, we affixed the network to the vehicle and navigated it along a narrow road, where a pedestrian walked at various distances. Figure 13 displays selected frames capturing the person in different positions and orientations during the walking sequence.

During the analysis of these measurements, the system operates in proximity to both its limits and optimal utilization. As the distance increases, the distinctions between the radial velocities estimated by each response become less pronounced. This poses a challenge in leveraging spatial diversity to accurately estimate the underlying information about the true velocity vector. Nevertheless, within the desired ranges, as illustrated in Figure 14, the process continues to effectively function for this particular target and environment. The radar system correctly detects the pedestrian in both cases with an SNR of 25 dB. In both instances of movement, the algorithm successfully estimates both the correct direction and magnitude of the velocity of 1.556 m/s for the first frame and 1.199 m/s for the second. For this scenario, the




Figure 13. The images depict frames extracted from a video capturing measurements from inside the car. In the first frame, the person is observed walking diagonally toward the car, while in the second frame, they exhibit a slight sideways movement, this time moving away from the car.




Figure 14. The left plot corresponds to the top frame, displaying the estimated velocity vector for a target at 10 m. On the right, the estimation is shown for the second frame, where the target moves away from the network at 12 m. The arrow lengths represent the estimated magnitudes of the vectors. Power values are normalized for improved clarity, relative to the scene's maximum.

results have been compared to measurements made with a Lidar system installed on the car. The estimations from our system closely follows the results provided by the ground truth system, demonstrating high accuracy.

Conclusion

Our recently developed process has proven highly effective within our specified use case. We have been able to separate, identify, and estimate the velocity vector of multiple targets at different distances from each other and the radar network. While its robust association across various responses and modules plays a crucial role in managing multi target scenarios, it also alleviates errors related to both thresholding and multi-path effects. It is able to both provide the velocity estimation and ease the weight on the detection process. Although we have currently tested pedestrians and metal targets, exploring how the system handles increasing size differences among targets is an intriguing avenue for future investigation.

We have conducted an in-depth study on Doppler-multiplex modulation, uncovering new relevant limitations. This analysis has shed light on unforeseen challenges associated with this modulation scheme, particularly in situations where the goal is to effectively distinguish responses within our radar systems. We have proposed a new modulation as a solution, effectively overcoming previous said limitations. It is evident that enhancing system performance involves increasing the number of modules, and we may explore alternative modulation schemes to facilitate response separation and enhance velocity estimation within the limits on current ADAS.

Nonetheless, the most significant advantage lies in the systematic organization of target itemization and response classification within the process's structure, positioning it as a valuable framework for future systems. This structured and comprehensive approach addresses diverse challenges in target estimation and classification, making it a promising foundation for future advancements.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.23919/EuRAD58043.2023.10289157>.

Acknowledgements. This work has been supported by the COMET-K2 Center of the Linz Center of Mechatronics (LCM) funded by the Austrian federal government and the federal state of Upper Austria.

Competing interests. The author(s) declare none.

References

1. Fernández SL, Samarasekera ACJ, Feger R, Stelzer A and Hanumegowda A (2023) Single cycle velocity vector estimation using a full-coherent MIMO radar network. In *20th European Radar Conference*, Berlin, Germany. IEEE, pp. 165–168.
2. Jones W (2001) Keeping cars from crashing. *IEEE Spectrum* 38(9), 40–45.
3. Meinel HH (2014) Evolving automotive radar—From the very beginnings into the future. In *The 8th European Conference on Antennas and Propagation*, The Hague, Netherlands. IEEE, pp. 3107–3114.
4. Roos F, Bechter J, Knill C, Schweizer B and Waldschmidt C (2019) Radar sensors for autonomous driving: Modulation schemes and interference mitigation. *IEEE Microwave Magazine* 20(9), 58–72.
5. Li M, Stolz M, Feng Z, Kunert M, Henze R and Küçükay F (2018) An adaptive 3D grid-based clustering algorithm for automotive high resolution radar sensor. In *IEEE International Conference on Vehicular Electronics and Safety*, Madrid, Spain. IEEE, pp. 1–7.
6. Ali M, Chen L, Tanyer G, Tang J, Graff A and Roy A (2023) Long range high resolution imaging radar with digital code modulation (DCM) and sparse array. In *20th European Radar Conference*, Berlin, Germany. IEEE, pp. 81–84.
7. Gottinger M, Gulden P and Vossiek M (2021) Coherent signal processing for loosely coupled bistatic radar. *IEEE Transactions on Aerospace and Electronic Systems* 57(3), 1855–1871.
8. Nguyen MQ, Feger R, Bechter J, Pichler-Scheder M and Stelzer A (2021) High angular resolution digital beamforming method for coherent FMCW MIMO radar networks. In *EEE MTT-S International Microwave Symposium*, Atlanta, GA, USA. IEEE, pp. 756–759.
9. Liu J, Chu M and Reich JE (2007) Multitarget tracking in distributed sensor networks. *IEEE Signal Processing Magazine* 24(3), 36–46.
10. Yan B, Paolini E, Xu L and Lu H (2022) A target detection and tracking method for multiple radar systems. *IEEE Transactions on Geoscience and Remote Sensing* 60, 1–21.
11. Doerry A, Milesosky B and Bickel D (2002) Tangential velocity measurement using interferometric MTI radar. *US Patent*.
12. Lee S-G, Jung J and Kim S-C (2022) Enhanced velocity vector estimation using distributed radar system. In *IEEE VTS Asia Pacific Wireless Communications Symposium*, Seoul, Republic of Korea. IEEE, pp. 75–79.

13. **Singh N, Sil D and Sharma A** (2021) Parallelized instantaneous velocity and heading estimation of objects using single imaging radar. In *IEEE Radar Conference*, Atlanta, GA, USA. IEEE, pp. 1–6.
14. **Wang P, Li H and Himed B** (2012) Target velocity estimation and CRB with distributed MIMO radar in non-homogeneous AR-modeled disturbances. In *13th International Radar Symposium*, Warsaw, Poland. IEEE, pp. 109–112.
15. **Kellner D, Barjenbruch M, Klappstein J, Dickmann J and Dietmayer K** (2014) Instantaneous ego-motion estimation using multiple Doppler radars. In *IEEE International Conference on Robotics and Automation*, Hong Kong, China. IEEE, pp. 1592–1597.
16. **Zhu S, Yarovoy A and Fioranelli F** (2023) DeepEgo: Deep instantaneous ego-motion estimation using automotive radar. *IEEE Transactions on Radar Systems* 1, 166–180.
17. **Grebner T, Riekenbrauck R and Waldschmidt C** (2024) Simultaneous localization and mapping (SLAM) for synthetic aperture radar (SAR) processing in the field of autonomous driving. *IEEE Transactions on Radar Systems* 2, 47–66.
18. **Willis N** (2005) *Bistatic Radar*. Raleigh, NC, US: Institution of Engineering and Technology.
19. **Wagner T, Feger R and Stelzer A** (2015) Modification of DBSCAN and application to range/Doppler/DoA measurements for pedestrian recognition with an automotive radar system. In *European Radar Conference*, Paris, France. IEEE, pp. 269–272.
20. **Farina A and Studer FA** (1987) A review of CFAR detection techniques in radar systems. In *Optimised Radar Processors*, pp. 179–185.
21. **Nguyen MQ, Feger R, Wagner T and Stelzer A** (2023) Analysis of 2D CA-CFAR for DDMA FMCW MIMO radar. In *20th European Radar Conference*, Berlin, Germany. IEEE, pp. 423–426.
22. **Yuan Y, Li W, Sun Z, Zhang Y and Xiang H** (2019) Two-dimensional FFT and two-dimensional CA-CFAR based on ZYNQ. *The Journal of Engineering* 2019(20), 6483–6486.
23. **Wang W, Wang R, Jiang R, Yang H and Wang X** (2019) Modified reference window for two-dimensional CFAR in radar target detection. *The Journal of Engineering* 2019(21), 7924–7927.
24. **Touati N, Sturm C, Imran M, Vanaev A, Kohler M, Krupinski K, Malik W and Lübbert U** (2021) High angle resolution automotive radar based on simultaneous 12Tx Doppler-multiplex MIMO. In *17th European Radar Conference*, Utrecht, Netherlands. IEEE, pp. 386–389.
25. **Sturm C, Sit YL, Li G, Vayghan HA and Lübbert U** (2018) Automotive fast-chirp MIMO radar with simultaneous transmission in a Doppler-multiplex. In *19th International Radar Symposium*, Bonn, Germany. IEEE, pp. 1–6.
26. **Moore M and Lu B** (2011) Autonomous vehicles for personal transport: A technology assessment. *SSRN Electronic Journal*. Available at <https://ssrn.com/abstract=1865047>.
27. **Rajasekhar MV and Jaswal AK** (2015) Autonomous vehicles: The future of automobiles. In *IEEE International Transportation Electrification Conference*, Chennai, India. IEEE, pp. 1–6.
28. **Patole SM, Torlak M, Wang D and Ali M** (2017) Automotive radars: A review of signal processing techniques. *IEEE Signal Processing Magazine* 34(2), 22–35.
29. **Xing Y, Lv C, Chen L, Wang H, Wang H, Cao D, Velenis E and Wang F-Y** (2018) Advances in vision-based lane detection: Algorithms, integration, assessment, and perspectives on acp-based parallel vision. *IEEE/CAA Journal of Automatica Sinica* 5(3), 645–661.
30. **Fernández SL, Samarasekera ACJ, Feger R and Stelzer A** (2022) Measurement-based analysis of a non-coherent MIMO radar network for automotive applications. In *19th European Radar Conference*, Milan, Italy. IEEE, pp. 1–4.
31. **Alves F, Cruz S, Ribeiro A, Bastos Silva A, Martins J and Cunha I** (2020) Walkability index for elderly health: A proposal. *Sustainability* 12(18), 7360.
32. **IEEE VT, ITS** (2022) Literature review on kinematic properties of road users for use on safety-related models for automated driving systems. IEEE, pp. 1–35.



Sergio López Fernández received the Master in Telecommunications Engineering (MEng) from the University of Oviedo, Oviedo, Spain in 2019. The same year he joined the Institute for Communications Engineering and RF-Systems at the Johannes Kepler University Linz, Linz, Austria, where he is currently pursuing his Dr. Techn. (Ph.D.) degree. His current research interests include FMCW radar sensors, automotive radar, and MIMO radar networks.



A. Chaminda J. Samarasekera earned his Master of Science degree from Blekinge Institute of Technology in Sweden in 2014. From 2014 to 2018, he conducted research at Kyung Hee University in South Korea. Since 2018, he has been pursuing his Ph.D. at Johannes Kepler University Linz in Austria, focusing on FMCW radar networks, radar network synchronization, phase noise, and non-linearity issues in radar networks.



Reinhard Feger was born in Kufstein, Austria, in 1980. He received the Dipl.-Ing. (M.Sc.) degree in mechatronics and the Dr. Techn. (Ph.D.) degree in mechatronics from Johannes Kepler University Linz, Linz, Austria, in 2005 and 2010, respectively. In 2005, he joined the Institute for Communications and Information Engineering, Johannes Kepler University Linz, Linz, Austria, as a Research Assistant. In 2007, he became a member of the Christian Doppler Laboratory for Integrated Radar Sensors, Johannes Kepler University Linz. He is currently an Associate Professor with the Institute for Communications Engineering and RF-Systems, Johannes Kepler University Linz. His research interests include radar signal processing, as well as radar system design for industrial and automotive radar sensors. Dr. Feger was recipient of the 2011 Microwave Prize and the 2011 German Microwave Conference Best Paper Award. In 2012, he received the Best Measurement Paper Prize at the European Conference on Antennas and Propagation.



Andreas Stelzer received the Diploma Engineer degree in electrical engineering from the Technical University of Vienna, Austria, in 1994, and the Dr. Techn. degree (Ph.D.) in mechatronics from the Johannes Kepler University (JKU) Linz, Austria, in 2000. In 2003, he became Associate Professor at Johannes Kepler University Linz. Since 2011, he is full Professor heading the Department for RF-Systems. His research focuses on microwave sensor systems for industrial and automotive applications, integrated sensor concepts, SiGe-based circuit design, RF and microwave subsystems, as well as digital signal processing for sensor signal evaluation. Dr. Stelzer received several awards including the European Microwave Association (EuMA) Radar Prize in 2003, the 2008 IEEE Microwave Theory and Techniques Society (IEEE MTT-S) Outstanding Young Engineer Award, the 2011 German Microwave Conference (GeMiC) Best Paper Award, the 2011 IEEE Microwave Prize, and the Best Paper Award of the International Journal of Microwave and Wireless Technologies (IJMWT) 2016.