Optimal Self-Calibration for Collaborative Sensing in mmWave Radar Networks

Anirban Banik†, Yasamin Mostofi†, Ashutosh Sabharwal§ and Upamanyu Madhow†
†University of California, Santa Barbara, CA, U.S.A. §Rice University, TX, U.S.A
{abanik, ymostofi, madhow}@ucsb.edu ashu@rice.edu

***Abstract*—The emergence of high-resolution millimeter-wave (mmWave) multi-input multi-output (MIMO) radar can enable a powerful framework for collaborative RF sensing with a radar network. Each node can use its range, Doppler, and angle information to track targets within its field of view (FOV), but collaborative *networked* sensing with multiple such nodes can provide several new capabilities for multi-target tracking, including "cellular-style" coverage of large areas, and robust performance under FOV limitations and line-of-sight (LoS) obstructions for individual nodes. However, collaborative target tracking and track-level fusion in a radar network requires knowledge of the radar nodes' *poses* (i.e., positions and orientations) relative to each other. In this paper, we propose an autocalibration strategy based on joint target tracking and pose estimation by fusing measurements corresponding to a moving target seen by multiple radars. We provide an optimal algorithm with a closed-form solution that enables any two nodes tracking a common target to determine their relative poses by matching their estimated tracks. Our preliminary results illustrate how this algorithm can be used as a building block for multi-node calibration, and target track association when tracking multiple targets.**

***Index Terms*—mmWave radar retwork, self-calibration, radar fusion**

I. Introduction

Advances in low-cost silicon implementations of mmWave MIMO radar in the 60 and 77 GHz bands are ushering in a revolution in RF sensing, with enhanced range and Doppler resolution due to the large operating bandwidths and high carrier frequency, and angular information obtained via multiple transmit and receiver antennas. While much of the recent work on integrated communication and sensing in cellular networks has focused on reusing communication waveforms and hardware for sensing, a complementary approach to developing advanced sensing systems is to integrate existing MIMO radar sensors into the cellular infrastructure, using communication as a means of coordinating and fusing measurements from these sensors.

Such "cellular radar" can provide robust, finely detailed environmental awareness through collaborative RF sensing. While radars can "see through" poor weather conditions better than optical modalities such as cameras and LIDAR, an individual radar node has limited range and FOV, and reduced visibility in non-line-of-sight (NLoS) conditions, and the angular resolution is far coarser than for optical sensing. Collaborative sensing using a network of such nodes can significantly alleviate these challenges, enabling sensing coverage over large areas and robust, high-precision multi-target tracking.

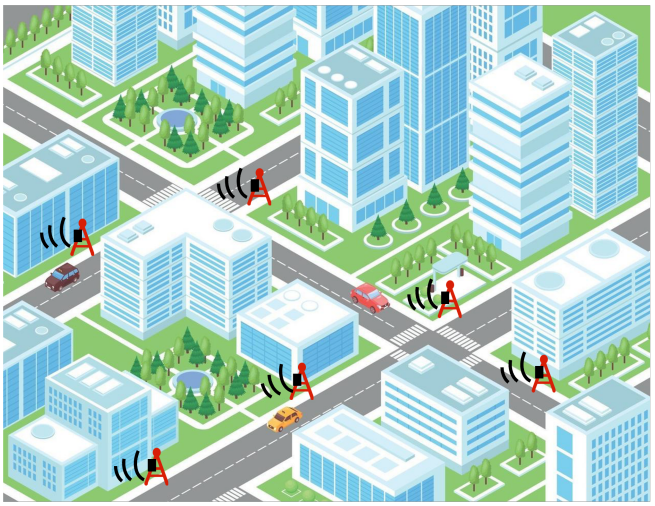

Fig. 1: Collaborative RF sensing using a cellular radar network

However, high-precision collaborative fusion requires accurate calibration, in which the relative poses (positions and orientations) of the radar nodes relative to each other are known across the network. Fine-grained manual calibration during deployment is typically impractical, especially for orientation. Coarse locations may be known using various position measurement technologies, such as Global Navigation Satellite Systems (GNSS) outdoors and WiFi-based positioning indoors, but the resolution is far coarser than the range resolution available with MIMO radars. Thus, self-calibration of radar networks is essential for cost-effective deployment and efficient operation.

We assume that neighboring radar nodes are coarsely synchronized at the frame level and that each radar performs its own range-Doppler-angle processing to obtain its "local" view of the scene. For a calibrated network with large enough communication bandwidth, neighboring nodes could potentially exchange range-Doppler-angle information for high-resolution fusion. However, for initial self-calibration, we consider fusion at the next layer of processing. Neighboring radar nodes seeing a common target each produce a track corresponding to their view of the target, and learn their relative poses by "matching" these tracks in a least squares sense.

Related work: There has been significant recent interest, including experimental validation, in the concept of collaborative sensing with mmWave radar nodes. The authors of [5] study a two-radar network to determine their relative position and orientation using a straight-line walking trajectory of a single moving person. The limitation of this method is that it necessitates a controlled environment, and is therefore unrealistic and inefficient. In [4], the authors utilize a multi-radar system to detect the respiration of multiple sitting subjects. By employing multiple radars, this approach overcomes the shadowing problem encountered in single radar systems, where the radar may not detect multiple people present in the scene if some of them obstruct others. Although this method addresses the multiple-target problem for self-calibration, it is only applicable to stationary subjects. In [6], the authors explore target-based calibration using corner reflectors for radars mounted on a vehicle. This approach addresses a point-set registration problem, where target detections are aligned with known locations in the vehicle coordinate system (VCS) to estimate calibration parameters. However, this work assumes that target locations are already known in the VCS, which is impractical. Joint self-calibration and multi-target tracking in indoor environments is considered in [1], [7], and robustness to reduced fusion rate is evaluated. Joint ego-motion and sensor orientation estimation has been investigated in [2] for radars with non-overlapping FOVs mounted on a vehicle. Joint 3D position and orientation estimation based on moving target tracking has been studied in [3]. These research works on collaborative sensing have explored the potential for self-calibration and multi-target tracking. However, in our work, we take a step back and develop a simplified abstraction that yields a building block for joint self-calibration and multi-target fusion based on an efficient, optimal algorithm.

II. Building block: two radar nodes, one target

Abstraction: We consider a 2D setting, and obtain an elegant mathematical framework by interpreting 2D positions as complex numbers. As a building block, we consider two nodes A and B synchronized at the frame level and tracking a common target (the exact tracking mechanism is abstracted out). The node positions with respect to an arbitrary, and potentially unknown, global frame are $p_A = x_A + jy_A$ and $p_B = x_B + jy_B$, respectively. For a given target seen by both nodes, Node A estimates (in its own coordinate frame) its *local* track $z_A[k] = u_A[k] + jv_A[k]$, $k = 1, ..., K$ while Node B estimates its *local* track as $z_B[k] = u_B[k] + jv_B[k]$, $k = 1, ..., K$, where k indexes successive radar frames. We wish to obtain estimates for the relative position p_{BA} (a complex number) and the real-valued scalar $\theta_{BA} \in [0, 2\pi)$, the rotation of the local coordinate frame of B relative to that of A. See Figure 2. Note that this formulation enables us to model the rotation of coordinate frames with a single scalar rather than with a rotation matrix in the orthogonal group $SO(2)$, and enables us to provide an optimal solution in closed form for a nonlinear least squares problem.

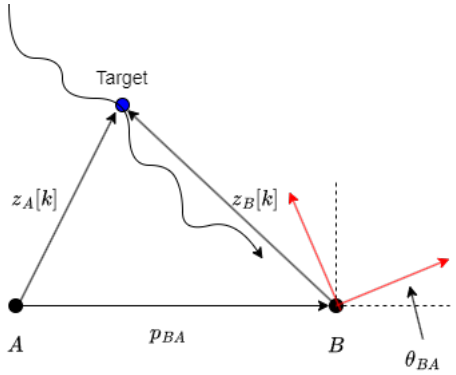

Fig. 2: Two nodes A and B seeing a common target can self-calibrate based on their local tracks.

For relative position p_{BA} and orientation θ_{BA}, the target tracks seen by nodes A and B are related as follows:

$$z_A[k] - p_{BA} \approx e^{j\theta_{BA}} z_B[k] \tag{1}$$

In order to find these parameters, therefore, we seek to minimize the least squares error:

$$J(p, \theta) = \sum_{k=1}^{K} \left| z_A[k] - p - e^{j\theta} z_B[k] \right|^2 \tag{2}$$

Thus, we wish to find

$$(\widehat{p}_{BA}, \widehat{\theta}_{BA}) = \arg\min J(p, \theta) \tag{3}$$

This is a nonlinear least squares problem in p and θ for which it is not *a priori* clear that a global optimum exists. However, as we state in the following theorem, there is indeed a unique global optimum, and we can provide a closed-form solution for both the optimizing arguments and the corresponding minimum value.

The geometry of the optimal solution is best expressed by representing the tracks as complex-valued vectors $\mathbf{z}_A = (z_A[1], ..., z_A[K])^T$ and $\mathbf{z}_B = (z_B[1], ..., z_B[K])^T$. We define the track centroids as

$$\begin{aligned} \overline{z}_A &= \tfrac{1}{K} \textstyle\sum_{k=1}^{K} z_A[k] = \tfrac{1}{K} \mathbf{1}^T \mathbf{z}_A \\ \overline{z}_B &= \tfrac{1}{K} \textstyle\sum_{k=1}^{K} z_B[k] = \tfrac{1}{K} \mathbf{1}^T \mathbf{z}_B \end{aligned} \tag{4}$$

where $\mathbf{1}$ denotes the all-one vector. We define centered tracks by subtracting out the centroids:

$$\begin{aligned} \widehat{\mathbf{z}}_A &= (z_A[1] - \overline{z}_A, ..., z_A[K] - \overline{z}_A)^T = \mathbf{z}_A - \overline{z}_A \mathbf{1} \\ \widehat{\mathbf{z}}_B &= (z_B[1] - \overline{z}_B, ..., z_B[K] - \overline{z}_B)^T = \mathbf{z}_B - \overline{z}_B \mathbf{1} \end{aligned} \tag{5}$$

We use the standard notation $\mathbf{x}^H = (\mathbf{x}^*)^T$ to denote the conjugate transposed of a vector $\mathbf{x}$. Note that $\mathbf{y}^H \mathbf{x}$ is the complex-valued inner product between $\mathbf{x}$ and $\mathbf{y}$, and we denote by $||\mathbf{x}||^2 = \mathbf{x}^H \mathbf{x}$ the energy, or square of the ℓ_2 norm, of a vector $\mathbf{x}$. We can now state the following theorem characterizing the solution of (2)-(3).

Theorem 1: The global minimum of the least squares problem (2)-(3) is attained by

$$\begin{aligned} \widehat{\theta}_{BA} &= -\angle \widehat{\mathbf{z}}_A^H \widehat{\mathbf{z}}_B \\ \widehat{p}_{BA} &= \overline{z}_A - e^{j\widehat{\theta}_{BA}} \overline{z}_B \end{aligned} \tag{6}$$

Furthermore, the minimum cost attained by this solution is given by

$$J_{min}(B,A) = ||\widehat{\mathbf{z}}_A||^2 + ||\widehat{\mathbf{z}}_B||^2 - 2\left|\widehat{\mathbf{z}}_A^H \widehat{\mathbf{z}}_B\right| \tag{7}$$

Proof: The cost function (2) can be written compactly as

$$J(p,\theta) = \left(\mathbf{z}_A - p\mathbf{1} - e^{j\theta}\mathbf{z}_B\right)^H \left(\mathbf{z}_A - p\mathbf{1} - e^{j\theta}\mathbf{z}_B\right) \tag{8}$$

For fixed θ, $J(p,\theta)$ is a quadratic function of the complex variable p that can be minimized by setting the derivative $\frac{\partial J}{\partial p^*} = 0$. This is easily seen to yield the optimal solution (as a function of θ)

$$\widehat{p}(\theta) = \overline{z}_A - e^{j\theta}\overline{z}_B \tag{9}$$

Plugging this into (8), we obtain upon simplification that

$$J(\widehat{p}(\theta),\theta) = \left(\widehat{\mathbf{z}}_A - e^{j\theta}\widehat{\mathbf{z}}_B\right)^H \left(\widehat{\mathbf{z}}_A - e^{j\theta}\widehat{\mathbf{z}}_B\right) \tag{10}$$

which further simplifies to

$$J(\widehat{p}(\theta),\theta) = ||\widehat{\mathbf{z}}_A||^2 + ||\widehat{\mathbf{z}}_B||^2 - 2\mathrm{Re}\left(e^{j\theta}\widehat{\mathbf{z}}_A^H \widehat{\mathbf{z}}_B\right) \tag{11}$$

Writing the complex inner product $\widehat{\mathbf{z}}_A^H \widehat{\mathbf{z}}_B = Re^{j\phi}$ (in polar form), the third term on the right-hand side can be written as

$$-2\mathrm{Re}\left(e^{j\theta}\widehat{\mathbf{z}}_A^H \widehat{\mathbf{z}}_B\right) = -2\mathrm{Re}\left(e^{j\theta}Re^{j\phi}\right) = -2R\cos(\phi+\theta)$$

which is minimized for $\theta = -\phi = -\angle \widehat{\mathbf{z}}_A^H \widehat{\mathbf{z}}_B$. This, together with (9), completes the computation of the optimal solution (6). Plugging into (11), we obtain the minimum cost (7).

Remarks:

(1) The solution (6) has an intuitively pleasing geometric interpretation. The coordinate frame rotation $\widehat{\theta}_{BA}$ geometrically aligns centered versions of the local tracks at the two radar nodes. The position translation $\widehat{p}_{BA}$ is then given by the difference of the centroids of the local tracks after bringing them both to the same frame of reference.

(2) In the next section, we provide numerical results showing how the pairwise solution provided by the preceding theorem can be used as a building block to calibrate multiple radar nodes tracking the same target.

(3) The value of the minimum cost (7) can be leveraged for track association: attempted calibration using local tracks corresponding to different targets would lead to larger cost than if the local tracks correspond to the same target.

(4) While our calibration framework is quite general, operating at a track-level abstraction, the parameters being calibrated (position and orientation) would be relevant for fusing measurements from MIMO radars providing range, Doppler and angle measurements.

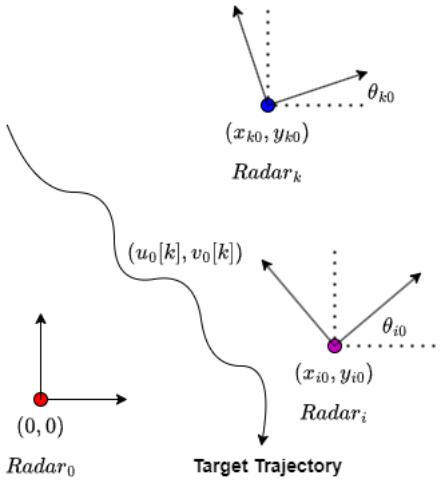

Fig. 3: Calibrating multiple nodes

III. Calibrating multiple nodes

Now, consider N stationary radar nodes with overlapping FOVs, as depicted in Figure 3. For a target visible to all nodes, enforcing least squares consistency among local tracks for all pairs decomposes into $N(N-1)/2$ instances of the problem solved in Theorem 1 due to the additive nature of the least squares cost. Thus, we obtain estimates of relative position $\widehat{p}_{ik}$ and relative orientation $\widehat{\theta}_{ik}$ for each distinct pair $0 \le i < k \le N-1$. These can be merged, for example, by choosing one of the nodes as anchor: this provides a common frame for expressing the results of collaborative sensing even if we do not know the anchor's position and orientation in a global frame. For example, if node 0 is the reference, we may set $p_0 = 0 + 0j$, and estimate the position of the ith node as

$$\widehat{p}_i = \frac{1}{N-1}\sum_{k\neq i}(\widehat{p}_{ik}e^{j\widehat{\theta}_{k0}} + \widehat{p}_{k0}) \ , \quad 1 \le i \le N-1 \tag{12}$$

where $\widehat{p}_0 = \widehat{p}_{00} = 0 + 0i$. For each node i, the kth term in (12) is the vector sum of its estimated position relative to k (expressed in k's coordinate frame) and the relative position of k from the anchor node 0.

The orientation of node i's coordinate frame with respect to node 0's can be estimated by similar averaging:

$$\widehat{\theta}_i = \frac{1}{N-1}\sum_{k\neq i}(\widehat{\theta}_{ik} + \widehat{\theta}_{k0}) \ , \quad 1 \le i \le N-1 \tag{13}$$

While we have a guarantee of optimality for two-node calibration in Theorem 1, we do not have similar assurances for the preceding merging procedure, since the estimates obtained from pairwise calibration are not independent. Thus, while our initial experiments show the efficacy of this procedure, it may be possible to improve upon it.

IV. Results and Discussion

Target tracking information for each radar is forwarded to a central fusion node (which may be located at one of the radar nodes) for calibration and fusion. We assume that all radars

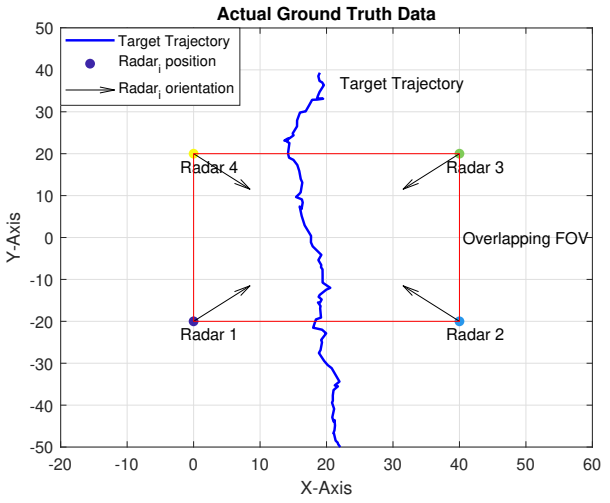

(a)

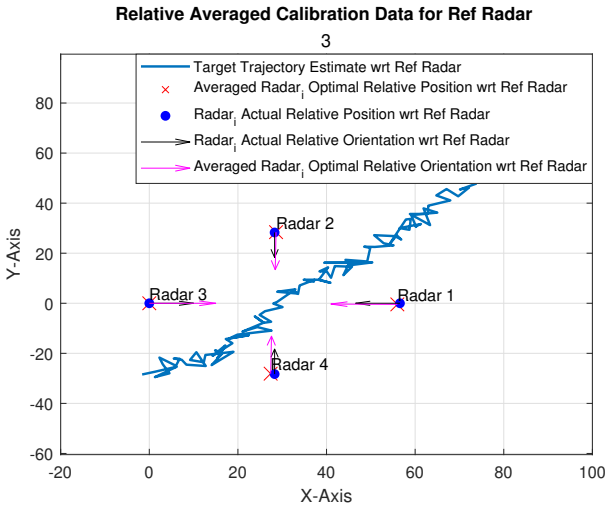

(b)

Fig. 4: (a) Ground truth calibration data for Scenario 1 with $N = 4$ radars and one moving target in their overlapping FOV, shown by the red lines. (b) Averaged relative calibration results for $Radar_3$ as reference. The target trajectory estimate relative to $Radar_3$ is shown here.

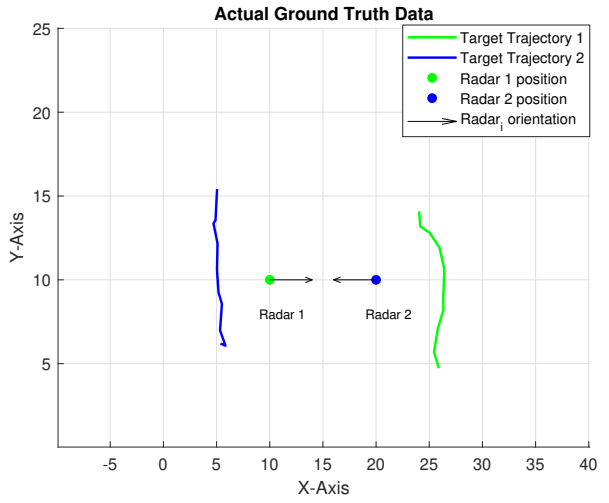

(a)

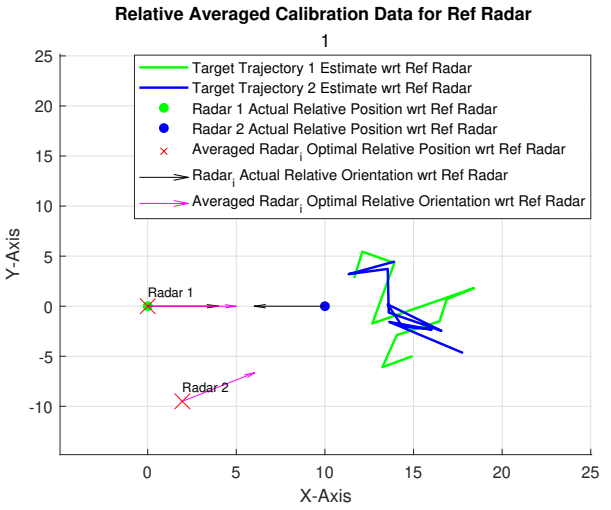

(b)

Fig. 5: (a) Ground truth calibration data for Scenario 2 with $N = 2$ radars and two moving targets. Target 1 is in the FOV of Radar 1 and Target 2 is in the FOV of Radar 2. (b) Relative calibration results for $Radar_1$ as reference. The target 1 and 2 trajectories w.r.t. $Radar_1$ after calibration (denoted by the green and blue lines respectively) indicate that they are not the same target.

have already estimated target trajectories in *local* Cartesian coordinates (x, y), with i.i.d. noise across each axis for each time sample of each local track. In our numerical results, we assume that the noise samples are distributed as $\mathcal{N}(0, 4)$. In practice, noise covariances could be obtained as the output of a tracker, which requires extension of our framework to non-white, time-varying noise covariances that are different for the different radars.

Fig. 4a shows simulation results for Scenario 1: $N = 4$ radar nodes with one moving target visible to all nodes. The efficacy of our self-calibration algorithm is confirmed by comparing the results to ground truth relative poses (Fig. 4b).

Scenario 2 (depicted in Fig. 5a) illustrates that the closed-form solution in Theorem 1 also provides a simple criterion for track association: when two radar nodes see two *different* targets but try to solve the least squares problem in Theorem 1, the minimum cost attained will be large, indicating that the tracks are not the same (incorrect target-track association). Fig. 5a shows two targets and two radar nodes: each target is in the FOV of one radar and out of the FOV of the other. Applying our self-calibration algorithm with the corresponding local tracks and referring the local tracks to the coordinate frame of radar 1, we see that the tracks do not match (Fig. 5b), which provides a pictorial illustration of the potential of

using the minimum value J_{min} in Theorem 1 as a criterion for track association in a radar network.

Scenario 3 (depicted in Fig. 6a) shows two radar nodes and five moving targets, three of which are in the overlapping FOV of both radars, while the other two are visible only to one radar each. Fig. 6b shows the relative calibration result for correctly associated tracks, while Fig. 6c shows the result for incorrect track association. From this we can conclude that correct target-track association is necessary for precise self-calibration in presence of multiple targets. As part of our future work, we aim to achieve joint self-calibration and target-track association for a multi-target scenario.

V. Conclusion

We propose a "cellular radar" approach to integrated communication and sensing that leverages existing radar technology within the cellular infrastructure. Networked self-calibration as considered here enables opportunistic deployment with minimal measurement requirements. We provide a simplified closed-form solution for self-calibration in a two-radar network, and extend it to multiple nodes. We also show that the minimum value of the residual obtained as a part of the closed-form solution can be leveraged for track association in multiple-target scenarios.

Acknowledgment

This work was supported in part by the Center for Ubiquitous Connectivity (CUbiC), sponsored by Semiconductor Research Corporation (SRC) and Defense Advanced Research Projects Agency (DARPA) under the JUMP 2.0 program, and in part by the National Science Foundation under grant CNS-2215646.

References

[1] Marco Canil, Jacopo Pegoraro, Anish Shastri, Paolo Casari, and Michele Rossi. Oracle: Occlusion-resilient and self-calibrating mmwave radar network for people tracking. *IEEE Sensors Journal*, 24(3):3157–3171, 2024.

[2] Timo Grebner, Vinzenz Janoudi, Pirmin Schoeder, and Christian Waldschmidt. Self-calibration of a network of radar sensors for autonomous robots. *IEEE Transactions on Aerospace and Electronic Systems*, 59(5):6771–6781, 2023.

[3] Timo Grebner, Matthias Linder, Nicolai Kern, Pirmin Schoeder, and Christian Waldschmidt. 6d self-calibration of the position and orientation of radar sensors in a radar network. In *2022 19th European Radar Conference (EuRAD)*, pages 157–160, 2022.

[4] Shunsuke Iwata, Takato Koda, and Takuya Sakamoto. Multiradar data fusion for respiratory measurement of multiple people. *IEEE Sensors Journal*, 21(22):25870–25879, 2021.

[5] Shuai Li, Junchen Guo, Rui Xi, Chunhui Duan, Zhengang Zhai, and Yuan He. Pedestrian trajectory based calibration for multi-radar network. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 1–2, 2021.

[6] Kunle T. Olutomilayo, Mojtaba Bahramgiri, Saeid Nooshabadi, and Daniel R. Fuhrmann. Extrinsic calibration of radar mount position and orientation with multiple target configurations. *IEEE Transactions on Instrumentation and Measurement*, 70:1–13, 2021.

[7] Anish Shastri, Marco Canil, Jacopo Pegoraro, Paolo Casari, and Michele Rossi. mmscale: Self-calibration of mmwave radar networks from human movement trajectories. In *2022 IEEE Radar Conference (RadarConf22)*, pages 1–6, 2022.

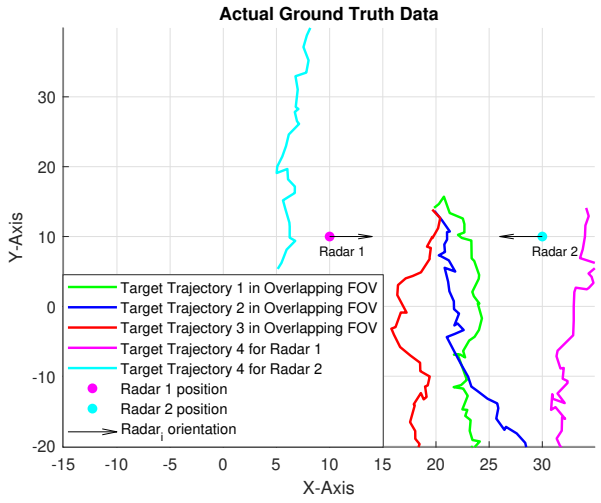

(a)

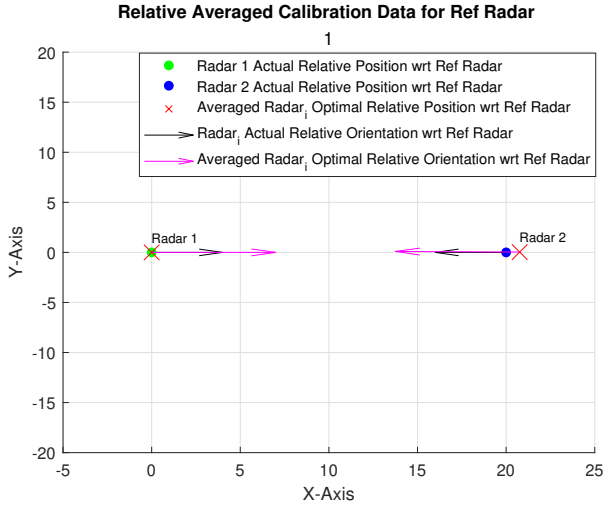

(b)

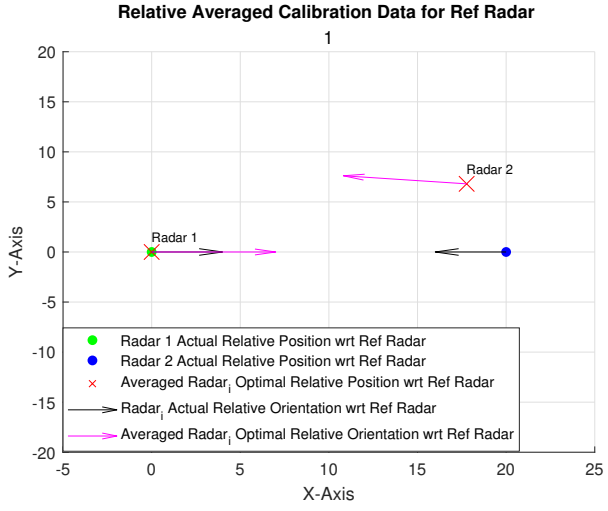

(c)

Fig. 6: (a) Ground truth calibration data for Scenario 3 with $N = 2$ radars and five moving targets. Three targets are in the overlapping FOV of the radars, while the other two are only visible to one radar. (b) Relative calibration results w.r.t $Radar_1$ for correct track association. (c) Relative calibration results w.r.t $Radar_1$ for incorrectly associated tracks.