

Cooperation and Federation in Distributed Radar Point Cloud Processing

S. Savazzi, V. Rampa, S. Kianoush
Consiglio Nazionale delle Ricerche (CNR)
IEIIT institute, Milano
Email: {name.surname}@ieiit.cnr.it

A. Minora, L. Costa
CogniMade S.r.l, <https://www.cognimade.com>
Via C. Colombo 10/A, I-20066 Melzo, Italy
Email: {name.surname}@cognimade.com

Abstract—The paper considers the problem of human-scale RF sensing utilizing a network of resource-constrained MIMO radars with low range-azimuth resolution. The radars operate in the mmWave band and obtain time-varying 3D point cloud (PC) information that is sensitive to body movements. They also observe the same scene from different views and cooperate while sensing the environment using a sidelink communication channel. Conventional *cooperation* setups allow the radars to mutually exchange raw PC information to improve ego sensing. The paper proposes a *federation* mechanism where the radars exchange the parameters of a Bayesian posterior measure of the observed PCs, rather than raw data. The radars act as distributed parameter servers to reconstruct a global posterior (i.e., federated posterior) using Bayesian tools. The paper quantifies and compares the benefits of radar federation with respect to cooperation mechanisms. Both approaches are validated by experiments with a real-time demonstration platform. Federation makes minimal use of the sidelink communication channel (20 ÷ 25 times lower bandwidth use) and is less sensitive to unresolved targets. On the other hand, cooperation reduces the mean absolute target estimation error of about 20%.

Index Terms—Distributed and federated radar networks, Point Cloud processing, localization, RF sensing, Bayesian estimation.

I. INTRODUCTION

The emerging paradigm of integrated radar sensing and communication represents a promising enabler for pervasive smart environments. These systems target both network performance improvements [1] and enhanced spatial discovery functions over which applications/verticals can be optimally implemented [2]. Targeting indoor environments, recent advances in hardware and digitalization technology enabled the deployment of low-cost, low-power COTS (Commercial Off-The-Shelf) radar devices of small dimension and suitable for smart home and industrial applications [3]. In this framework, a crucial challenge is related to the development of accurate modelling tools to optimize the interactions between radar devices and environment as well as among the radars themselves.

Radar technology has recently evolved to support a multitude of functions, ranging from detection, classification, multi-subject localization [3], human-robot cooperation [4], human gait, gesture, and activity recognition [5], [6]. Indoor




Fig. 1. Model of the proposed federated radar system and its point cloud representation.

classification and identification of human subjects is based on the processing of point cloud (PC) information extracted from millimeter-wave radars with relative low angular resolution and number of antennas [7], [8]. PCs or raw data returned by the radar can be used as input to machine learning [7], [9] to extract the features of the target. Merging of PC data from multiple radar sensors have been reported in [10] for better detection. A deep learning approach to fuse the features of both PCs and range-Doppler for classifying activities has been proposed in [9]. A Bayesian approach for fusing/grouping of radar data [11] is shown to help in angular resolution and provide added information about the object which can be used for subsequent target/object detection. On the other hand, the real-time exchange of radar data can be infeasible in typical localization and tracking applications found in smart home environments.

The paper investigates radar cooperation models where a sidelink communication channel of limited bandwidth is used to coordinate the sensing task. According to the scenario in Fig. 1, a network of radars is deployed to monitor the same scene S_t at time t . Each radar is equipped with multiple antennas, operates in frequency modulated continuous wave (FMCW) mode [5], and obtains time-varying 2D/3D Point Cloud (PC) information, through local/independent Digital Signal Processing (DSP) stages. Usage of multiple radars in the same environment can mitigate shadowing effects of the sensed PCs as well as improve resolution that would be obtained from ego behavior. On the other hand, effective denoising algorithms are needed to limit parallel/cross chirp interference sources, Signal to Noise Ratio (SNR) degradation,




Fig. 2. Cooperation v. Federation PC processing architectures: examples from real data. From left to right: (a) scene S_t consist of 2 targets co-present and moving in opposite directions with two radars monitoring the same scene; (b) *cooperation mode*: the radars exchange pre-processed PCs $\hat{Q}_{t,k}$ and obtain a target probability measure via local fusion of the received PCs; (c) *federation mode*: the radars exchange the parameters $\mathbb{P}_k$ characterizing the local posterior measure approximated as Gaussian mixture (number of Gaussian components M_k , weights $\beta_{k,m}$, mean $\mu_{k,m}$, and covariance $\Gamma_{k,m}$ of each component m).

blind spots, and ghost object effects. The joint analysis of the spatio-temporal structure of the radar PCs is here explored with the goal of improving both accuracy and resolution of the individual radars.

Contributions. Unlike classical distributed radar processing that merge/fuse raw data and/or PC information, the paper proposes a *federation model* that does not rely on data intensive fusion of the radars raw samples on a server. The model is rooted at federation tools [12], and exploits sidelink communication channels of limited bandwidths. The radars keep the raw data on the devices, thus protecting data ownership as in as in classical federation approaches. On the other hand, rather than exchanging the machine learning model parameters, we propose an architecture where the radars share the parameters of a Bayesian local posterior probability measure that describes the (local) “beliefs” on the observed scene S_t . Local posteriors follow a Gaussian mixture model that is chosen for point clouds (PC) representation. Proposed federation model is evaluated experimentally targeting a building automation scenario. It is compared with classical cooperation, where the radars exchange raw or pre-processed PC info. The benefits of federation vs. cooperation are quantified considering accuracy, resolution, and communication requirements as well.

The paper is organized as follows. Sect. II introduces the system model and the PC processing tools. Cooperation and Bayesian federation settings are analyzed in Sect. III, while Sect. IV describes the developed prototype and presents the results of the on-field experiments. Finally, future activities are discussed in Sect. V.

II. POINT CLOUD MODEL AND PROCESSING

As depicted in Fig. 1, the focus is on an indoor scenario where K radio devices operate as MIMO radars to sense and track the same time-varying scene S_t from different points of view. The scene might be indicative of specific body trajectories followed by one or more users (i.e., the targets) in the monitored space, or specific actions performed by a user/subject in a specific location(s). The observed scene is represented in terms of (true) PCs $\mathbf{q}_i \in S_t$ relative to the moving targets in a 2D or 3D space: these points can be generally

described through an empirical probability mass function on a 2D or 3D pixel/voxel grid. At any time $t > 0$, the scene S_t follows a stochastic random process with probability function $p(\cdot)$ where $p(\mathbf{q}_i \in S_{t+1} | S_t)$ is assumed to be memoryless (i.e., Markovian); therefore, the PC distribution of the scene at time $t + 1$ depends only on the scene at present time t .

Each k -th device is equipped with a radar and measures time-varying (3D) noisy PC information $\hat{\mathbf{q}}_{i,k} \in Q_{t,k}$. In particular, we adopt the following PC formation model

$$Q_{t,k} = \{\hat{\mathbf{q}}_{i,k}\} = \{\mathbf{q}_i + \mathbf{n}_{i,k}\}_{\mathbf{q}_i \in S_t} \cup \{\mathbf{o}_j\}_{\mathbf{o}_j \in O_{t,k}}, \quad (1)$$

where $\mathbf{n}_{i,k}$ is the additive noise produced by k -th radar PC processing hardware, and $\mathbf{o}_j \in O_{t,k}$ is a set of outlier points, typically caused by chirp interference, and SNR degradation. In what follows, the PCs are obtained locally using the Cell Averaging (CA) Constant False Alarm Rate (CA-CFAR) method. Other approaches may be implemented as in [13].

The radar devices operate autonomously and independently to process and extract local point clouds $Q_{t,k}$. On the other hand, a sidelink communication channel is used by the radars to share selected information with the goal of improving ego knowledge. The topology of the radar network is modelled as a directed graph $\mathcal{G} = (\mathcal{V}, \xi)$ with the set of nodes $\mathcal{V} = \{1, \dots, K\}$ and edges, or links, $\xi = \{1, \dots, L\}$. The neighbor set of the k -th device is denoted as $\mathcal{N}_k = \{h \in \mathcal{V}, h \neq k : (h, k) \in \xi\}$, with cardinality $|\mathcal{N}_k|$.

The goal of the k -th radar at time t is to output a higher quality reconstructed scene $\hat{S}_{t,k}$ by using side information obtained from the radar neighborhood. In the following, we define the ensemble of PCs in the neighbor set $\mathcal{N}_k$, as $Q_{t,k} \cup \{Q_{t+\Delta t_h, h}\}_{h \in \mathcal{N}_k}$, where $Q_{t+\Delta t_h, h}$ collects the PCs obtained from neighbor $h \in \mathcal{N}_k$ and Δt_h accounts for the relative time offset w.r.t. the k -th radar. Compared to the observed points $Q_{t,k}$, the estimated scene $\hat{S}_{t,k}$ should be closer to the true one S_t , and suitable for further machine learning based processing, e.g. based on PointNet methods [14], [15].

III. DISTRIBUTED RADAR MODELS AND SETTINGS

In the following, we compare two distributed settings as summarized in the Fig. 2. The example refers to 2 radars but

can be easily generalized. For the cooperation case (Fig. 2(b)), the radars mutually share a subset $\hat{\mathbb{Q}}_{t,k} \subseteq \mathbb{Q}_{t,k}$ of the local PCs (after displacement correction and outlier removal). Scene $\hat{\mathbb{S}}_{t,k}$ reconstruction at radar k is obtained from the global Bayesian *global posterior probability* measure $\pi_t^{(C)}$. Considering the federation case (Fig. 2(c)), scene $\hat{\mathbb{S}}_{t,k}$ reconstruction is solved by a loose cooperation (federation) of radars that share the parameters of a *local* posterior probability measure, rather than exchanging the raw points $\hat{\mathbb{Q}}_{t,k}$. The local posterior measures are approximated by a Gaussian mixture function: therefore, the radars exchange the parameters $\mathbb{P}_k$ characterizing the mixture model (number of Gaussian components M_k , weights $\beta_{k,m}$, mean $\boldsymbol{\mu}_{k,m}$, and covariance $\boldsymbol{\Sigma}_{k,m}$ of each component m). Scene $\hat{\mathbb{S}}_{t,k}$ reconstruction is now obtained from the *federated posterior probability* $\pi_t^{(F)}$. Notice that the proposed federation model has roots in Bayesian federated learning [16], [17] and decentralized setups [12].

A. Cooperation: point cloud data exchange

Radars engaged in device cooperation share the PCs obtained from the local DSP processor. Before transmission, the local points $\mathbb{Q}_{t,k}$ in (1) are pre-processed

$$\hat{\mathbb{Q}}_{t,k} = \left\{ \hat{\mathbf{q}}_{i,k} + \hat{\mathbf{d}}_{i,k} \right\}_{\hat{\mathbf{q}}_{i,k} \in \mathbb{Q}_{t,k} \setminus \hat{\mathbb{Q}}_{t,k}} \quad (2)$$

to correct the displacement errors $\hat{\mathbf{d}}_{i,k}$ (i.e., depending on the relative position of the radar) and to remove the outliers $\hat{\mathbb{Q}}_{t,k}$ detected in $\mathbb{Q}_{t,k}$. The dBscan method [18], [19] is adopted for outlier estimation: considering the neighborhood $\mathcal{N}_k$, the PC ensemble $\mathbb{Q}_{t,\mathcal{N}_k}$ received by the k -th radar is

$$\mathbb{Q}_{t,\mathcal{N}_k} = \hat{\mathbb{Q}}_{t,k} \cup \left\{ \hat{\mathbb{Q}}_{t+\Delta t_h,h} \right\}_{h \in \mathcal{N}_k}. \quad (3)$$

Scene reconstruction is based on the *global* measure $\pi_t^{(C)}$ of the posterior probability, namely the posterior obtained from the PCs observed in the radar neighborhood:

$$\pi_t^{(C)} = p(\mathbf{q}_i \in \mathbb{S}_t | \mathbb{Q}_{t,\mathcal{N}_k}). \quad (4)$$

Finally, the PCs representing the true scene are here identified as those with largest posterior probability measure:

$$\hat{\mathbb{S}}_{t,k} = \left\{ \mathbf{q}_i : \pi_t^{(C)} > \tau \right\}, \quad (5)$$

with threshold τ defined at calibration time (see Sect. IV). As analyzed in Sect. IV, the posterior measure (5) is used to track the positions of the targets $(\hat{x}_t, \hat{y}_t)$ in the space: these correspond to the local maxima of the posterior, according to maximum a-posteriori (MAP) estimation.

According to Bayesian filtering tools, the posterior $\pi_t^{(C)}$ is updated at time t by an iterative approach that uses the PCs representing the previous reconstructed scene $\hat{\mathbb{S}}_{t-1,k}$ as

$$\pi_t^{(C)} \propto p(\hat{\mathbf{q}}_i \in \mathbb{Q}_{t,\mathcal{N}_k} | \mathbb{S}_t) \cdot p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{S}}_{t-1,k}), \quad (6)$$

with the conditional likelihood $p(\hat{\mathbf{q}}_i \in \mathbb{Q}_{t,\mathcal{N}_k} | \mathbb{S}_t)$ and the prior measure $p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{S}}_{t-1,k})$. Prior and conditional likelihood models suitable for tracking of human subjects are discussed in

the following sections. The interested reader may also refer to [20], [22]. Finally, notice that $p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{S}}_{t-1,k})$ approximates the a-priori probability term $p(\mathbf{q}_i \in \mathbb{S}_t | \mathbb{Q}_{t-1,\mathcal{N}_k}) = p(\mathbf{q}_i \in \mathbb{S}_t | \mathbb{S}_{t-1}) \cdot \pi_{t-1}^{(C)}$.

B. Federation: local posterior exchange

Rather than exchanging the PC set $\hat{\mathbb{Q}}_{t,k}$, in the proposed federation model the radars share the parameters of the *local posterior* probability measure $p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{Q}}_{t,k})$. The PC samples are kept on the individual radars and not transferred: therefore, the global posterior $\pi_t^{(C)}$ in (4) is now approximated by a mixture of the local posteriors to obtain $\pi_t^{(F)}$ as

$$\pi_t^{(F)} = \alpha_k \cdot p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{Q}}_{t,k}) + \sum_{h \in \mathcal{N}_k} \alpha_h \cdot p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{Q}}_{t,h}), \quad (7)$$

where $\alpha_k + \sum_{h \in \mathcal{N}_k} \alpha_h = 1$ and $\alpha_k = \frac{Q_k}{Q_k + \sum_{h \in \mathcal{N}_k} Q_h}$. Q_k is the number of points measured by radar k . The *federated posterior* $\pi_t^{(F)}$ can be evaluated either centrally by a server collecting the local posterior parameters, namely the Parameter Server, (PS) or distributedly. Scene reconstruction is represented similarly as in (5), now with $\hat{\mathbb{S}}_{t,k} = \left\{ \mathbf{q}_i : \pi_t^{(F)} > \tau \right\}$. Notice that, the estimated scene $\hat{\mathbb{S}}_{t,k}$ from the federated posterior is used to update the *local* posterior on the next step $t+1$, thus augmenting the k -th radar knowledge about the monitored space. Likewise equation (6), the local posterior can be also evaluated iteratively as $p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{Q}}_{t,k}) \propto p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{S}}_{t-1,k}) \cdot p(\hat{\mathbf{q}}_i \in \hat{\mathbb{Q}}_{t,k} | \mathbb{S}_t)$ using the prior $p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{S}}_{t-1,k})$ and the (local) conditional likelihood $p(\hat{\mathbf{q}}_i \in \hat{\mathbb{Q}}_{t,k} | \mathbb{S}_t)$, respectively.

Cooperation (4) versus federation (7) tradeoffs are discussed in Sect. IV targeting the problem of multi body tracking. In what follows, we adopt the multivariate Gaussian mixture distributions to represent both the local, the global posterior probability measures, and the priors.

C. Gaussian mixture approximations

Local/global posteriors follow a multivariate Gaussian mixture model. The local posterior is therefore approximated as

$$p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{Q}}_{t,k}) \simeq \sum_{m=1}^{M_k} \beta_{k,m} \cdot \mathcal{N}(\boldsymbol{\mu}_{k,m}, \boldsymbol{\Sigma}_{k,m}), \quad (8)$$

with parameters $\mathbb{P}_k = \{\beta_{k,m}, \boldsymbol{\mu}_{k,m}, \boldsymbol{\Sigma}_{k,m}\}_{m=1}^{M_k}$ for each k -th radar. M_k is the number of components observed by radar k , while $(\boldsymbol{\mu}_{k,m}, \boldsymbol{\Sigma}_{k,m})$ are the mean and covariance of the m -th component¹ and $\sum_{m=1}^{M_k} \beta_{k,m} = 1$. Number of components M_k depends on the post-processed PCs $\hat{\mathbb{Q}}_{t,k}$ after outlier removal [19] and will be assessed in the following experimental analysis. The federation model (7) requires the radars share the parameters $\mathbb{P}_k$ of the Gaussian mixture (8). Federated posterior measure $\pi_t^{(F)}$ is thus obtained by substituting (8) in (7). For the cooperation model, the global posterior $\pi_t^{(C)}$ is

$$\pi_t^{(C)} = p(\mathbf{q}_i \in \mathbb{S}_t | \mathbb{Q}_{t,\mathcal{N}_k}) \simeq \sum_{m=1}^N \beta_{C,m} \cdot \mathcal{N}(\boldsymbol{\mu}_{C,m}, \boldsymbol{\Sigma}_{C,m}), \quad (9)$$

¹for 3D PCs, $\boldsymbol{\mu}_{k,m}$ has dimension 3×1 and covariance $\boldsymbol{\Sigma}_{k,m}$ is 3×3 .




Fig. 3. Proposed architecture and system model. Radar k monitoring the sidelink (S) and receiving information from radar h . The radar MIMO HW supports range-azimuth and elevation tracking and it is equipped with a DSP to extract time-varying 2D/3D point cloud information. The radar CFAR configuration parameters can be controlled via UART communication with the MCU.




Fig. 4. Scene example featuring 2 subjects moving from position L to H and from M to F. Posteriors observed over three consecutive time samples ($t = 1, 2, 3$): a) no-cooperation case with $k = 1, 2, 3$; and b) federation and cooperation cases for $K = 3$ radars with deployment described above.

with $\sum_{m=1}^N \beta_{C,m} = 1$, and typically $N \gg M_k$. Global and federated posteriors are compared in the following section in terms of the Kullback-Leibler (KL) divergence.

As far as the prior $p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{S}}_{t-1,k})$ is concerned, we assume that the objects/targets in the scene consist of smooth and well separated patches moving at maximum speed v . For




Fig. 5. Kullback-Leibler (KL) $D_{KL}[\cdot]$ divergence sample probability functions $\Pr[D_{KL}]$: solid line compares global and federated posteriors ($D_{KL}[\pi_t^{(C)} || \pi_t^{(F)}]$), dashed lines compare global and local posteriors for the three deployed radars. Probability measure is obtained by collecting 900 consecutive sample divergences (9 sec. with 10ms per localization update) corresponding to 2 subjects moving from position L to H and from M to F (same scenario as in Figure 4).

S points in the scene $\hat{\mathbb{S}}_{t-1,k}$, the prior measure

$$p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbb{S}}_{t-1,k}) \simeq \frac{1}{S} \sum_{\mathbf{q}_i \in \hat{\mathbb{S}}_{t-1,k}} \mathcal{N}[\mathbf{q}_i, \sigma^2(v) \times \mathbf{I}] \quad (10)$$

is obtained as the superposition of Gaussian distributions each modelling the motion of individual particle $\mathbf{q}_i$ with $\sigma^2(v)$ obtained from a random walk model [23]. Other priors can be also exploited to quantify the smoothness and sharpness of the targets as proposed in [20], [21] for imaging applications.

IV. EXPERIMENTAL DATA AND VALIDATION

This section discusses cooperation and federation models based on a prototype implementation and experimental data. Proposed architecture is summarized in Fig. 3: in the example, the radar k is monitoring the sidelink (S) and receiving information from radar h . The radars work in the 77 GHz band and support MIMO-FMCW radio featuring an array of 3 TX and 4 uniformly spaced RX antennas. Each radar has an azimuth Field Of View (FOV) of $\pm 60^\circ$ angle and range resolution of 25° and 4.2cm, 3.9 GHz band (sweep time $36\mu s$). It integrates a C674x DSP and an ARM R4F MCU with TX output power set to 12 dBm. PC samples are obtained from the radar DSP using the CA-CFAR method configured with range and Doppler detection thresholds of 15dB. Then, PCs are further processed using a ARM-Cortex-A57 SoC (Jetson Nano device model). Notice that the radar CFAR configuration parameters can be controlled via UART communication with the MCU.

Three wall-mounted radars are deployed as detailed in Fig. 4 (positions are relative to radar $k = 1$). The scene $\mathbb{S}_t$, detailed in Fig. 4 (on top), consists of 2 co-present human

subjects moving inside an indoor lab environment. Trajectories follow 6 landmark points (A-M), with corresponding relative ground-truth locations detailed in the same figure. Radars exchange either raw PCs (cooperation) or local model parameters $\mathbb{P}_k$ (federation) via WiFi links while synchronization is implemented using NTP (Network Time Protocol). Impact of inaccurate clock source on performance [24] is here assessed experimentally: however, a dedicated analysis is out of scope.

A. Global, federated and local posterior analysis

Figure 4 compares the local posterior distributions $p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbf{Q}}_{t,k})$ for three (isolated) radars $k = 1, 2, 3$ (Fig. 4.a) with the global $\pi_t^{(C)}$ and federated posterior $\pi_t^{(F)}$ (Fig. 4.b). Posteriors are observed at 3 consecutive time instants $t = 1, 2, 3$. The radar $k = 1$ correctly identifies the target T2 following the path $M \rightarrow G \rightarrow F$, but it is less effective for T1, moving in the opposite direction, from $L \rightarrow H$, due to low SNR. On the other hand, the radar $k = 2$ identifies T1, while T2 is now shadowed. Cooperation and federation cases refer to a scenario where the radar $k = 1$ obtains side information from neighbors $k = 2, 3$: in both cases, the radar is able to effectively resolve both targets.

Figure 5 now assess in more detail the capability of the federated posterior $\pi_t^{(F)}$ in (7) to effectively reproduce the global posterior $\pi_t^{(C)}$ as obtained by fusion of the PCs on the radar. We thus analyze and compare the Kullback-Leibler (KL) distance D_{KL} [25] between the global, the federated and the local posterior measures. Notice that the global posterior in (4) approximates the posterior that would be obtained by letting the radar exchange the (full) raw 3D radar data. However, exchanging radar data over sidelinks is not considered in this paper as infeasible due to bandwidth constraints introduced by the real-time data sharing process (Fig. 3). Considering that the posterior is computed on each localization update, we collect consecutive KL distances corresponding to 2 subjects moving from position L to H and from M to F (see Figure 4) and obtain a sample probability $\Pr[D_{KL}]$. Solid line compares global and federated posteriors ($D_{KL}[\pi_t^{(C)} || \pi_t^{(F)}]$), while dashed lines compare global and local posteriors ($D_{KL}[\pi_t^{(C)} || p(\mathbf{q}_i \in \mathbb{S}_t | \hat{\mathbf{Q}}_{t,k})]$, $k = 1, 2, 3$) for the three radars. Proposed federation setup confirms the property of being a loose form of cooperation that makes the federated posterior closer to the global one, compared with the local measures, but with no exchange of the raw PCs.

B. MAE, unresolved targets and communication overhead

Fig. 6 and Tab. I provide a quantitative evaluation of the localization accuracy and communication overhead. T1 and T2 are now located in positions C, I and move towards the corresponding landmarks E, G, by following the trajectory $C \rightarrow D \rightarrow E$ and $I \rightarrow M \rightarrow G$, respectively. The estimated positions of the subjects ($\hat{x}_t, \hat{y}_t$) in the space correspond to the local maxima of the posteriors ($\pi_t^{(C)}, \pi_t^{(F)}$), according to MAP estimation. Notice that before obtaining the local maxima, we apply a thresholding (5) of the posteriors, with $\tau = 0.45$ to limit the search time of the local maxima to the PCs




Fig. 6. Probability distribution of estimated positions $\hat{x}_t$ for targets T1, T2 located in C, I (red bars) moving to G, E (blue bars). P_u is the probability of unresolved targets. Cooperation, federation and isolated ($k = 1$) cases are compared.

TABLE I
MAE $\sigma(x), \sigma(y)$ [M], PROBABILITY OF UNRESOLVED TARGETS (P_u) AND SIDELINK OVERHEAD B [MB] CONSIDERING T1 AND T2 IN POSITIONS (C, I) AND (G, E).

	$k = 1$	Cooperation	Federation
σ (acc.)	I: 0.08, 0.15 $\sigma(x) \quad \sigma(y)$	I: 0.08, 0.13 $\sigma(x) \quad \sigma(y)$	I: 0.07, 0.15 $\sigma(x) \quad \sigma(y)$
	C: 0.09, 0.16	C: 0.09, 0.07	C: 0.08, 0.09
	G: 0.23, 0.14	G: 0.07, 0.09	G: 0.11, 0.09
	E: 0.12, 0.15	E: 0.08, 0.08	E: 0.1, 0.12
P_u	(C, I): 10%	(C, I): 1%	(C, I): 1%
	(G, E): 59.3%	(G, E): 40.9%	(G, E): 19.8%
B [MB]	(C, I): 0.0	(C, I): 12 Mbit/s 62.500 PCs/s	(C, I): 0.28 Mbit/s 4.400 $\mathbb{P}_k$ /s
	(G, E): 0.0	(G, E): 15.6 MB/s 81.500 PCs/s	(G, E): 0.28 Mbit/s 4.400 $\mathbb{P}_k$ /s

that belong to the reconstructed scene $\hat{\mathbb{S}}_{t,k}$. Fig. 6 shows the probability distribution of estimated positions $\hat{x}_t$ obtained by radar $k = 1$ alone, via cooperation and federation, respectively. Low azimuth resolution of the radar $k = 1$ causes targets in positions G, E be unresolved with large probability (P_u). Cooperation limits these impairments, while federation shows superior performance being less sensitive to the number of observed PCs. For the same scene, Tab. I shows the localization mean absolute error (MAE), namely $\sigma(x) = \mathbb{E}_t[|\hat{x}_t - x|]$, $\sigma(y) = \mathbb{E}_t[|\hat{y}_t - y|]$, the probability of unresolved targets P_u , and the communication overhead B [Mbit/s], quantified here as the sum of bits exchanged over the sidelink per second (localization updates are issued on every 10 ms). The number of exchanged PCs (cooperation) and the mixture model parameters $\mathbb{P}_k$ (federation) per second are also reported.

Cooperation generally improves the average accuracy ($\sigma(x) = 0.09m$) compared with federation ($\sigma(x) = 0.11m$). On the other hand, federation is more effective in minimizing unresolved targets (P_u) as less sensitive to outliers: this effect can be appreciated for co-located subjects G-E. Federation also keeps the required communication overhead

about $20 \div 25$ times lower than cooperation. In cooperation mode, the number of PCs exchanged by each radar depend on the subject position. Subjects in positions C and I make each radar send 625 PCs per localization update on average (62.500 PC/s), or 12 Mbit/s, considering that 3D PCs are 64-bit quantized and formatted in JSON [4]. Subjects in positions G and E are represented on average by 815 PCs per update with required sidelink throughput of 15.6 MB/s. In federation mode, each radar exchanges (on average) 44 local posterior parameters ($\mathbb{P}_k$) per localization update, namely 281.6 Kbit/s. Parameters exchanged contain: the points Q_k measured by radar k and used to represent the federated posterior as in (7), $M_k = 3$ of components (on average), and, for each component, 14 parameters namely, $\beta_{k,m}$, $\mu_{k,m}$ (3 terms), $\Sigma_{k,m}$ (9 terms) and the number of represented PCs $Q_{k,m}$, being $\sum_{m=1}^{M_k} Q_{k,m} = Q_k$.

V. CONCLUSION AND FUTURE ACTIVITIES

The paper proposed and compared cooperation and federation models suitable for distributed radar point cloud processing. The considered radars are characterized by limited range-angular resolution when operating in isolation; however, they exchange side information using a communication channel of limited bandwidth. The paper discusses a prototype and related experiments that quantify the benefits of federation vs. cooperation targeting the localization and tracking of two people in an indoor environment. Federation and cooperation benefits are quantified in terms of latency/sidelink bandwidth use, resolution, probability of unresolved targets and mean absolute localization error performance.

Cooperation generally provides improved average mean absolute localization error (0.09m) than federation (0.11m). Federation forces the radar to exchange model parameters of observed PCs probability (local posterior measure) and it is thus more effective in minimizing unresolved targets and outliers. Federation is a promising solution for scaling up the system to support the deployment of more radars as it keeps the required communication overhead about $20 \div 25$ times lower than cooperation. Required sidelink throughput is also independent from the number of observed PCs. Future activities will consider the impact of synchronization, radar network topologies and heterogeneity in terms of angular resolution and field of view.

REFERENCES

- [1] M. Bennis, et al., "Ultrareliable and Low-Latency Wireless Communication: Tail, Risk, and Scale," *Proc. of the IEEE*, vol. 106, no. 10, pp. 1834–1853, Oct. 2018.
- [2] C. De Lima, et al., "Convergent Communication, Sensing and Localization in 6G Systems: An Overview of Technologies, Opportunities and Challenges," *IEEE Access*, vol. 9, pp. 26902–26925, 2021.
- [3] S. Savazzi, et al., "On the Use of Stray Wireless Signals for Sensing: A Look Beyond 5G for the Next Generation of Industry," *Computer*, vol. 52, no. 7, pp. 25–36, July 2019.
- [4] S. Kianoush, et al., "A Multisensory Edge-Cloud Platform for Opportunistic Radio Sensing in Cobot Environments," *IEEE Internet of Things Journal*, vol. 8, no. 2, pp. 1154–1168, Jan. 2021.
- [5] D. Salami, et al., "Tesla-Rapture: A Lightweight Gesture Recognition System from mmWave Radar Sparse Point Clouds," *IEEE Transactions on Mobile Computing*, 2022.
- [6] S. Palipana, et al., "Pantomime: Mid-Air Gesture Recognition with Sparse Millimeter-Wave Radar Point Clouds," *Proc. of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 5, no. 1, Mar. 2021.
- [7] X. Jiang, et al., "Millimeter-Wave Array Radar-Based Human Gait Recognition Using Multi-Channel Three-Dimensional Convolutional Neural Network," *Sensors*, 20, 5466, 2020. <https://doi.org/10.3390/s20195466>
- [8] Y. Kim, et al., "Human Activity Classification Based on Point Clouds Measured by Millimeter Wave MIMO Radar with Deep Recurrent Neural Networks," *IEEE Sensors Journal*, vol. 21, no. 12, pp. 13522–13529, June, 2021.
- [9] Y. Huang, et al., "Activity Recognition Based on Millimeter-Wave Radar by Fusing Point Cloud and Range-Doppler Information," *Signals* 3, no. 2: 266–283, 2022. <https://doi.org/10.3390/signals3020017>
- [10] I. Baftiu, et al., "Multi-Mode Surround View for ADAS vehicles", *Proc. of 2016 IEEE International Symposium on Robotics and Intelligent sensors (IRIS2016)*, Tokyo, Japan, Dec 2016.
- [11] S. Raj et al., "Bayesian Grouping of Multi Sensor Radar Fusion for effective Pedestrian Classification in Automotive Surround View," *Proc. of IEEE MTT-S International Conference on Microwaves for Intelligent Mobility (ICMIM)*, Linz, Austria, pp. 1–4, 2020. doi: 10.1109/ICMIM48759.2020.9299033.
- [12] S. Savazzi, et al., "Opportunities of Federated Learning in Connected, Cooperative, and Automated Industrial Systems," *IEEE Communications Magazine*, vol. 59, no. 2, pp. 16–21, Feb. 2021.
- [13] B. M. Keel, "Constant False Alarm Rate Detectors," in *Principles of Modern Radar: Basic Principles*, M. A. Richards, et al. (Eds.), SciTech Publishing, 2010.
- [14] L. Barbieri, et al., "Decentralized federated learning for extended sensing in 6G connected vehicles," *Vehicular Communications*, vol. 33, 100396, 2022.
- [15] Y. Cheng, et al., "A New Automotive Radar 4D Point Clouds Detector by Using Deep Learning," *Proc. of IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP'21)*, pp. 8398–8402, 2021.
- [16] K., Rahif, et al., "Federated Generalized Bayesian Learning via Distributed Stein Variational Gradient Descent," *IEEE Transactions on Signal Processing*, vol. 70, pp. 2180–2192, 2022.
- [17] L. Liu, et al., "A Bayesian Federated Learning Framework with Online Laplace Approximation," [Online]. Available: <https://arxiv.org/abs/2102.01936>.
- [18] M. Ester et al., "A density-based algorithm for discovering clusters in large spatial databases with noise," *Proc. of the 2nd Intl. Conf. on Knowledge Discovery and Data Mining*, pp. 226231, Aug. 1996.
- [19] J. Jang, et al., "DBSCAN++: Towards fast and scalable density clustering," *Proc. of the 36th Intl. Conf. on Machine Learning*, no. 97, pp. 3019–3029, 2019. [Online]. Available: <https://arxiv.org/abs/1810.13105>
- [20] P. Jenke, et al., "Bayesian point cloud reconstruction," *Computer Graphics Forum*, vol. 25, pp. 379–388, 2006.
- [21] O. Schall, et al., "Robust filtering of noisy scattered point data," *Proc. Eurographics/IEEE VGTC Symposium Point-Based Graphics*, Stony Brook, NY, USA, pp. 71–144, 2005. doi: 10.1109/PBG.2005.194067.
- [22] M-J. Rakotosaona, et al., "PointCleanNet: Learning to Denoise and Remove Outliers from Dense Point Clouds," *Computer Graphics Forum*, vol. 39, no. 1, 2019. [Online]. Available: <https://arxiv.org/abs/1901.01060>.
- [23] P. Salamon, et al., "A review and numerical assessment of the random walk particle tracking method," *Journal of Contaminant Hydrology*, vol. 87, no. 3–4, pp. 277–305, Oct. 2006.
- [24] A. N. Novick et al., "Practical limitations of NTP time transfer," *Joint Conference of the IEEE International Frequency Control Symposium & the European Frequency and Time Forum*, pp. 570–574, 2015.
- [25] Kullback S, Leibler RA "On information and sufficiency," *Ann Math Stat* 22:79–86, 1951.