

A Belief Propagation Solution for Beam Coordination in MmWave Vehicular Networks

Original

A Belief Propagation Solution for Beam Coordination in MmWave Vehicular Networks / LIMANI FAZLIU, Zana; Malandrino, Francesco; Chiasserini, Carla Fabiana; Nordio, Alessandro. - In: IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS. - ISSN 1536-1276. - STAMPA. - 21:12(2022), pp. 10809-10823. [10.1109/TWC.2022.3187468]

Availability:

This version is available at: 11583/2968287 since: 2022-12-10T08:19:50Z

Publisher:

IEEE

Published

DOI:10.1109/TWC.2022.3187468

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

A Belief Propagation Solution for Beam Coordination in MmWave Vehicular Networks

Zana Limani Fazliu, *Member, IEEE*, Francesco Malandrino, *Senior Member, IEEE*,
Carla Fabiana Chiasserini, *Fellow, IEEE*, Alessandro Nordio, *Member, IEEE*

Abstract—Millimeter-wave communication is widely seen as a promising option to increase the capacity of vehicular networks, where it is expected that connected cars will soon need to transmit and receive large amounts of data. Due to harsh propagation conditions, mmWave systems resort to narrow beams to serve their users, and such beams need to be configured according to traffic demand and its spatial distribution, as well as interference. In this work, we address the beam management problem, considering an urban vehicular network composed of gNBs. We first build an accurate, yet tractable, system model and formulate an optimization problem aiming at maximizing the total network data rate while accounting for the stochastic nature of the network scenario. Then we develop a graph-based model capturing the main system characteristics and use it to develop a belief propagation algorithmic framework, called CRAB, that has low complexity and, hence, can effectively cope with large-scale scenarios. We assess the performance of our approach under real-world settings and show that, in comparison to state-of-the-art alternatives, CRAB provides on average a 50% improvement in the amount of data transferred by the single gNBs and up to 30% better user coverage.

Index Terms—Vehicular networks, mmwave communications, message passing.

I. INTRODUCTION

Vehicular networks and their users have long been identified as great consumers of data, for applications including safety [1], [2], map updates [3], content downloading [4], and onboard entertainment [5]. The issue has been further exacerbated by the emergence of connected and autonomous vehicles: such vehicles need frequently-updated and detailed information on the topology and conditions of the road [6], in addition to providing their users with even richer multimedia content, especially for automated and autonomous vehicles [6]. The effect of such trends is a further increase of the requirement posed on the infrastructure serving the vehicles.

Whenever more network capacity for wireless networks is needed, moving towards higher frequencies is an appealing

Z. Limani Fazliu is with University of Prishtina, Prishtina, Kosovo. F. Malandrino, C. F. Chiasserini and A. Nordio are with CNR-IEIIT, Italy. F. Malandrino and C. F. Chiasserini are with CNIT, Italy. C. F. Chiasserini is with Politecnico di Torino.

This work was supported by the EU Commission under the RAINBOW project (Grant Agreement no. 871403). The views expressed are those of the authors and do not necessarily represent the project. The Commission is not liable for any use that may be made of any of the information contained therein.

This publication was made possible by NPRP-S 13th Cycle grant no. NPRP13S-0205-200265 from the Qatar National Research Fund (a member of Qatar Foundation). The findings herein reflect the work, and are solely the responsibility, of the authors.

option. Indeed, millimeter-wave (mmWave) networks, which operate at frequencies of tens of gigahertz and were originally envisioned for quasi-static, indoor scenarios, are becoming an appealing option also for vehicular networks. However, although mmWave technology allows for very large bandwidths and high data rate, it is also characterized by harsh propagation conditions, with high path loss and virtually no connectivity in non-line-of-sight conditions.

To address these shortcomings, directional antennas capable of *beamforming* are employed. Unlike antennas used at lower frequencies, mmWave base stations (gNBs) serve their users through *beams*, each concentrating the available power along a given direction in order to achieve higher values of received power as well as lower interference [7]. This also means that swift, high-quality beamforming decisions are critical to the performance, and indeed the very usefulness, of mmWave networks. The task of beamforming is especially challenging in vehicular networks, owing to the fast mobility of the users; on the positive side, such mobility is constrained by the road topology and can be forecast with good accuracy [8].

In this work, we formulate the problem of beamforming in vehicular networks as an optimization problem, where the decisions to make concern the beam configuration at each gNB, and the objective is to maximize the total network data rate. Owing to the multiple sources of variability in our scenario and the stochastic nature of the wireless medium, we then introduce a *randomized* solution strategy. Under such an approach, we do not choose directly a beam configuration, but rather (i) we set the *probabilities* of each configuration to be selected, and (ii) we enact the beam configurations over time according to those probabilities. Similar randomized approaches have been successfully used in many fields, including routing [9], resource provision in cloud computing [10], and network orchestration [11]. The intuition behind the success of randomized approaches is that, in very complex scenarios, different sources of randomness tend to cancel one another, in a manner similar to errors in Fermi approximation. Therefore, randomly choosing the beam configuration to enact can actually reduce the negative effects of the variability in vehicular traffic or wireless channel conditions.

We further introduce a distributed heuristic algorithm, called coverage-rate aware belief propagation (CRAB), to make effective and efficient decisions about the beam configuration probabilities. CRAB is based upon the belief propagation approach [12], and leverages scenario-specific information and insights. In CRAB, gNBs exchange *messages* about their *local* beam configurations, until converging to a situation where the

global data rate is maximized. We compare CRAB against state-of-the art approaches based upon clustering or graph matching, and find it to provide remarkably better solutions with a very low complexity.

In summary, our main contributions are as follows:

- We provide a detailed, yet mathematically tractable, model of a mmWave network infrastructure for the support of vehicular communications, which is based upon established standards and cutting-edge research studies. Using this model, we formulate both a centralized and distributed optimization problem making beamforming decisions that maximize the overall network data rate. Unlike previous work and motivated by the highly dynamic scenario under study, we leverage a randomized approach whereby decisions concern the probability with which a given configuration is used at each gNB.
- We then focus on the distributed formulation, so as to make high-quality decisions while exploiting solely local information. In so doing, we develop a graph-based representation of the network infrastructure, which accounts for the inter-gNB conflicts, i.e., interference and overlapping user coverage. Through such a model, we define a novel belief propagation-based approach, named CRAB, that efficiently yields beamforming configurations at the gNBs that effectively avoid inter-gNB conflicts.
- We assess the performance of the CRAB scheme and compare it against state-of-the art alternatives, under real-world settings. Our results show that under CRAB over 30% of the gNBs experience a rate increase of over 100%, and 55% experience a gain of at least 50%, while serving up to 30% more vehicular users and exhibiting higher fairness.

The remainder of this paper is organized as follows. We begin by discussing some relevant related work in Sec. II, highlighting the novelty of our approach, and, then, we introduce our system model in Sec. III. In Sec. IV, we formulate a centralized and distributed beamforming optimization problem. Sec. V presents the CRAB algorithm, highlighting how it combines the mechanics of message-passing algorithms with scenario-specific knowledge and insights. The performance of CRAB is compared against that of state-of-the-art alternatives in Sec. VI, while Sec. VII concludes the paper.

II. RELATED WORK

Beam management in mmWave networks, and in particular initial access, beam alignment and configuration have been active areas of research in the last few years. While earlier works focused on exhaustive and iterative search techniques to identify and configure mmWave beam directions [13], [14], later works turned to more intricate approaches which were often driven by data and based on advanced learning techniques [15]–[19]. In particular, [15] adopts a data-driven approach and uses convolutional neural networks to reduce the necessary coordination between transmitter and receiver when configuring their beam settings. A similar data-driven approach is used in [16] wherein a convolutional neural network architecture is applied in tandem with LIDAR preprocessing

technique to optimize beam selection. The proposed model is trained to exploit LIDAR and positional data in order to identify the best beam directions and reduce the beam search overhead in vehicle-to-infrastructure communication.

In [17], the authors tackle the optimal beam selection problem by formulating the decision-making process as a partially observable Markov decision process. They also propose a point-based value iteration method to design an approximately optimal policy, wherein the goal is to select the strongest beam pair that maximizes the beamforming gain between a single base station-user pair. The study in [18] envisions an optimal beam association policy in mmWave vehicular networks using a lightweight alternative to the Q-learning algorithm, while modeling the dynamics of the mmWave communication link using a semi-Markov decision process framework. [18], however, deals only with straight-road scenarios, with the assumptions that infrastructure nodes cover separate and distinct segments of the road. The algorithm is therefore independently applied by each node to identify the optimal beam association strategy for vehicles under their coverage. No coordination between nodes is foreseen.

To predict the best beam choice for a vehicle, [19] introduces the usage of channel charting. The proposed approach consists of two stages: one offline during which the channel charts are constructed for each beam, and one online, during which live collected data is used to make online predictions for the best beam combination. Again, the work in [19] focuses on a single infrastructure node covering a single straight road segment.

Learning-based techniques, however, are known to be computationally taxing and time consuming, which is why most of these works consider limited scenarios with a single infrastructure node or address the best beam selection with respect to a single vehicle. In addition, they address a highway or a straight road scenario, thereby largely ignoring the interactions and interference potential that can be found in an urban setting.

Furthermore, all of the above studies focus on beam aligning for a single base station-vehicle pair, implicitly assuming that each mmWave beam is employed to transmit to a single user only. However, in ultra dense scenarios, a narrow beam can cover several users simultaneously, users that can be multiplexed within the same beam. There are only few works that consider that *a mmWave link can be used to establish communication with several end users simultaneously* [20], [21]. In particular, [20] considers a dense urban scenario, and uses traffic light information to guide the beam directions chosen by the infrastructure nodes; however, [20] does not consider coordination between nodes as we do in this work. In [21], instead, the authors focus on vehicle-to-vehicle networks and propose an adaptive beamforming scheme based on K-means clustering for point-to-multipoint communications for message dissemination. The work in [21], however, is tailored to highway scenarios and message dissemination therein is enabled by data relaying performed by individual vehicles, and it cannot be easily extended to vehicle-to-infrastructure communication scenarios.

We also remark that, to our knowledge, few works have applied graph theory to address beam management in mmWave

communications [22]–[25]. Both [22] and [23] apply graph techniques to reduce inter-cell interference, which is different from our preliminary work in [24] and this work. We recall that, in [24] as well as this work, the goal of the proposed graph-based approaches is to maximize the network data rate, the difference being that in [24] a centralized approach is proposed, while in this work we consider a distributed approach enabled through the coordination between connected nodes in the graph.

As for message passing applied to mmWave networks, studies that leverage such an approach can be found in [25], [26]. In [25], the authors use a graph approach to tackle the user association and power control in mmwave HetNets, by modeling the network as a coordination graph with edges between base stations and users. Using this graph, they apply a message passing algorithm combined with reinforcement learning to achieve a solution that maximizes the overall time averaged risk averse rate of the network. The authors do not address the management of the beams in such a network, rather they consider the beam pattern of the base station to be fixed and, consequently, users transition from an aligned to a non-aligned state as they move. The work in [26], instead, uses a dynamic compressed sensing-approximate message passing algorithm, to leverage the sparsity and correlation in subchannels for channel estimation and propose an alternative technique that exploits information about the antenna geometry and the range of the transceiver distance, for compressive beam alignment. The message passing algorithm is used to establish an individual short-range link between one access point and its users, and, thus, it does not take into account the behavior of other transmitters in the network.

Novelty. Our work represents an improvement over existing literature along three main directions. First, our graph-based representation of the mmWave infrastructure is a complete and compact way to account for the non-trivial outcome of beamforming decisions, without the intrinsic complexity of data-driven approaches. Second, our distributed message-passing solution strategy allows for swifter convergence compared to centralized algorithms, without the need to share and transfer large amounts of data. Third, by embedding domain-specific knowledge into the messages being passed, we are able to obtain higher-quality solutions compared to general-purpose approaches, e.g., based upon Markov decision processes, which need to blindly “learn” the problem structure.

III. SYSTEM MODEL

To develop a system model that captures all the main aspects of a mmWave vehicular network, we consider a reference scenario based on real-world mobility and infrastructure traces, as per [27], [28]. Such traces contain information about the topology of the city of Luxembourg, the road layout (e.g., regulated intersections), as well as the mobility traces of around several thousands of vehicles traveling within the city center, accumulated over a 12-hour window. Based on this data, we construct a scenario as the one depicted in Fig. 1, in which a set of gNBs, denoted by \mathcal{G} , are co-located with traffic lights to serve a set of vehicles, i.e., the mmwave *users*.




Fig. 1. Real-world scenario: Luxembourg city center. The red circles represent the locations of the traffic lights, i.e., of the gNBs.

The service is divided into a set \mathcal{Z} of discrete zones: each vehicle is, at any given time, into exactly one zone. The set of discrete zones is divided in two distinct subsets, the set of non-empty zones \mathcal{Z}_n and the set of empty zones \mathcal{Z}_e . The subset of non-empty zones is composed of zones which are occupied by at least one vehicle at any given time, while empty-zones are those zones which are not occupied by any vehicles. This allows us to focus the beam optimization efforts only on those zones that require service, and ignore those spatial areas that do not contain any vehicles. Below, we detail the characteristics of our scenario, and present the assumptions we make to build an accurate, yet tractable, system model.

Antenna arrays and spatial signatures: We assume that all network nodes (gNBs and vehicles) have the same height and consider a 2D network topology, which allows for a simpler mathematical analysis and a lighter notation. Indeed, while easily generalizable to 3D, the 2D network model is already able to describe and capture all important features of the system. Both gNBs and vehicles are equipped with a uniform planar array (UPA) of antennas, with the vehicle UPA being capable of analog beamforming only¹. We assume the elements of a UPA arranged in a square grid and spaced by $\lambda/2$, with λ denoting the signal wavelength. In particular, the gNB’s UPA has size $N_t \times N_t$ elements, while the vehicle’s UPA has size $N_r \times N_r$ elements. The surface of a generic UPA is vertically placed; thus, the vector normal to the surface points to the horizon and has azimuth ψ (measured with respect to a global coordinate system).

If the UPA has $N \times N$ antenna elements, then its spatial signature in the direction, defined by the azimuth ϕ (measured with respect to its normal), is the N^2 -size vector

$$\mathbf{s}(N, \phi) = \mathbf{1}_N \otimes \tilde{\mathbf{s}}(N, \phi), \quad (1)$$

where $\mathbf{1}_N$ is a vector of length N with all elements equal to 1, \otimes denotes the Kronecker product, and the n -th component of $\tilde{\mathbf{s}}(N, \phi)$ is given by $[\tilde{\mathbf{s}}(N, \phi)]_n = e^{j\pi n \sin \phi}$.

¹While deriving our numerical results, the direction of the vehicle’s beam is rounded to 1° accuracy, while the beamwidth is fixed to 13° , which is the value that can be obtained with a 8x8 UPA.

Beamforming: A beam, b , is generated by a gNB using a subset of antenna elements from its UPA. In practice, the beam is obtained by coordinating (phasing) the signals emitted by the antenna elements so that they globally act as a single directional antenna whose main radiation lobe is characterized by a specific half-power beam width (HPBW), α , and direction δ . Beams with different HPBW can be obtained by varying the number of antennas involved in the beam generation. Specifically, if the array generating the beam b has size $N \times N$ elements, the HPBW, α , depends on N through the relation [29]:

$$\alpha \approx \frac{1.78}{N} [\text{rad}] \approx \frac{102^\circ}{N}. \quad (2)$$

The direction of the beam can be set by properly choosing the *beamforming vector* \mathbf{v} i.e., the vector of phases applied to the array elements in order to emit the considered beam, given by

$$\mathbf{v} \triangleq \frac{1}{N} \mathbf{s}(N, \varphi) \quad (3)$$

where the angle φ represents the direction of the generated beam, in the azimuth plane, with respect to the normal to the gNB UPA. In practical systems, the values that the angle φ can take are limited to a discrete set. The direction of the beam in a global horizontal coordinate system is then defined by the angle $\delta = \varphi + \psi$, where we recall that ψ is the azimuth of the normal to the gNB UPA.

Remark. In our scenario, we consider each gNB equipped with a single UPA and able to generate beams with arbitrary direction. However, in a practical scenario, UPA hardware limitations prevent a single beam to span the entire azimuth range $[0, 2\pi]$. To do so, gNBs should be equipped with multiple UPAs, each one covering a sector. Also, on the vehicle side, UPA could be integrated into the vehicle roof whereas in our scenario we consider it vertically placed. Our assumptions, however, have the advantage of allowing for a simpler and lighter mathematical notation, while being general enough to encompass practical scenarios. Specifically, in our study a beamforming configuration depends only upon the number of beams available at each gNB, along with their direction and beamwidth, regardless of the number or position of the UPAs they are generated by.

We denote by \mathcal{B} the set of possible beams, which is common to all gNBs and whose cardinality is equal to the number of possible beam directions multiplied by the number of HPBW values available at the gNB. The set \mathcal{B} contains also an extra element, i.e., the null-beam denoted by \emptyset . In the following, the notation $b = \emptyset$ denotes that the beam b is not emitted.

Let us define as B the maximum number of beams that can be activated at a gNB g ($g = 1, \dots, G$), and denote with vector $\mathbf{b}_g = [b_{g,1}, \dots, b_{g,B}]$ the generic beam-configuration adopted at gNB g . The i -th beam, $i = 1, \dots, B$, is characterized by direction $\delta_{g,i}$ and HPBW $\alpha_{g,i}$. However, not all beam-configurations \mathbf{b}_g are possible since two main constraints need to be accounted for. First, the sum of the number of antennas elements simultaneously involved in the beam generation should not exceed the number of antenna elements of the gNB's UPA (i.e., N_t^2). Second, we must ensure that beams do not overlap with each other, i.e., for every two beams

$b_{g,i}, b_{g,j} \neq \emptyset$ simultaneously emitted on the same frequency band by gNB g , the following condition must hold:

$$|\delta_{g,i} - \delta_{g,j}| \geq \frac{\alpha_{g,i} + \alpha_{g,j}}{2}. \quad (4)$$

We then define $\mathcal{F} \subseteq \mathcal{B}$ as the set of feasible beam configurations at a generic gNB. Finally, a network beam configuration, \mathbf{B} , can be described as an array of G beam-configurations, one for each gNB, i.e., $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_g, \dots, \mathbf{b}_G] \in \mathcal{F}^G$.

Transmitted signal: Radio resources available for communication are organized into an $N_b \times N_\tau$ frequency-time matrix, whose time size is called frame. A single element of such matrix, named resource block, is characterized by a bandwidth W and a time fraction $\tau = 1/N_\tau$. The total bandwidth assigned for communication is then $N_b W$. We assume that every beam transmits over all resource blocks; therefore, given a network beam configuration, \mathbf{B} , the baseband signal transmitted by the i -th beam of gNB $g \in \mathcal{G}$, in resource block q can be modeled by vector:

$$\mathbf{t}_{g,i,q} = \mathbf{v}_{g,i} x_{g,i,q} \quad (5)$$

where $\mathbf{v}_{g,i}$ is the beamforming vector in (3) particularized to beam $b_{i,g}$, and $x_{g,i,q}$ is a complex random symbol with zero mean. Further, assuming a uniform power allocation over all resource blocks, we associate to $x_{g,i,q}$ a power equal to: $\mathbb{E}[|x_{g,i,q}|^2] = P_{g,i}/N_b$. Note that the available transmit power at gNB g , P_g , is shared among the beams simultaneously emitted therein. Therefore, the values $P_{g,i}$ are subject to power allocation constraints that also depend on the adopted beam configuration.

Mmwave communication channel: In a typical mmwave urban scenario, the channel between a gNB g and a zone z in resource block q , can be modeled as described in [30]–[32]. Such models consider $L_g(z)$ clusters of paths, each described by a complex coefficient, $h_{g,\ell,q}(z)$, and two angles, $\phi_{g,\ell}(z)$ and $\theta_{g,\ell}(z)$, $\ell = 1, \dots, L_g(z)$, which represent, respectively, the departure and arrival direction of the signal, measured with respect to the normal to the transmitting and receiving UPAs. Since we consider zones to be sufficiently small so that vehicles therein (if any) experience the same propagation channel for a given beam and gNB, in the following we associate a zone with a unique UPA and receiver, and we detail the channel model by referring to a gNB-zone communication link. For a given network beam configuration, \mathbf{B} , the time-domain channel experienced by beam $b_{g,i}$, connecting gNB g to zone z , in resource block q is given by the matrix:

$$\mathbf{H}_{g,i,q}(z) = \sqrt{\frac{1}{L_g(z)}} \sum_{\ell=1}^{L_g(z)} h_{g,\ell,q}(z) \mathbf{u}_{g,\ell}(z) \boldsymbol{\mu}_{g,i,\ell}(z)^H \quad (6)$$

where $\boldsymbol{\mu}_{g,i,\ell}(z) \triangleq \mathbf{s}(N_{g,i}, \phi_{g,\ell}(z))$ and $\mathbf{u}_{g,\ell}(z) \triangleq \mathbf{s}(N_\tau, \theta_{g,\ell}(z))$ are the signatures of the transmit and receive antenna arrays, and $N_{g,i}^2$ is the number of antenna elements used by beam $b_{g,i}$. Note that the channel model in (6) is frequency-flat since we refer to a 5G scenario where the bandwidth of a resource block is about 1.4 MHz, while the channel delay spread is in the range of 30 to 300 ns [33].

Received signal: For a given network beam configuration $\mathbf{B} \in \mathcal{F}^G$, the signal carried by beam $b_{g,i}$ in resource block q and received within zone z can be represented by the following $N_r^2 \times 1$ vector:

$$\mathbf{y}_{g,i,q}(z) = \mathbf{H}_{g,i,q}(z)\mathbf{t}_{g,i,q} + \boldsymbol{\eta}_{g,i,q}(z) \quad (7)$$

where $\mathbf{H}_{g,i,q}(z)$ and $\mathbf{t}_{g,i,q}$ are given by (6) and (5), respectively, and $\boldsymbol{\eta}_{g,i,q}(z)$ is a term accounting for noise and interference in zone z , e.g., caused by beams generated by nearby gNBs (in general, all beams in \mathbf{B} except for $b_{g,i}$). We assume that vector $\boldsymbol{\eta}_{g,i,q}(z)$ has complex Gaussian independent entries with zero mean and covariance $\mathbb{E}[\boldsymbol{\eta}_{g,i,q}(z)\boldsymbol{\eta}_{g,i,q}(z)^H] = (N_0W + I_{g,i,q}(z))\mathbf{I}$ where N_0 is the thermal noise power spectral density, W is the signal bandwidth, and $I_{g,i,q}(z)$ is the interference power. Since only analog beamforming is possible at the receiver, the receiver applies to $\mathbf{y}_{g,i,q}(z)$ the vector of weights $\mathbf{w}(z)$ given by $\mathbf{w}(z) \triangleq \frac{1}{N_r}\mathbf{s}(N_r, \varphi(z))$. Notice that such vector has norm 1 and allows the receiver to generate a beam in the direction specified by the angle $\varphi(z)$. After weighting the UPA output, the receiver obtains

$$\mathbf{w}^H(z)\mathbf{y}_{g,i,q}(z) = \tilde{h}_{g,i,q}(z)x_{g,i,q} + \mathbf{w}^H(z)\boldsymbol{\eta}_{g,i,q}(z) \quad (8)$$

where the scalar $\tilde{h}_{g,i,q}(z) = \mathbf{w}^H(z)\mathbf{H}_{g,i,q}(z)\mathbf{v}_{g,i}$ represents the overall communication channel summarizing the effects of the signal propagation and of the antenna and beam design. The term $\mathbf{w}^H(z)\boldsymbol{\eta}_{g,i,q}(z)$ is a Gaussian random variable with zero mean and variance $N_0W + I_{g,i,q}(z)$, and the received signal power is given by $\mathbb{E}[|\tilde{h}_{g,i,q}(z)x_{g,i,q}(z)|^2] = P_{g,i}|\tilde{h}_{g,i,q}(z)|^2/N_b$. Note that the values of the transmit power, $P_{g,i}$, and of channel $\tilde{h}_{g,i,q}(z)$ depend on the specific network beam configuration, \mathbf{B} . To stress this dependence, in the following, we will write $P_{g,i}(\mathbf{B})$ and $\tilde{h}_{g,i,q}(z, \mathbf{B})$.

Serving and interfering beams: Each zone $z \in \mathcal{Z}_n$, is allocated a set of resource blocks, $\mathcal{Q}(z)$, and a set of serving beams. In turn, a beam can serve one or more zones, depending on its direction and HPBW. Specifically, for a given network beam configuration, \mathbf{B} , the set of beams *serving* zone z is denoted by $\mathcal{S}(z, \mathbf{B})$. In other words, the elements of $\mathcal{S}(z, \mathbf{B})$ are pairs (g, i) ; and $(g, i) \in \mathcal{S}(z, \mathbf{B})$ if, under network beam configuration \mathbf{B} , the i -th beam emitted by gNB g serves zone z . Note that, in CoMP-like communications, we can have $1 \leq |\mathcal{S}(z, \mathbf{B})| \leq G_c$ where G_c is the maximum number of gNBs that can partake in the coordinated transmission. Clearly, when no CoMP is enabled, $G_c = 1$.

Similarly, $\mathcal{I}_q(z, \mathbf{B})$ denotes the set of 2-tuples (g, i) identifying the beams $b_{g,i}$ *interfering* in zone z , under network beam configuration \mathbf{B} and resource block q . By using these definitions, the received signal and interference powers for the resource block $q \in \mathcal{Q}(z)$ in zone z can be written, respectively, as

$$\begin{aligned} \tilde{P}_q(z, \mathbf{B}) &= \left| \sum_{(g,i) \in \mathcal{S}(z, \mathbf{B})} \sqrt{\frac{P_{g,i}(z, \mathbf{B})}{N_b}} \tilde{h}_{g,i,q}(z, \mathbf{B}) \right|^2 \\ I_q(z, \mathbf{B}) &= \sum_{(g,i) \in \mathcal{I}_q(z, \mathbf{B})} \frac{P_{g,i}(z, \mathbf{B})}{N_b} |\tilde{h}_{g,i,q}(z, \mathbf{B})|^2 \end{aligned}$$

TABLE I
MAIN NOTATION

Variable	Description
\mathcal{G}	set of gNBs
\mathcal{Z}_n	set of non-empty zones
$\mathcal{Q}(z)$	set of resource blocks assigned to zone z
$\mathcal{F} \subseteq \mathcal{B}$	set of feasible beam configurations in a gNB
\mathbf{B}	network beam configuration
\mathbf{b}_g	beam configuration at gNB g
$\alpha_{g,i}$	HPBW of the i -th beam when the local configuration \mathbf{b}_g is adopted at gNB g
$\delta_{g,i}$	direction of the i -th beam when the local configuration \mathbf{b}_g is used at gNB g
$\Pi(\mathbf{B})$	probability that the network-wide configuration \mathbf{B} is adopted
$\pi_g(\mathbf{b}_g)$	probability that the local configuration \mathbf{b}_g is adopted at gNB g
$T(\mathbf{B})$	network data rate achieved under beam configuration \mathbf{B}
$\chi^{(g,h)}(\mathbf{b}_g, \mathbf{b}_h)$	joint compatibility function between gNBs g and h when configurations \mathbf{b}_g and \mathbf{b}_h are adopted

and the achievable rate as

$$R_q(z, \mathbf{B}) = W\tau \log_2 \left(1 + \frac{\tilde{P}_q(z, \mathbf{B})}{N_0W + I_q(z, \mathbf{B})} \right). \quad (9)$$

Given \mathbf{B} , the total network data rate is given by:

$$T(\mathbf{B}) = \sum_{z \in \mathcal{Z}_n} \sum_{q \in \mathcal{Q}(z)} R_q(z, \mathbf{B}). \quad (10)$$

IV. A RANDOMIZED APPROACH TO NETWORK THROUGHPUT MAXIMIZATION

Given the set of gNBs, \mathcal{G} , the maximum number of supported beams at each gNB, B , and the set of zones, \mathcal{Z} , our goal is to determine the best beam configuration to be used at each gNB. More specifically, we aim at jointly addressing the following questions while maximizing the overall network data rate:

- (i) how many beams, of what width and direction, each gNB should set up, and
- (ii) which zones should be associated to which gNB, and scheduled on which beam.

Owing to the highly dynamic scenarios we target and to the stochastic nature of the wireless channel, we adopt a *randomized approach*, whereby (i) decisions concern the *probability* with which a given configuration is adopted, and (ii) the actual configuration to enact is chosen according to such probabilities. Randomized approaches similar to ours have been long used in many fields, including routing [34], [35], resource allocation [36], and security [37]. It follows that, in our problem formulation (i) the decision variables correspond to *probabilities* that each beam configuration is

adopted, and (ii) the objective function is the *expected value* of the actual target metric. Below, first we present a centralized formulation and discuss its complexity, then we introduce a distributed version of the problem that allows for local decision making based on local information. The main notation we use in our formulations is summarized in Tab. I.

A centralized formulation: Let us first formulate a centralized optimization problem that aims at a globally optimal solution. Such a problem needs to be time-dependent and solved periodically². Upon solving the problem, the solution is represented by the probability values associated to each beam configuration, which are delivered to the set of gNBs. The gNB nodes set their beams according to the received configuration probabilities, by devoting to each of them a fraction of time that is proportional to the probability values. Since the problem formulation holds at every decision period, to simplify the notation, in the following we do not highlight the dependency on time.

Let $\Pi(\mathbf{B})$ represent the probability that network beam configuration \mathbf{B} is selected. Then the total network data rate is given by:

$$\mathbb{E}_{\Pi}[T] = \sum_{\mathbf{B} \in \mathcal{F}^G} \Pi(\mathbf{B})T(\mathbf{B}) \quad (11)$$

and our problem can be formulated as:

$$\max_{\Pi(\mathbf{B})} \mathbb{E}_{\Pi}[T(\mathbf{B})]. \quad (12)$$

Note that (12) is maximized when $\Pi(\mathbf{B}) = 1$ for $\mathbf{B} = \mathbf{B}^{\max}$ and 0 otherwise, where $\mathbf{B}^{\max} = \arg \max_{\mathbf{B}} T(\mathbf{B})$. The optimal beam configuration \mathbf{B}^{\max} is however difficult to compute due to the interference among beams, the need for centralized global knowledge, and the non-linear nature of the problem. Indeed, non-linear problems are notoriously complex to solve, and solution strategies only find local optima in the general case.

A decentralized formulation: In light of the issues above, we are interested in a decentralized solution where each gNB independently chooses its own beam configuration according to a local probability distribution $\pi_g(\mathbf{b}_g)$. To this end, we consider as objective function the following expression:

$$\mathbb{E}_{\pi_1, \dots, \pi_G}[T(\mathbf{B})] = \sum_{\mathbf{b}_1 \in \mathcal{F}} \dots \sum_{\mathbf{b}_G \in \mathcal{F}} \prod_{g=1}^G \pi_g(\mathbf{b}_g) T([\mathbf{b}_1, \dots, \mathbf{b}_G]) \quad (13)$$

and write our problem as:

$$\max_{\{\pi_g\}} \mathbb{E}_{\pi_1, \dots, \pi_G}[T(\mathbf{B})]. \quad (14)$$

Under this decentralized approach, once a decision is made, each gNB g will implement the optimal policy by applying in each of the K time steps one of the possible configurations, with such probability as specified by the densities $\pi_g(\mathbf{b}_g)$. Importantly, the following holds:

Lemma 1: The joint probability obtained using the solutions to (14) coincides with the optimal solution to (12).

²The extension to an event-driven decision making process is however straightforward.

Proof: The proof comes from the well-known fact that, if variables are independent, then the joint distribution of a set of random variables is equal to the product of individual, marginal distributions. In our scenario, independence is guaranteed by the fact that each individual gNB g performs its own randomized strategy, i.e., chooses the concrete strategy to enact according to π_g , with no influence from other gNBs. Then, it is sufficient to observe that the marginal densities maximizing (13) are given by $\pi_g(\mathbf{b}_g) = 1$ for $\mathbf{b}_g = \mathbf{b}_g^{\max}$ and 0 otherwise, for $g = 1, \dots, G$ and $[\mathbf{b}_1^{\max}, \dots, \mathbf{b}_G^{\max}] = \mathbf{B}^{\max}$, as defined in the centralized approach. ■

Moving from the maximization over the joint distribution in (12) to the one over the marginal distributions in (14) does not, *per se*, change the complexity of the problem, nor its solution. However, it does provide us with valuable insights on a possible solution strategy, namely, one leveraging *belief propagation* (BP) algorithms.

In general, BP allows multiple agents to cooperatively estimate the marginal distributions of a set of random variables. Specifically, each agent is associated with a random variable, and an *influence graph* expresses which variables (hence, which agents) influence one another. The algorithm works iteratively, with agents that are neighbors on the influence graphs exchanging messages (indeed, BP belongs to the family of *message passing* algorithms). BP is guaranteed to converge to the centralized solution if the influence graph is a tree, but works remarkably well under a much wider set of conditions.

In our case, agents correspond to gNBs, the random variables to estimate are the local decisions $\pi_g(\mathbf{b}_g)$, and messages express the extent to which decisions of different gNBs conflict with one another. By leveraging BP, we are able to make swift, high-quality decisions that (i) for particular network topologies, *match* the optimal ones, i.e., the ones we would obtain by solving (12), and (ii) in general, are very close to the optimal decisions in most practical cases.

V. CRAB: COVERAGE-RATE AWARE BELIEF PROPAGATION

Our solution strategy, named Coverage-Rate Aware Belief propagation (CRAB), is predicated on allowing each gNB to make *local* decisions about the number, direction, and width of its beams. Such decisions are aimed to avoid *conflicts* among distributed beam management decisions, thus yielding high serving rate and coverage to vehicular users. Indeed, in our scenario, improving the coverage of users is a very good way towards optimizing the total data rate (10), as good coverage implies less interference and more served users.

As discussed above, we follow a *randomized* approach, where the decision variables are represented by probability distributions; specifically, we associate to each gNB g and beam configuration \mathbf{b}_g a probability $\pi_g(\mathbf{b}_g)$. Then, at every time step k , each gNB adopts one of the possible configurations with probability $\pi_g(\mathbf{b}_g)$. Such decisions are local in nature, therefore, we want to allow each gNB to choose the distributions $\pi_g(\mathbf{b}_g)$ that maximize its performance. At the same time, we have to avoid *conflicts* among decisions made by different gNBs. We define a conflict as two or more beams

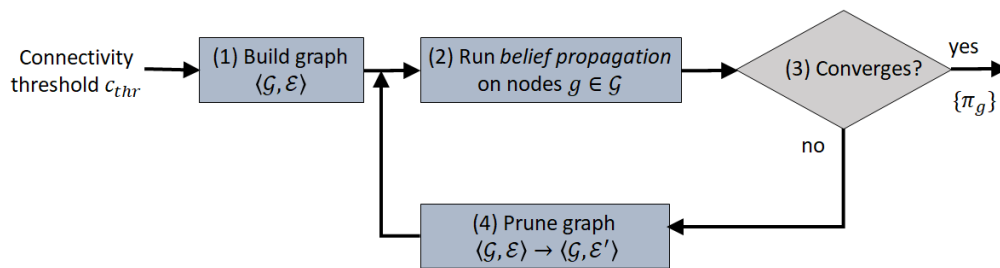


Fig. 2. Schematic representation of the CRAB framework: (1) the interaction graph is built using the given connectivity threshold as detailed in Sec. V-A; (2) the *belief propagation* algorithm is run on the nodes of the graph, as described in Sec. V-B; (3) if the algorithm does not converge within a maximum number of iterations, in step (4) the graph is pruned as described in Sec. V-C. Steps (2)-(4) are repeated until convergence is reached and the marginal probabilities π_g are obtained.

of distinct gNBs, not involved in CoMP communications, covering the same zone(s): such a situation is doubly wasteful, as (i) beams may interfere with each other, and (ii) one of the them could cover different, hitherto unserved, zones instead. Note that the system would greatly benefit from anticipating such conflicts and avoiding them, as detecting a conflict a posteriori would entail that the involved gNBs have to change their configuration so as to remove the conflict itself, with the intrinsic overhead as well as the risk to create a new conflict.

We tackle this conundrum by:

- creating an *interaction graph*, modeling the mmwave network and summarizing which gNBs' decisions may conflict with each other;
- apply a belief propagation algorithm to solve the distributed problem introduced in Sec. IV;
- if the algorithm does not converge, *pruning* such a graph and re-applying the belief propagation algorithm until convergence is reached and a solution is obtained.

These steps are depicted in the general scheme of the CRAB process in Fig. 2.

Below, we first present how the interaction graph is built, in such a way that the mmwave network characteristics are accounted for (Sec. V-A). Then, we associate with each gNB a *state*, i.e., a beam configuration, and, given the objective in (14), we design a belief propagation algorithm which applies the message passing approach and yields a probability distribution over the local beam configuration decisions (Sec. V-B). Finally, we detail how the interaction graph can be pruned if necessary, so as to make the algorithm reach stable and mutually beneficial solutions to be adopted at the gNBs (Sec. V-C).

A. Building the interaction graph

To address the beam management problem, we model the network as a graph, composed of a set of nodes coinciding with \mathcal{G} , i.e., with each node representing a single gNB, and a set of edges \mathcal{E} between said nodes. As exemplified in Fig. 3, an edge $e_{g,h}$ between nodes g and h exists if the value of the connectivity criterion between the two vertices, $c_{g,h}$, is above a certain threshold c_{thr} , i.e.,

$$\mathcal{E} = \{e_{g,h} : \forall g, h \in \mathcal{G} \wedge c_{g,h} > c_{thr}\}.$$

As detailed below, such criterion should reflect the gNB's potential to interfere with each other's transmissions and can include, but it is not limited to, the inter-gNBs distance, the existence of line-of-sight (LoS) conditions between gNBs, or the gNBs coverage overlapping in terms of number of non-empty zones.

Examples of how the graph of the mmWave network in Luxembourg city would look like using different criteria is shown in Fig. 4. In the leftmost figure, the value of the connectivity criteria is expressed as a function of the distance between the nodes, i.e., $c_{g,h} = \frac{1}{d_{g,h}}$, and the threshold is related to a fixed distance $c_{thr} = \frac{1}{d_{thr}}$ where $d_{thr} = 400$ m, as it has been shown that the probability of maintaining a mmwave communication link over longer distances is negligible [31]. In the center figure, the criteria is the existence of LoS between two nodes, which can be expressed as $c_{g,h} = \text{LoS}(g, h)$, where $\text{LoS}(g, h)$ is a binary variable indicating whether there is LoS between the two nodes, and $c_{thr} = 0$. Finally, in the right-hand plot the connectivity criteria is based on the fraction of the number of non-empty zones under the coverage of the two gNBs with respect to the total number of zones covered by the two gNBs separately, and the threshold values has been set to $c_{thr} = 0.1$.

The various criteria can also be combined, and identifying the most suitable criterion to be considered is one of the goals of this work. It should be noted, however, that while the first two criteria are topology based, and we can assume that the structure of the graph will be fixed, the third one depends on the vehicle traffic patterns since the number of non-empty zones may change over time, as traffic patterns shift.

B. Configuring the belief propagation algorithm

Belief propagation is an algorithm allowing to infer the (local) marginal distributions of a set of random variables taking into account their mutual correlation. The algorithm is defined over graphs where each random variable is associated to a node, and works by letting *messages* flow along the graph edges. A message is a real valued function that measures the influence that a random variable (i.e., a node) exerts on the neighboring ones. The algorithm works iteratively and, at every iteration, the marginal distributions of the random variables at each node are computed and updated, until convergence is reached.

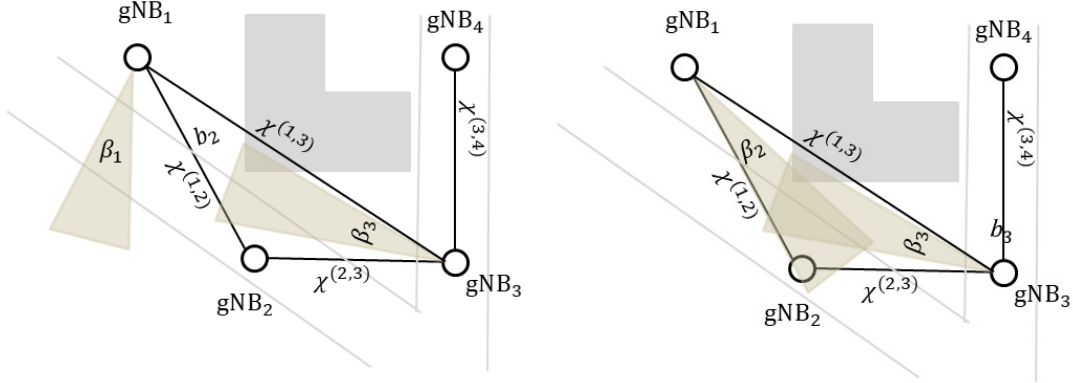


Fig. 3. Examples of interaction graph modeling, highlighting two possible beam configurations at each gNB (i.e., graph node).

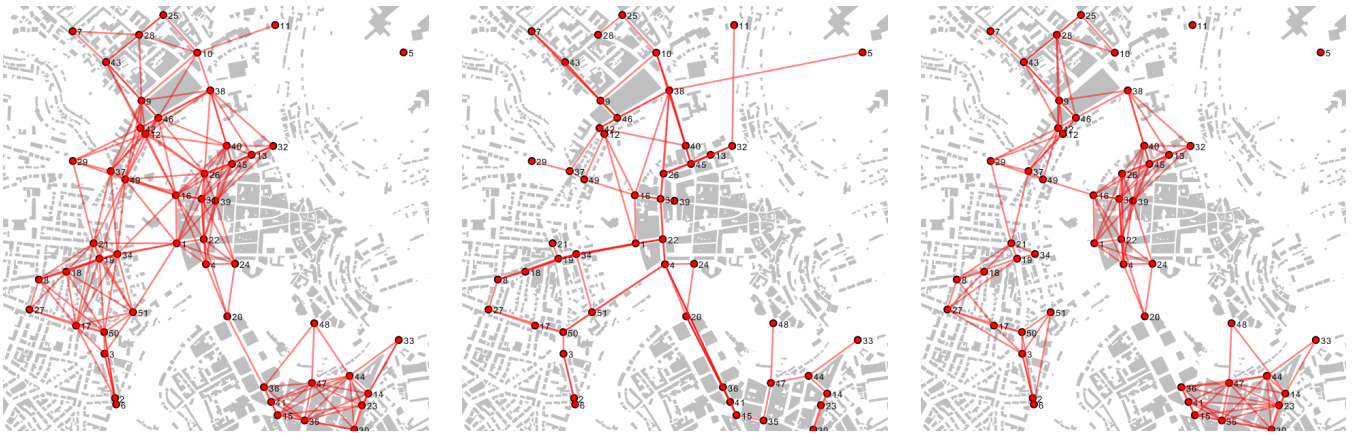


Fig. 4. Mmwave network graph using different connectivity criteria: (left) distance-based criterion with $c_{g,h} = \frac{1}{d_{g,h}}$ and $d_{g,h} = 400$ m; (center) LoS-based criterion; (right) overlapping coverage criterion with $c_{\text{thr}} = 0.1$.

In our scenario the graph is the network interaction graph obtained under one of the connectivity criteria mentioned in Sec. V-A. The random variables are the beam configurations at the nodes, \mathbf{b}_g ($g = 1, \dots, G$), and their marginal distributions are the $\pi_g(\mathbf{b}_g)$'s previously introduced.

Joint compatibility function: The interaction among random variables is described by the joint compatibility function $\chi^{(g,h)}(\mathbf{b}_g, \mathbf{b}_h)$, which, for every pair of nodes $g, h \in \mathcal{G}$, measures the compatibility between beam configurations \mathbf{b}_g and \mathbf{b}_h when they are simultaneously activated by nodes g and h , respectively.

We design our compatibility function based upon the intuition that two configurations *interfere* with each other when they cover the same set of users. Such a situation hurts the objective in (14) in two ways, namely:

- it results in fewer users being served, hence, the total data rate decreases;
- it creates interference for the users that do get served, further reducing the total data rate.

For the above reasons, improving *coverage* is strongly linked with increasing the total data rate in (10).

For example, if we look at Fig. 3 where the maximum number of activated beams is $B = 1$, we expect that the

compatibility value between the beam configurations $b_1 = \beta_1$ at node gNB_1 and beam $b_3 = \beta_3$ at node gNB_3 , to be higher than the value obtained when $b_1 = \beta_2$ at node gNB_1 and $b_3 = \beta_3$ at node gNB_3 , i.e., $\chi^{(1,3)}(\beta_1, \beta_3) > \chi^{(1,3)}(\beta_2, \beta_3)$, since in the first case the beams do not interfere with each other. Accordingly, we define the joint compatibility function as:

$$\chi^{(g,h)}(\mathbf{b}_g, \mathbf{b}_h) = \sum_{z \in \mathcal{Z}_n} \sum_{q \in \mathcal{Q}(z)} R_q(z, \mathbf{B}^{(g,h)}) \quad (15)$$

where³ $\mathbf{B}^{(g,h)} = [\emptyset, \dots, \emptyset, \mathbf{b}_g, \emptyset, \dots, \emptyset, \mathbf{b}_h, \emptyset, \dots, \emptyset]$, $\emptyset = [\emptyset, \dots, \emptyset] \in \mathcal{F}$ is the null beam set, and $R_q(z, \mathbf{B})$ is the rate defined in (9). Note that, in order to calculate $R_q(z, \mathbf{B})$, we require the knowledge of the channels linking the gNB to the users. Since beam configuration decisions are made for longer time periods⁴, instead of instantaneous channel information, CRAB uses averages of the square magnitude of the channel coefficients $\tilde{h}_{g,i,q}$. Such averages are calculated both over a longer period of time, as well as over the area of the zone.

In practice, $\chi^{(g,h)}(\mathbf{b}_g, \mathbf{b}_h)$ returns the rate achieved by

³Without loss of generality, we assume $g < h$.

⁴In highly mobile scenarios the channel coherence time is very small, e.g. few ms whereas the CRAB algorithm requires to be run at much lower rate, e.g. every second.

the network when beam-configurations \mathbf{b}_g and \mathbf{b}_h have been selected at gNBs g and h , respectively, and all other gNBs are silent. It is interesting to observe how the compatibility function in (15) is not directly derived from the objective (14); rather, it incorporates domain- and scenario-specific knowledge about which situations ought to be avoided in order to improve performance.

Message passing process. As mentioned above, the belief propagation algorithm works by exchanging messages along graph edges as shown in Fig. 5. Let us denote by $m^{g \rightarrow h}(\mathbf{b}_h)$ the message that node g sends to node h about beam configuration, \mathbf{b}_h .

Once the interaction graph is built, we define $\nu(g)$ as the set of neighbors of node g in the graph. We then design a belief propagation algorithm that, thanks to the compatibility function defined above, yields probabilities $\pi_g(\mathbf{b}_g)$'s that maximize the network data rate. Specifically, at each iteration, a node g computes the messages to be sent to each of its neighbors, h , and for each of the neighbor's beam configuration \mathbf{b}_h , according to:

$$m^{g \rightarrow h}(\mathbf{b}_h) = \sum_{\mathbf{b}_g \in \mathcal{F}} \chi^{(g,h)}(\mathbf{b}_g, \mathbf{b}_h) \prod_{k \in \nu(g) \setminus h} m^{k \rightarrow g}(\mathbf{b}_g) \quad \forall h, \mathbf{b}_h \quad (16)$$

where $m^{k \rightarrow g}(\mathbf{b}_g)$ is the last message received by g from k about its beam configuration \mathbf{b}_g . At the first iteration of the algorithm messages are initialized to a constant, arbitrary chosen value.

In the first iteration all the nodes send their initialized messages to the respective outgoing nodes. In the following iterations, a node will send an outgoing message, once it has received all incoming messages necessary to compute it⁵. An iteration is complete once the node calculates and sends all of its outgoing messages. Further, notice that since the joint compatibility function can take any value, we must ensure a proper normalization of the values $\pi_g(\mathbf{b}_g)$ at each iteration of the algorithm. Such normalization is obtained by imposing:

$$C_g = \sum_{\mathbf{b}_g \in \mathcal{F}} \prod_{h \in \nu(g)} m^{(h \rightarrow g)}(\mathbf{b}_g), \quad \forall g = 1, \dots, G.$$

Convergence is reached when the difference between all consecutively outgoing messages is negligible. If convergence is not reached within a maximum number of iterations, we proceed as explained in Sec. V-C. Upon convergence, the estimated marginal distributions of the beam configuration at node g is given by [12]:

$$\pi_g(\mathbf{b}_g) = \frac{1}{C_g} \prod_{h \in \nu(g)} m^{(h \rightarrow g)}(\mathbf{b}_g) \quad (17)$$

for all $\mathbf{b}_g \in \mathcal{F}$. Each node can then locally decide regarding the beam configuration by randomly selecting the state according to the obtained marginal probability.

C. gNB graph pruning

Belief propagation does not guarantee convergence in graphs that contain cycles [12], [38], and, in general, the

⁵We consider that inter-gNB communication is enabled through reliable links, hence no messages are lost during the exchange.

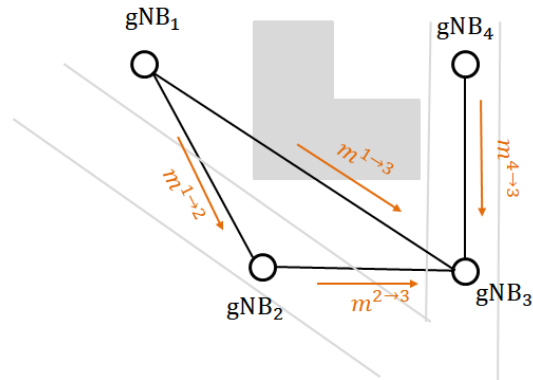


Fig. 5. An example of message passing in the CRAB framework.

conflict graph of a mmwave network is not guaranteed not to contain cycles. The application of belief propagation in such graphs, commonly referred to as *loopy belief propagation*, is known to converge in most practical cases, although no guarantee of convergence can be provided [38]. Rather, the existing relevant literature provides upper bounds conditions under which the graph will converge, which are quite loose and often pessimistic compared to actual performance [38], as confirmed by our own results.

In general, the convergence in such graphs depends heavily on the irregularity of the geometry of the graph itself and the variability in the strength, i.e., the range of the values of the compatibility function along the different edges, as shown in [38]. It follows that the criteria used to determine the connectivity between the nodes plays a significant role. As already discussed, in the case of mmwave networks, several factors can be taken into account when determining whether two gNBs are connected; nonetheless, it is quite apparent that in an urban scenario, loops are unavoidable. To ensure convergence, we therefore take the following approach: if the graph does not converge within a maximum number of iterations, we prune the graph by using Kruskal's reverse delete algorithm [39], which removes the least significant edge that does not affect the overall connectivity of the graph. In the worst case, the pruning process leads to a tree-like graph for which convergence is guaranteed [12], but our results (see Sec. VI) show that usually the graph converges much earlier than that. In the rare instances in which it does not converge, only a few prunings are required to obtain a graph that does converge.

VI. NUMERICAL RESULTS

We evaluate the effectiveness of our approach by considering the real-world network layout of the Luxembourg city center, as described in Sec. III, and the realistic vehicular mobility trace in [27]. The 2×2 km² service area is divided into a discrete set of zones, each of 10×10 m² size.

The system parameters are configured as follows. The center frequency available for mmwave communication is set to 52 GHz, while the available bandwidth is $W = 400$ MHz. The latter corresponds to the maximum allowed bandwidth in 5G New Radio (NR) using numerology $\mu = 3$ with subcarrier

spacing of 120 kHz and 264 resource blocks [33]. Further, we assume that all gNBs are equipped with a 64×64 uniform planar array (UPA) with up to 4 RF chains transmitting at a maximum power of $P_g = 33$ dBm, while users are equipped with a 8×8 UPA with a single RF chain. To simulate the mmwave channel, we use the statistical approximation of the 3GPP channel model accounting for the Doppler effect, shadowing and multipath fading, and we set the large-scale parameters used for modeling as in [20].

Beam directions can take any integer value between 0° and 360° , but we limit the number of possible beamwidth configurations to $\{5^\circ, 10^\circ, 15^\circ\}$. The maximum number of beams at the gNB is limited by the number of available RF chains, therefore $B = 4$ for all gNBs. The CRAB algorithm is executed every second. We assume that channel information is updated every ms, therefore the time averages of the channel gains are computed over a 1-second time window, i.e., over 1,000 realizations of the channel.

Due to the level of complexity required to simulate a full scale network consisting of 51 gNBs and over 2,500 vehicles, the performance of the proposed approach is evaluated using a single simulation that is 10s long. The longer simulation period was adopted to capture the realistic movement of the vehicles in an urban environment. We assume that all vehicles in the network are requesting data for the entire duration of the simulation (the so-called full buffer model), and we focus on the downlink performance of the network. The resource allocation is performed according to the proportional fair algorithm, while user association is based on the strongest received reference signal. The effective data rate for each vehicle is derived from the calculated SINR, by using the 4-bit channel quality indicator (CQI) table in [40], which maps the reported CQI onto a particular modulation coding scheme (MCS) and spectral efficiency value.

For the purposes of this study, the SINR to CQI mapping has been performed using the spectral-efficiency based approach presented in [41]. To take account for the effects of frequent changes in beam-configurations as well as handover of users from one gNB to another, the overhead values reported in [42] were used when calculating the effective data rate of the individual vehicles. In particular, a beam-reselection overhead of 23 ms was applied for each configuration change, and a 43 ms cell handover overhead was applied for each instance of user handover between gNBs.

The performance of the CRAB algorithm is assessed against three state-of-the-art benchmarks, namely,

- A Conflict-Aware Weighted Bipartite Matching (CAWBM) algorithm - an iterative approach based on avoiding overlapping between beams leveraging conflict-aware weighted bipartite graph matching, presented in [24] and labeled CAWBM in the plots;
- DBSCAN - a clustering-based approach [43] aiming to serve as many users with the same beam by identifying largest vehicle clusters of fixed width;
- Upper Confidence Bound (UCB) algorithm - a simple reinforcement learning approach [44] which balances exploration-exploitation as it gathers more knowledge about the best beam configuration.

The first two benchmark solutions offer a limited flexibility when selecting beam configurations. Specifically, CAWBM requires a fixed number of beams, which we set to 4, chosen as the best performing option in line with the results in [24]. DBSCAN, on the other hand, requires a fixed beamwidth, which we set to 10° , selected as the option providing the best trade off between performance and complexity, based on the results in [24], [45]. On the positive side, DBSCAN is a distributed algorithm which can be executed dynamically to account for mobility in the network, while CAWBM is a static, centralized, algorithm which takes advantage of the comprehensive knowledge about the network to better match gNBs with zones. The third benchmark algorithm, UCB, does not impose any limits on flexibility and is a fully decentralized solution. However, since it is based on reinforcement learning principles, it requires a significant number of iterations to learn the beam configuration that provides the highest value of reward. The reward of each gNB is defined, similar to the CRAB optimization objective, in terms of the achievable rate, i.e., for gNB g the reward is defined as $\sum_{z \in \mathcal{Z}_n} \sum_{q \in \mathcal{Q}(z)} R_q(z, \mathbf{b}_g)$. For a detailed implementation of the algorithm, the reader is referred to [44, Chapter 2].

For CRAB, we employ the coverage criterion with a threshold of $c_{\text{thr}} = 0.15$. Further, all messages are initialized to random values uniformly drawn between 0 and 1, and we deem that convergence at a single node is reached when the difference between consecutive outgoing messages is lower than 10^{-5} times the minimum over all values carried by the most recent messages. A node will not send out a message that has not changed from the previous iteration. A node that has reached convergence will not send out any outgoing messages. Finally, a node, which has stopped receiving a message from a neighboring node, will use the latest received message from that node. The network converges once all the nodes reach convergence, and there are no more messages to be exchanged between nodes. Once an estimate of the marginal distribution for the gNB is obtained, the gNB randomly draws the beam configuration to apply every 100 ms. Again, each of the considered algorithms is executed periodically every second, and the total simulation duration is 10 s.

A. Connectivity analysis

Due to the influence of the interaction graph on the overall performance of the CRAB algorithm, we begin by assessing how the topology of such a graph is influenced by the connectivity criterion used to determine whether or not to draw an edge between two gNBs. We compare two of the criteria discussed in Sec. V-A, namely, the distance between gNBs (hereinafter referred to as distance-based criterion) and the fraction of the non-empty zones that are covered by *both* g and h (hereinafter referred to as coverage-based criterion).

It is important to point out that, regardless of the criterion employed, there is an inherent trade-off between the complexity of the CRAB algorithm (which in turn depends upon the connectivity of the interaction graph) and the performance of the resulting solution, i.e., the value of the objective (10). We can therefore compare the two connectivity criteria we

consider by characterizing the trade-offs that can be reached by adopting either of them.

The results are summarized in Fig. 6(left): each marker therein corresponds to a different value of a different criterion, with different criteria corresponding to different colors. The position of each marker along the x - and y -axes corresponds, respectively, to the degree of the resulting interaction graph and the total data rate it yields. Ideally, we would like solutions with a low complexity (i.e., markers closer to the left-hand side of the plot) and a high data rate (i.e., markers closer to the top of the plot). From Fig. 6(left) it is clear that red markers, referring to the results obtained under the coverage-based criterion, correspond to markedly better trade-offs than blue markers, referring to the result achieved under the distance-based criterion. Intuitively, this is equivalent to saying that the coverage-based criterion creates *better*, i.e., more meaningful edges in the interaction graph, therefore, it yields better performance with the same graph degree (or, equivalently, the same performance with a simpler graph), hence, less overhead and faster convergence of the CRAB algorithm.

In terms of convergence, as mentioned in the previous section, CRAB tends to reach the state well before the maximum number of iterations, which is set at 50. This holds true especially as the connectivity threshold is increased and the average node degree is reduced, as shown in Fig. 6(center), in blue. In particular, we can see that for threshold values around $c_{thr} = 0.15$, which offer the best trade-off between complexity and performance, the number of iterations required to converge is between 8 and 15. Recall that in those instances in which CRAB does not converge within a maximum number of iterations, we perform the pruning procedure. The number of prunings required to obtain a converging graph with respect to the threshold is plotted in red in Fig. 6(center), and we see that on average the number of edges that are pruned does not exceed 3. In contrast, UCB, the other iterative-based learning algorithm that we use as benchmark, requires significantly more iterations to converge, in the range of several thousands, as shown in Fig. 6(right) for each gNB. The number of iterations is capped at 10,000; therefore, for most gNBs UCB fails to converge, oscillating between several beam configurations that provide similar performance. Importantly, this behaviour further justifies our randomized approach to the beam management problem.

B. Benchmark evaluation

Next, we move to assessing the performance of our CRAB algorithm against the benchmark solutions.

A first aspect we are interested in is the absolute performance, i.e., which of the four solutions yields the highest effective data rate. As shown in Fig. 7(left), CRAB outperforms the benchmark solutions, by delivering in total 41% more data than DBSCAN, and 20% more than CAWBM. The performance improvement over UCB is small, as CRAB delivers only around 7% more data; however, as further discussed below, we recall that UCB exhibits a very long convergence time. It should also be noted that the overhead due




Fig. 6. Performance of the CRAB algorithm when using the distance-based criterion (blue) and the coverage-based criterion (orange): amount of data downloaded vs. average node degree (left); Average number of iterations as a function of the coverage-based connectivity threshold (center); Average number of iterations per gNB required for the convergence of the UCB algorithm (right).

to beam configuration changes and handovers was taken into account when processing the results shown in Fig. 7 and Fig. 8. Indeed, losses in the amount of effective data downloaded due to overhead are shown in orange in Fig. 7(left).

In order to gather more insight about the performance difference for individual vehicles, Fig. 7(center) highlights the fraction of vehicles covered by each solution. We can immediately notice that one of the reasons for DBSCAN's poor performance is the fact that it only manages to serve about 50% of the vehicles in the network, while both CRAB and CAWBM serve more than 80%. Focusing on the users that do get served, Fig. 7(right) shows the distribution of the data rate obtained by individual vehicles. We can see that the difference between CRAB and CAWBM is more clear for the vehicles in the top highest percentiles: under CRAB 20% of the top users can reach data rates over 500 Mb/s, while with CAWBM only the top 10% can reach such data rates. In other words, CRAB can serve *more* users than DBSCAN, and offer to them a *better rate* than CAWBM. Lastly, we note that in terms of these metrics, CRAB and UCB perform quite similarly. Indeed, we can consider UCB performance




Fig. 7. Total amount of data downloaded in blue, losses due to overhead in orange (left); percentage of vehicles covered (center); cumulative density function (CDF) of the data rate offered to the vehicles in Gb/s (right).




Fig. 8. CDF of number of vehicles covered by individual beams (left); CDF of number of vehicles served by (i.e., receiving data from) individual beams (center); CDF of the service duration per vehicle in seconds (right).

as an upper-bound for a fully decentralized solution, with an unlimited amount of time resources to learn from the environment and choose the optimal configuration in terms of the reward, i.e., effective data rate. However, as we saw in the previous section, UCB requires a significant number of iterations, and it is infeasible to implement it periodically as often as 1 s, as envisioned in this setup.

The high-level reason for such improved performance of CRAB over the two more practical solutions lies in the better coordination between beams and gNBs enabled by the message-passing approach. Fig. 8(left) and Fig. 8(center) highlight two of the main effects CRAB's superior beam configuration, namely, that CRAB is able to cover and serve more users with each beam than its counterparts, respectively, twice and four times more. Also, Fig. 8(right) shows how CRAB can serve each vehicle for a longer time, up to four times longer than CAWBM. In summary, CRAB is able to serve more users, give them a higher data rate, and cover them for a longer time, and at a feasible level of complexity.

Since CRAB is a randomized approach, the beam configurations tend to change more often, as expected. The CDF of the number of beam configuration changes per second per gNB is shown in Fig. 9(left). Note that most gNBs tend to change the beam configuration every second with DBSCAN and CAWBM, which coincides with the algorithm update period. With CRAB, on the other hand, 50% of the gNBs require less than 0.75 changes per second, even though a randomized approach is used. Only 40% of the gNBs in fact change the

beam configuration more often than once per second, however the benefits in terms of data rate outweigh the overhead incurred, as shown in Fig. 9(center), where we have plotted the amount of data served by the gNB versus the number of beam configuration changes per second, which shows clearly that gNBs with higher number of beam configuration changes also serve more data. Finally, as observed by looking at Fig. 9, while some individual gNBs may see a drop in capacity due to coordination resulting from CRAB, more than 80% of the nodes experience significant gains, often exceeding 100%. Specifically, when compared to DBSCAN over 30% of the gNBs experience improvement over 100%, and 55% of the nodes experience a gain of 50% or more. When compared to CAWBM, these numbers are expectedly lower, respectively, around 10% of the gNBs experience a data rate increase of 100% or more, while 30% of them experience a gain of 50% or more.

Fig. 10 summarizes how the performance obtained by different zones of the topology changes as a consequence of moving from DBSCAN to CRAB (left plot) and from CAWBM to CRAB (right plot). Compared to DBSCAN (left plot), CRAB can better serve the areas around the center of the city, where both the density of the gNB and vehicles is higher. This is due to DBSCAN's fully decentralized nature, whereby gNBs decide the configuration of the beams solely on their local knowledge about vehicle mobility. When the number of gNBs is relatively high, the configured beams at different gNBs may easily target overlapping areas, resulting




Fig. 9. Number of beam configuration changes per second (left); Amount of data downloaded using CRAB versus number of configuration changes per second, per gNB (center); Improvement of individual gNBs due to CRAB in terms of data rate compared to DBSCAN (blue circles) and CAWBM (red diamonds) (right).

in inefficient beam usage and higher interference. Looking at the right plot and the comparison with CAWBM, we can see that the improvement is more evenly distributed across the road topology, including the peripheral regions. This is due to CAWBM's centralized approach, which works best in dense areas, but may overestimate the interference between far-away nodes, thus providing a worse coverage in less dense areas. Overall, CRAB emerges as an effective compromise between a fully centralized approach like CAWBM, an approach like DBSCAN that exploits only local information and a learning approach which requires significant amount of time to converge, like UCB.

VII. CONCLUSIONS

We identified mmwave as a promising technology to enhance the capacity of vehicular networks. However, the performance of mmwave networks depends on the number, the alignment, and the width of beams between gNBs and vehicles, and these have to be carefully configured in order to maximize the data rate and avoid coverage overlapping among distinct gNBs.

To address this problem, we adopted a randomized approach and formulated an optimization problem, providing both a centralized and a decentralized version thereof. To efficiently

find a solution in a distributed manner, we devised an algorithmic framework leveraging belief propagation, called CRAB. The proposed framework (i) models the vehicular mmwave network as graph, capturing the interference that the possible beam configurations at the gNB may generate, (ii) adopts a message passing approach on such a graph, pruning it if necessary, and (iii) effectively finds a high-quality solution.

Our performance evaluation, based on real-world topology and realistic mobility traces, shows that CRAB significantly outperforms practical state-of-the-art alternatives, delivering in total from 20% up to 41% more data, while making, on average, a single gNB transfer 50% more data. Further, CRAB provides up to 30% better user coverage and an improved level of fairness in data rate performance across users.

REFERENCES

- [1] B. Bloessl, C. Sommer, F. Dressler, and D. Eckhoff, "The scrambler attack: A robust physical layer attack on location privacy in vehicular networks," in *IEEE ICNC*, 2015.
- [2] C. Sommer, S. Joerer, M. Segata, O. K. Tonguz, R. L. Cigno, and F. Dressler, "How shadowing hurts vehicular communications and how dynamic beaconing can help," *IEEE Transactions on Mobile Computing*, 2014.
- [3] I. Leontiadis, P. Costa, and C. Mascolo, "Extending access point connectivity through opportunistic routing in vehicular networks," in *IEEE INFOCOM*, 2010.
- [4] F. Malandrino, C. Casetti, C.-F. Chiasserini, and M. Fiore, "Optimal content downloading in vehicular networks," *IEEE Transactions on Mobile Computing*, 2012.
- [5] R. Yu, Y. Zhang, S. Gjessing, W. Xia, and K. Yang, "Toward cloud-based vehicular networks with efficient resource management," *IEEE Network*, 2013.
- [6] E. Uhlemann, "Autonomous vehicles are connecting," *IEEE Vehicular Technology Magazine*, 2015.
- [7] W. Roh, J.-Y. Seol, J. Park, B. Lee, J. Lee, Y. Kim, J. Cho, K. Cheun, and F. Aryanfar, "Millimeter-wave beamforming as an enabling technology for 5g cellular communications: Theoretical feasibility and prototype results," *IEEE Communications Magazine*, 2014.
- [8] R. Ketabi, M. Al-Qathrady, B. Alipour, and A. Helmy, "Vehicular traffic density forecasting through the eyes of traffic cameras: a spatio-temporal machine learning study," in *ACM DIVANet*, 2019.
- [9] A. Mukhopadhyay, R. R. Mazumdar, and F. Guillemin, "The power of randomized routing in heterogeneous loss systems," in *IEEE ITC*, 2015.
- [10] L. Zhang, Z. Li, and C. Wu, "Dynamic resource provisioning in cloud computing: A randomized auction approach," in *IEEE INFOCOM*, 2014.
- [11] Z. Zhou, Q. Wu, and X. Chen, "Online orchestration of cross-edge service function chaining for cost-efficient edge computing," *IEEE Journal on Selected Areas in Communications*, 2019.
- [12] M. Mézard and A. Montanari, *Information, Physics, and Computation*. Oxford University press, 2009.
- [13] L. Wei, Q. Li, and G. Wu, "Exhaustive, iterative and hybrid initial access techniques in mmwave communications," in *2017 IEEE Wireless Communications and Networking Conference (WCNC)*, 2017, pp. 1–6.
- [14] M. Giordani, M. Mezzavilla, and M. Zorzi, "Initial access in 5g mmwave cellular networks," *IEEE Communications Magazine*, vol. 54, no. 11, pp. 40–47, 2016.
- [15] M. Polese, F. Restuccia, and T. Melodia, "Deepbeam: Deep waveform learning for coordination-free beam management in mmwave networks," in *Proceedings of the Twenty-Second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, ser. MobiHoc '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 61–70. [Online]. Available: <https://doi.org/10.1145/3466772.3467035>
- [16] M. Zecchin, M. Boloursaz Mashhadi, M. Jankowski, D. Gund, M. Kountouris, and D. Gesbert, "A Novel Look at LIDAR-aided Data-driven mmWave Beam Selection," *arXiv preprint arXiv:2104.14579*, 2021.
- [17] M. Hussain and N. Michelusi, "Learning and Adaptation for Millimeter-Wave Beam Tracking and Training: a Dual Timescale Variational Framework," *arXiv preprint arXiv:2107.05466*, 2021.




Fig. 10. Improvement in absolute terms of amount of data downloaded at individual zones: compared to DBSCAN (left); compared to CAWBM (right).

- [18] N. Van Huynh, D. N. Nguyen, D. T. Hoang, and E. Dutkiewicz, "Optimal beam association for high mobility mmwave vehicular networks: Lightweight parallel reinforcement learning approach," *IEEE Transactions on Communications*, vol. 69, no. 9, pp. 5948–5961, 2021.
- [19] T. Ponnada, H. Al-Tous, and O. Tirkkonen, "Location-free beam prediction in mmwave systems," in *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*, 2021, pp. 1–6.
- [20] Z. Limani Fazliu, F. Malandrino, C. F. Chiasserini, and A. Nordin, "mmWave beam management in urban vehicular networks," *IEEE Systems Journal*, 2020.
- [21] X. Zhang, S. Pan, and Q. Miao, "Adaptive beamforming-based gigabit message dissemination for highway vanets," *IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS*, 2021.
- [22] Z. Sha, Z. Wang, S. Chen, and L. Hanzo, "Graph theory based beam scheduling for inter-cell interference avoidance in mmwave cellular networks," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 4, pp. 3929–3942, 2020.
- [23] H. Shokri-Ghadikolaei, L. Gkatzikis, and C. Fischione, "Beam-searching and transmission scheduling in millimeter wave communications," 2015.
- [24] Z. Limani Fazliu, F. Malandrino, C. F. Chiasserini, and A. Nordin, "Graph-based model for beam management in mmwave vehicular networks," *ACM Mobihoc (Workshops)*, 2020.
- [25] Y. Fan, Z. Zhang, and H. Li, "Message passing based distributed learning for joint resource allocation in millimeter wave heterogeneous networks," *IEEE Transactions on Wireless Communications*, vol. 18, no. 5, pp. 2872–2885, 2019.
- [26] N. J. Myers, J. Kaleva, A. Tolli, and R. W. J. Heath, "Message passing-based link configuration in short range millimeter wave systems," *IEEE Transactions on Communications*, vol. 68, pp. 3465 – 3479, 2020.
- [27] L. Codeca, R. Frank, S. Faye, and T. Engel, "Luxembourg sumo traffic (LuST) scenario: Traffic demand evaluation," *IEEE Intelligent Transportation Systems Magazine*, 2017.
- [28] M. Rapelli, C. Casetti, and G. Gagliardi, "Vehicular traffic simulation in the city of turin from raw data," *IEEE Transactions on Mobile Computing*, 2021.
- [29] M. A. Richards, *Fundamentals of Radar Signal Processing*. McGraw-Hill Professional, 2005.
- [30] T. S. Rappaport, G. R. MacCartney, M. K. Samimi, and S. Sun, "Wide-band millimeter-wave propagation measurements and channel models for future wireless communication system design," *IEEE Trans. on Comm.*, 2015.
- [31] M. R. Akdeniz, Y. Liu, M. K. Samimi, S. Sun, S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter wave channel modeling and cellular capacity evaluation," *IEEE JSAC*, 2014.
- [32] 3GPP, "5G; Study on channel model for frequencies from 0.5 to 100 GHz - Release 14," 3rd Generation Partnership Project (3GPP), Tech. Rep. 38.901, 2017.
- [33] —, "5G;NR; Base Station (BS) radio transmission and reception," 3rd Generation Partnership Project (3GPP), Tech. Rep. 38.104, 2019.
- [34] R. Flury and R. Wattenhofer, "Randomized 3d geographic routing," in *IEEE INFOCOM*, 2008.
- [35] S. Zhao, Z. Lu, and C. Wang, "How can randomized routing protocols hide flow information in wireless networks?" *IEEE Transactions on Wireless Communications*, 2020.
- [36] M. Rost and S. Schmid, "Virtual network embedding approximations: Leveraging randomized rounding," *IEEE/ACM Transactions on Networking*, 2019.
- [37] T. Shu, M. Krunz, and S. Liu, "Secure data collection in wireless sensor networks using randomized dispersive routes," *IEEE Transactions on Mobile Computing*, 2010.
- [38] A. T. Ihler, J. W. F. III, and A. S. Willsky, "Loopy belief propagation: Convergence and effects of message errors," *Journal of Machine Learning Research*, 2005.
- [39] R. E. Tarjan, *Data Structures and Network Algorithms*. CBMS-NSF Regional Conference Series in Applied Mathematics, 1983.
- [40] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA) – Physical layer procedures – Release 15," 3rd Generation Partnership Project (3GPP), Tech. Rep. 36.213, 2018.
- [41] M. Mezzavilla, M. Miozzo, M. Rossi, N. Baldo, and M. Zorzi, "A lightweight and accurate link abstraction model for the simulation of lte networks in ns-3," in *ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2012, pp. 55–60.
- [42] S. S. Kalamkar, F. Baccelli, F. M. Abinader, A. S. Marcano Fani, and L. G. Uzeda Garcia, "Beam management in 5g: A stochastic geometry analysis," *IEEE Transactions on Wireless Communications*, pp. 1–1, 2021.
- [43] M. Ester, H.-P. Kriegel, J. Sander, X. Xu *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise." in *KDD*, vol. 96, no. 34, 1996, pp. 226–231.
- [44] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. The MIT press, 2018.
- [45] Z. Limani Fazliu, F. Malandrino, and C. F. Chiasserini, "mmWave in vehicular networks: Leveraging traffic signals for beam design," in *IEEE WoWMoM Workshop on Communication, Computing, and Networking in Cyber Physical Systems (CCNCPS)*, 2019.

Zana Limani Fazliu (M'17) received the B.Sc. degree from the University of Prishtina, in 2008, the M.Sc. degree from Universitat Politècnica de Catalunya (Spain) / Politecnico di Torino (Italy) in 2010 and the PhD degree from Politecnico di Torino in 2017. Her research interests include wireless communications, mmWave communications, next-generation communication networks, game theory and machine learning applications. Currently she is an Assistant Professor at the University of Prishtina, Faculty of Electrical and Computer Engineering.

Francesco Malandrino (M'09, SM'19) earned his Ph.D. degree from Politecnico di Torino in 2012 and is now a researcher at the National Research Council of Italy (CNR-IEIIT). His research interests include the architecture and management of wireless, cellular, and vehicular networks.

Carla Fabiana Chiasserini (F'18) worked as a visiting researcher at UCSD, and as a Visiting Professor at Monash University in 2012 and 2016 and at TUB in 2021 and 2022. She is currently a Professor at Politecnico di Torino and EIC of Computer Communications.

Alessandro Nordio (M'99) is a Senior Researcher at the Institute of Electronics, Information Engineering and Telecommunications of the Italian National Research Council (CNR-IEIIT). In 2002 he received the Ph.D. in Telecommunications from EPFL, Lausanne, Switzerland. His research interests are in the field of signal processing, Smart Radio Environments, 6G and beyond communications, theory of random matrices, and crowdsourcing systems.