



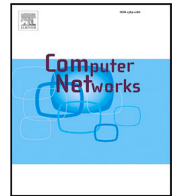
## Human-centric decision-making in cell-less 6G networks

Downloaded from: <https://research.chalmers.se>, 2025-09-25 23:29 UTC

Citation for the original published paper (version of record):

Chiaramello, E., Chiasserini, C., Malandrino, F. et al (2025). Human-centric decision-making in cell-less 6G networks. *Computer Networks*, 270. <http://dx.doi.org/10.1016/j.comnet.2025.111522>

N.B. When citing this work, cite the original published paper.



# Human-centric decision-making in cell-less 6G networks

Emma Chiaramello<sup>a,b</sup>, Carla Fabiana Chiasserini<sup>c,d,a,b</sup>, Francesco Malandrino<sup>a,b,\*</sup>,  
Alessandro Nordio<sup>a,b</sup>, Marta Parazzini<sup>a,b</sup>, Alvaro Valcarce<sup>e</sup>

<sup>a</sup> CNR-IEIIT, Torino, Italy

<sup>b</sup> CNIT, Italy

<sup>c</sup> Politecnico di Torino, Torino, Italy

<sup>d</sup> Chalmers University of Technology, Göteborg, Sweden

<sup>e</sup> Nokia Bell Labs, Paris, France

## ARTICLE INFO

### Keywords:

Cell-free networks

EMF

Human-centric communications

Optimization

## ABSTRACT

In next-generation networks, *cells* will be replaced by a collection of points-of-access (PoAs), with overlapping coverage areas and/or different technologies. Along with a promise for greater performance and flexibility, this creates further pressure on network management algorithms, which must make joint decisions on (i) PoA-to-user association and (ii) PoA management. We solve this challenging problem through an efficient and effective solution concept called Cluster-then-Match (CtM). While state-of-the-art approaches tend to focus on performance-related metrics, e.g., network throughput, CtM makes *human-centric* decisions, where pure network performance is balanced against energy consumption and electromagnetic field exposure. Importantly, such human-centric metrics concern all humans in the network area — including those who are not network users. Through our performance evaluation, which leverages detailed models for EMF exposure estimation and standard-specified signal propagation models, we show that CtM outperforms state-of-the-art network management schemes that solely focus on network performance, including those utilizing machine learning, reducing energy consumption by over 80% in indoor scenarios, and over 36% in outdoor ones.

## 1. Introduction

The recent trend in mobile communication networks is to make them evolve away from the cellular paradigm. Indeed, traditional cells – areas exclusively served by one base station, with overlap between cells being minimized – are becoming increasingly rare. After the introduction of multi-layer coverage via small cells [1] in 4G, and multi-RAT (radio access technology) networks in 5G [2], 6G will almost completely dispense with cells, and replace them with a collection of points-of-access (PoAs), thus becoming *cell less*. As exemplified in Fig. 1, PoAs may use different technologies (e.g., 5G-NR and Wi-Fi) and overlapping coverage areas, and each end user (including user-terminals like smartphones, but also non-human devices like robots) can be served by multiple PoAs.

Another relevant feature of 5G-and-beyond networks is the increasing attention to sustainability, understood in a wide sense as the impact of networking on the environment; such impact can be evaluated through aspects like power consumption [3] and electromagnetic field (EMF) exposure [4]. This has generated a trend toward *human-centric* networking, a new paradigm where pure network performance (e.g., throughput) is balanced against sustainability metrics. In many

cases, sustainability metrics become the objective of the optimization (e.g., minimize the energy consumption), while performance becomes a constraint (e.g., give to each user the minimum required throughput).

Compared to network performance metrics, human-centric ones tend to be more of a moving target: taking EMF exposure as an example, next-generation networks are moving towards higher frequencies, for which different metrics for EMF exposure assessment are being considered [5]. This means that our general goal is to endow human-centric network with a new capability – *a priori* sustainability management. This entails satisfying a set of constraints that are (i) different from the traditional ones, as they include human-centric metric as well as performance-related ones, and (ii) dynamic over time, as they account for factors like demand evolution and user mobility. Decision variables are similarly dynamic, as they reflect the fact that the network configuration can be changed over time to adapt to external conditions. This represents a marked departure from – and a significant improvement over – the state-of-the-art approach of ensuring that network configuration complies with a static, given set of key performance and value indicator (KPI, KVI) targets.

\* Correspondence to: corso Duca degli Abruzzi 24, 10129 Torino, Italy.

E-mail address: [francesco.malandrino@cnr.it](mailto:francesco.malandrino@cnr.it) (F. Malandrino).

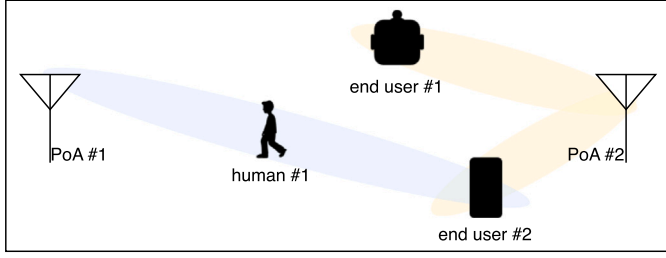


Fig. 1. A simple indoor cell-less network. Two points-of-access (PoAs) serve their end users through one or more beams; one of the users can be served by two PoAs. A human is irradiated by the beam from PoA #1 to User #2; humans will nonetheless incur EMF exposure whether or not they are network users.

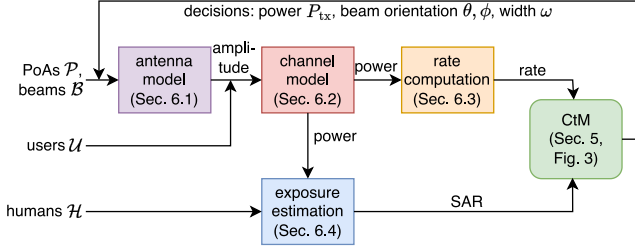


Fig. 2. The main components of our solution strategy and their mutual interactions..

As a result of the above, managing cell-less and human-centric networks will be substantially more complex than with current-generation ones. Reasons behind such complexity include (i) the number and diversity of the PoAs to manage; (ii) the fact that some technologies (especially higher-frequency ones) support beamforming [6], resulting in further decisions to make, and (iii) the fact that multiple PoAs may be available to serve a given user. Any decision made in cell-less and human-centric networks will impact, jointly and often in a counter-intuitive manner, both performance-related and human-centric metrics. The traditional (and popular, as detailed in Section 2) approach of ignoring human-centric aspects is of course not an option; also, straightforward optimization approaches are ruled out by the complexity of the problem to solve, as per Section 4.2.

Motivated by the need to combine both performance-related and human-centric metrics in the dynamic management of next-generation networks, we propose an efficient solution strategy called *Cluster-then-Match* (CtM), which can jointly decide on the association between users and PoAs, and the management of the PoAs themselves, including beam steering and power management aspects at the PoAs. As depicted in Fig. 2, CtM leverages the main features of cell-less networks – to wit, the availability, heterogeneity, and configurability of PoAs – to improve over the state-of-the-art in three main ways. More specifically:

1. CtM jointly manages all PoAs, hence, it is able to prevent conflicts between different decisions;
2. CtM considers all aspects of PoA management, from beam steering and beamwidth selection to power allocation;
3. CtM considers both performance-based and human-centric metrics, and the latter are evaluated for all humans in the network area, including those who are not network users.

To ensure both the decision quality and the fact that they can be made swiftly enough, CtM performs two main steps: users are first *clustered* into groups, and then groups are *matched* with the PoAs. This allows us to exploit the flexibility of the cell-free networking paradigm, where the PoA (and beam) serving each end user can be chosen dynamically and there is no *a priori* association between users and PoAs.

Our contributions and structure of the paper can thus be summarized as follows:

- We propose a concise and comprehensive model of cell-less networks, the decisions they require, and their effects on both performance-related and human-centric metrics, most notably, EMF exposure (Section 3);
- We formulate the resulting optimization problem, revealing how different decisions interact (Section 4), and analyzing its complexity, which rules out solving realistically-sized instances thereof (Section 4.2);
- We present our CtM algorithm and highlight how it is able to ensure both the quality of network management decisions and the speed with which they are made (Section 5);
- We describe how to evaluate CtM's performance under standardized channel models and scenarios introduced in [7], as well as detailed anatomic models for EMF exposure assessment (Section 6);
- We introduce the three benchmark we compare CtM against, including two based upon machine learning (ML), and present our performance evaluation results, showing that CtM outperforms the alternatives by over 80% (Section 7).

## 2. Related work

The *human-centric* expression is applied to 5 GB/6G networks in multiple contexts and with different meanings. Networks that are expected to closely interact with humans, e.g., in the tactile Internet, are called human-centric in [8], which investigates the additional capabilities and performance (most notably, latency) required by these applications. [9] takes a more technical view on the same issue, identifying novel 6G use cases (e.g., telepresence and augmented reality) and the associated requirements. At the same time, the societal impact of 6G networks is evaluated and assessed in [10], focusing on aspects like empowerment. Finally, since artificial intelligence (AI) is a key part of 6G, research efforts on human-centric networking are looking into making AI explainable [11] and to better control the data used for its training [12].

A second research area our work is related to is the so-called green networking, exploring the trade-offs between network performance and the energy consumption (hence, carbon footprint) they incur. Most approaches act on the network infrastructure and optimize its configuration, e.g., enlarging the coverage area of some base stations so that others can be switched off [13]. To achieve the former objective, novel transmission techniques are often employed, e.g., passive intelligent reflective surfaces [14]. Many recent approaches leverage AI to either predict the demand for coverage and capacity [15], evaluate the impact of switching decisions [16], or predict spectrum behavior [17]. At the same time, assessing and ensuring the sustainability of AI itself is a highly active research area [18], with approaches targeting both deep neural network (DNN) architectures [19] and distributed learning paradigms [20].

Finally, several works try to combine EMF exposure assessment and wireless (mostly, cellular) networking. Many early approaches [21] focus on network *planning*, working under the assumption that the infrastructure is given and impossible to control beyond setting power levels. Under such conditions, they check whether EMF limits are likely to be exceeded. More recent works look at network management, seeking to adapt the configuration of PoAs so as to account for both EMF limits and network performance; as an example, [22] focuses on power control in cell-free massive multiple-input-multiple-output systems. We seek to further advance with respect to these works by (i) fully modeling and exploiting the configurability of modern-day networks (e.g., beamforming and PoA selection), and (ii) estimating the levels of EMF exposure considering both the presence of users and non-users and their anatomical variability. Furthermore, the cell-less scenario targeted by CtM is different from the cell-free scenario targeted by works like [22]: while in cell-free scenarios multiple PoAs can be used to jointly serve the same user, CtM assigns each user to one PoA

only, and looks for the best configuration that allows each PoA to serve multiple users.

Last, a preliminary version of this work has been presented in [23]. With respect to our conference paper, this work compares CtM against several benchmarks, including ML-based solutions. Additionally, it evaluates CtM and its alternatives in a urban canyon network and propagation scenario, considering new phantoms for EMF exposure assessment.

### 3. System model

Our system model aims at representing the main elements of cell-less networks, their features, and the main decisions to make therein.

**Model elements and parameters.** Our system includes the PoAs, represented by elements  $p \in \mathcal{P}$ , and the end users (either devices such as IoTs, or user terminals), represented by elements  $d \in \mathcal{D}$ . Each PoA  $p \in \mathcal{P}$  is associated with one or more *beams*  $b \in \mathcal{B}$ . The number of beams, i.e., the size  $|\mathcal{B}|$  of set  $\mathcal{B}$ , depends upon the technology and features of the PoAs, and is an input to our problem. Further, we consider a set of *humans*  $h \in \mathcal{H}$ , representing people who might incur the EMF exposure resulting from the network. Humans in  $\mathcal{H}$  include both users of the network and simple passers-by, as depicted in Fig. 1.

End users, humans, and PoAs, are associated with a position and height, described through parameters:  $x_H(h)$ ,  $y_H(h)$ , and  $z_H(h)$  for humans  $h \in \mathcal{H}$ ;  $x_D(d)$ ,  $y_D(d)$ , and  $z_D(d)$  for end users  $d \in \mathcal{D}$ ;  $x_P(p)$ ,  $y_P(p)$ , and  $z_P(p)$  for PoAs  $p \in \mathcal{P}$ . People with a user terminal (e.g., smartphone, tablet, laptop, wearable) are represented by *both* an end user in  $\mathcal{D}$  and a human in  $\mathcal{H}$ , sharing the same location. IoT devices like robots, on the other hand, are only represented by an element in  $\mathcal{D}$ .

Concerning PoAs  $p \in \mathcal{P}$ , they operate at a frequency  $f(p)$  (which is known) and feature a set of beams  $\mathcal{B}(p) \subseteq \mathcal{B}$ , each of which must have a minimum width  $\omega^{\min}(p)$ . Furthermore, we indicate with  $\pi(b) \in \mathcal{P}$  the PoA originating beam  $b$ . If a single PoA can operate at different frequencies, then two distinct elements in  $\mathcal{P}$  are created to represent it, sharing the same location.

**Decisions and their effects.** Our decisions concern the joint aspects of (i) end user-to-PoA assignment, and (ii) PoA and beam management. The former is addressed through binary variables  $y(b, d) \in \{0, 1\}$ , expressing whether beam  $b \in \mathcal{B}$  (hence, PoA  $\pi(b) \in \mathcal{P}$ ) serves device  $d \in \mathcal{D}$ . Concerning the latter, we have to make four decisions, namely:

- the transmission power  $P_{tx}(p)$  of PoAs  $p \in \mathcal{P}$ ;
- the azimuth angle  $\phi(b)$  of beam  $b \in \mathcal{B}$ ;
- its elevation angle  $\theta(b)$ ;
- its width  $\omega(b)$ .

All the above decisions are then fed as input to the following three functions:

- $\text{Rate}(y, P_{tx}, \phi, \theta, \omega, d)$ , computing the data rate obtained by each end user  $d \in \mathcal{D}$ , accounting for interference (both intra- and inter-cluster), mobility, and noise;
- $\text{Energy}(y, P_{tx}, \phi, \theta, \omega)$ , computing the total energy consumption;
- $\text{Exposure}(y, P_{tx}, \phi, \theta, \omega, h)$ , estimating the EMF exposure levels incurred by each human  $h \in \mathcal{H}$ , quantified through any of the metrics discussed in Section 6.4.

It is important to highlight that the shape of the beam directly impacts the rate, energy, and exposure it results into. Specifically, a beam is a specific shape of an antenna radiation pattern characterized by a main lobe, i.e., a direction in space where most of the energy is radiated. The direction of the beam in spherical coordinates is characterized by a pair of angles,  $\theta(b)$ ,  $\phi(b)$ , and an angular width (called the beamwidth). These effects are captured by the functions Rate, Energy and Exposure defined above, which, therefore, depend on the orientation and width of all network beams. These functions are then combined to express the objective we seek to optimize and the system constraints. From the

viewpoint of our scheme, they are considered to be given and known. It is important however to remark how rate, energy, and exposure can be computed through different techniques – and, in the case of exposure, even be quantified through different metrics.

As an example, the received power can be determined through simple path-loss formulas or through much more complex channel models, accounting for clutter and mobility, such as the one we employ in Section 6. Similarly, EMF exposure can be quantified through metrics that focus on the electromagnetic field at a given location (e.g., electric field strength or power density [24]), or through more complex metrics like the specific absorption rate (SAR), which account for the interaction between electromagnetic fields and biological tissues [25]. Different scenarios and conditions call for different metrics, hence, there is no *right* way to compute the rate, energy, and exposure.

Accordingly, our system model and problem formulation – as well as the CtM solution strategy – can accommodate *any* technique to do so, from the simplest to the most realistic ones. This, coupled with the modular architecture of our implementation (see Section 7.2), greatly improves the flexibility of the proposed CtM solution, as well as its applicability to many real-world scenarios and conditions.

### 4. Optimization formulation

This section first presents the formulation of the problem we address, and then proves its NP-hardness.

#### 4.1. Objective and requirements

In general, we aim to optimize one of the metrics introduced above, i.e., one of the functions Rate, Energy or Exposure, while using the other two as constraints. The traditional approach would be to maximize the performance (quantified, typically, through the sum-rate [26]), subject to constraints about energy and EMF exposure.

We instead opt, in accordance with the human-centric paradigm, for a formulation where:

- minimizing the *global* energy consumption is the objective;
- constraints ensure that (i) each end user  $d \in \mathcal{D}$  obtains the required data rate, and (ii) exposure limits are honored for all humans in  $\mathcal{H}$ .

This formulation allows us to combine global and local quantities, as well as performance-related and human-centric metrics. Notice that our approach works unmodified for other options, e.g., maximizing the minimum rate.

More formally, we have:

$$\min_{y, P_{tx}, \phi, \theta, \omega} \text{Energy}(y, P_{tx}, \phi, \theta, \omega) \quad (1)$$

$$\text{Rate}(y, P_{tx}, \phi, \theta, \omega, d) \geq \text{Rate}^{\min}(d), \forall d \quad (2)$$

$$\text{Exposure}(y, P_{tx}, \phi, \theta, \omega, h) \leq \text{Exposure}^{\max}, \forall h. \quad (3)$$

The problem above reflects the complexity of our scenario, as discussed earlier: the objective (1) concerns a global variable, the rate requirements are to be guaranteed for each single end user, while meeting the guidelines about EMF exposure has to be ensured for all humans.

Also notice how minimum rates may be different for all end users, reflecting the fact that performance requirements vary wildly across applications and services. With regard to the EMF exposure, the international guidelines have been defined by international scientific bodies, e.g., the International Commission on Non-Ionizing Radiation Protection (ICNIRP), to protect people from all substantiated harmful effects of EMF exposure. The estimation of the levels of EMF exposure (i.e., the Exposure function in this study) is performed here in terms of the rate of energy absorbed by the biological tissues; thus, its values will be different for humans with different morphologies, consistently with the human-centric approach of our study. We remark that the levels of EMF exposure is not used as optimization objective, as the goal is not to obtain the lowest level possible, but to be compliant with the international exposure guidelines.

#### 4.2. Problem analysis

We can prove the following property.

**Property 1.** *The problem of minimizing (1) subject to (2) and (3) is NP-hard.*

**Proof.** We show how any instance of the generalized assignment problem (GAP), which is known to be NP-hard [27], can be transformed into an instance of our problem. In GAP, a set of  $I = \{i\}$  of items must be assigned to containers (bins)  $c \in C$ . Items have a value  $v_i$  and a weight  $w_i$ ; containers can support a maximum weight of  $W_c$ . Any GAP instance can be transformed in a heavily simplified instance of our problem, where:

- PoAs in  $\mathcal{P}$  correspond to items in  $I$ ;
- end users in  $\mathcal{D}$  correspond to containers in  $C$ ;
- $|\mathcal{B}|$  is such that all PoAs can serve all end users;
- the rate experienced by end user  $d$  when served by PoA  $p$  is equal to  $R(p, d)$ , regardless of all other decisions, and such a quantity is equal to the weight of the item corresponding to the PoA;
- the function Rate is such that end users served by multiple PoAs simply enjoy the sum of the individual rates  $R(p, d)$ , i.e., there is no scheduling and no interference;
- no EMF exposure is incurred (or, alternatively, there are no humans in  $\mathcal{H}$ ).

Notice that these are highly simplified assumptions, i.e., we are transforming a full instance of GAP into a very simple (hence, intuitively, relatively easy to solve) instance of ours.

We start from a condition where all PoAs serve all end users, and assign to each end user a minimum rate of  $\text{Rate}^{\min}(d) = \sum_{p \in \mathcal{P}} \text{Rate}(p, d) - W_c$ , where  $c$  is the container associated with end user  $d$ . Similarly to the rate, serving end user  $d$  with PoA  $p$  incurs an energy cost of  $E(c, d)$ , and the total energy (i.e., the value of objective (1)) is simply given by  $\sum_{c,d} E(c, d)$ . Also, initially, all PoAs serve all end users with one beam. Then, assigning item  $i$  to container  $c$  (in the GAP) means *stopping* serving the end user  $d$  associated through  $c$  with the PoA associated with  $i$ . Thus, every time we assign an item, we decrease the energy cost in our problem (which corresponds to increasing the value in the GAP). At the same time, the rate obtained by the end user decreases (i.e., the weight in the container increases) and gets closer to the minimum rate (i.e., the maximum container weight). Finally, the optimal solution of our problem is also the optimal solution for the GAP. Through the above procedure, we can reduce any instance of the GAP problem to an instance of ours. The reduction requires polynomial (linear) time, hence, the thesis is proved.  $\square$

There are two aspects of the proof of Property 1 that are worth highlighting. First, we convert an instance of the GAP problem to a *very simple* instance of our own; intuitively, this suggests that our problem might be even more complex than GAP. Indeed, if simple instances like the ones we create in the proof are NP-hard, that holds *a fortiori* for more complex, realistic instances.

Furthermore, the proof makes heavily simplified assumptions on the form of the functions Rate and Energy, while altogether ignoring Exposure. As exemplified by the channel model presented in Section 6, these functions can be very complex to compute, non-linear, and non-convex; indeed, they may even lack a closed-form expression. All such elements add further complexity to our problem; specifically, traditional solvers require a closed-form expression of the objective and constraints of the problem to solve, and many – critically, those that are guaranteed to converge to the optimum – only work for convex problems.

As a consequence of Property 1 and the remarks above, it is impractical – and, often, impossible – to solve our problem using conventional numerical solvers. This is a consequence of the size of the solution space to explore and the combinatorial nature of the solutions themselves, and motivates us to propose an efficient and effective heuristic strategy called cluster-then-match (CtM), as described next.

#### 5. The CtM strategy

As discussed in Section 4.2, it is prohibitively complex to solve our problem to optimality, mostly due to the size of its *solution space*. Accordingly, our CtM heuristic strategy is predicated on (i) sequentially considering the main elements of our system model (end users, PoAs, humans), and (ii) at each step, restricting our attention to the most promising possible decisions. In doing so, CtM can leverage the flexible relationship between users and PoAs, as well as the configurability of PoAs themselves, and explore a *subset* of the original solution space; although small, such subspace will contain most of (often, all) the highest-quality solutions. In so doing, CtM is able to balance the need for speed (because only a small number of solutions are explored) and the pursuit of effectiveness (since most of the good-quality solutions are among those explored).

Specifically, as depicted in Fig. 3, the CtM strategy consists of three main steps:

1. clustering the end users in  $\mathcal{D}$ , so as to make subsequent decisions on a per-cluster, rather than per-end user, basis;
2. assigning clusters to PoAs in  $\mathcal{P}$ , thus determining beam steering;
3. optimizing beam width and transmission power levels.

Steps 1 and 2 above directly map, respectively, to the twin objectives of CtM: exploring a small number of solutions, whilst ensuring their quality. In the following, we describe each of the steps separately.

**Step 1: Clustering.** A major reason for the complexity of the original problem formulated in Section 4 is the many-to-many relationship between end users in  $\mathcal{D}$  and PoAs in  $\mathcal{P}$ . On the one hand, we need to take this relationship into account to exploit the potential of cell-less networks; on the other hand, considering *all* possible end user-to-PoA associations is unnecessarily complex. Our intuition is to leverage the beams in  $\mathcal{B} = \bigcup_{p \in \mathcal{P}} \mathcal{B}(p)$ , and make the key observations that:

- since there will be more end users than beams, each beam will serve multiple end users, and
- to keep beams as narrow as possible, it is desirable that end users served by the same beam are close to each other. The latter is motivated by the fact that wider beams result in a higher likelihood that additional users – other than those in the cluster being served –, as well as more humans, to be affected in the beam itself. This, in turns, results in higher energy consumption and higher interference.

Following the observations above, in Step 1 of the CtM strategy we *cluster* the end users in  $\mathcal{D}$  into as many clusters as the number of possible beams in  $\mathcal{B}$ , exploiting their position information (i.e., the  $x_p$  and  $y_p$  parameters) to make clusters as small – in terms of area occupied by their end users – as possible. Although any clustering algorithm can be used in this step, for concreteness we adopt the  $k$ -means algorithm [28].  $k$ -means indeed yields very good results in scenarios like ours, also thanks to the fact that (unlike hierarchical clustering and later approaches such as DBSCAN [29]) it takes as an input the target number  $k$  of clusters, which corresponds to the number of beams in our case. End users in the same cluster can be served either via pairing in multi-user, multiple-input, multiple-output MIMO (MU-MIMO) systems [30], or via time-division multiplexing (TDM) if MU-MIMO is not used in the current scenario.

The number  $k$  of clusters to create must match the number  $|\mathcal{B}|$  of beams in order for the subsequent matching step to be ran efficiently, hence, that quantity does not need optimization. Alternative clustering algorithms could be explored in principle, so long as they accept the target number of clusters as a parameter; however, we have found that most recent clustering approach aim at discovering the optimal number of clusters, hence,  $k$ -means remains very well-suited to CtM in spite of its age.

**Step 2: Matching.** We have now as many clusters of end users as the number of beams in  $\mathcal{B}$ ; our next task is to *match* the beams and the clusters they serve. This is a *bipartite matching* problem; importantly, thanks to the clustering performed in Step 1, the matching to perform is one-to-one instead of many-to-many. To efficiently solve to optimality

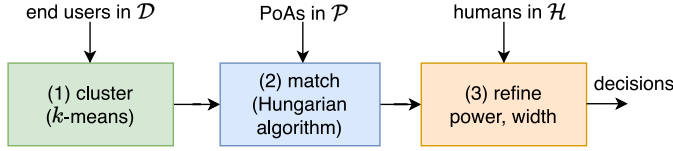


Fig. 3. The three main steps of the CtM strategy: *clustering* of the end users, *cluster-to-PoA assignment*, *optimizing* beam width and transmission power levels.

our one-to-one bipartite matching problem, we apply the Hungarian algorithm [31]. The input to the algorithm consists of a square matrix  $C = \{c_{kb}\}$ , where element  $c_{kb}$  is the *cost* of serving the  $k$ th cluster of devices through the  $b$ th beam; the output is the assignment minimizing the sum of the costs.

As our cost metric, we simply consider the *distance* from the PoA originating the beam to the centroid of the cluster of devices to serve. The main reason for choosing distance is that path loss depends upon distance, hence, serving clusters that are farther away would entail higher transmission powers and/or lower data rates. More specifically, making decisions based upon the distance has a twofold advantage:

- shorter distances between PoAs and devices make it possible to use lower power levels and narrower beams;
- using distances (as opposed to, e.g., the achievable rate) does not require any knowledge of the specific implementation of the Energy, Rate, or Exposure functions.

The latter item also implies that Step 2 of CtM is easy and quick to perform even in scenarios and cases where evaluating the rate, energy, or exposure is costly and/or time consuming (e.g., if those quantities are evaluated through simulations).

**Step 3: Optimizing beam width and power.** In this step, we set the width of each beam and the transmission power of each PoA, keeping into account both the objective (1) and the constraints (2) and (3). Specifically, we set the width of beams to the minimum value necessary to serve the end users assigned to them. Additionally, we reduce their power as much as possible to decrease the energy consumption and the EMF exposure, while ensuring the required performance.

Concerning beam widths, let  $D(b) \subseteq D$  be the set of end users belonging to the cluster assigned to beam  $b$ . Then the angle  $\alpha(b, d)$  from each end user  $d \in D(b)$  is given by:

$$\alpha(b, d) = \text{sgn}(\Delta_y) \arccos \frac{\Delta_x}{\sqrt{\Delta_x^2 + \Delta_y^2}},$$

where  $\Delta_y = y_D(d) - y_P(\pi(b))$ ,  $\Delta_x = x_D(d) - x_P(\pi(b))$  and the width of the beam is set to the angular width of the cluster, as observed from the PoA generating beam  $b$ , i.e.,

$$\omega(b) \leftarrow \max \left\{ \omega^{\min}(\pi(b)), \pi - \left| \alpha(b)^{\max} - \alpha(b)^{\min} - \pi \right| \right\},$$

where  $\alpha(b)^{\max} = \max_{d \in D(b)} \alpha(b, d)$  and  $\alpha(b)^{\min} = \min_{d \in D(b)} \alpha(b, d)$ . Setting the beam width to the angular width of the cluster means making the beam as narrow as possible (whilst serving all users of the cluster). Recall that narrow clusters mean a smaller likelihood of a beam affecting users not belonging to its cluster, hence, making clusters as narrow as possible minimizes inter-cluster interference. Also notice how the above expression also accounts for the minimum beam width  $\omega^{\min}$ , reflecting the fact that there are technological limits – specific to each PoA – preventing beams from being as narrow as one might desire.

The last part of Step 3 concerns power levels. Our key observations are as follows:

- reducing power might endanger feasibility (by resulting in a violation of constraint (2)), but will never affect adversely energy consumption (1) or EMF exposure (3);

- reducing the power of a PoA will not impact feasibility for end users served by another PoA.

Accordingly, we process one PoA at a time, as follows:

- set an amount of power  $\delta$ ;
- consider one of the PoAs  $p \in P$ ;
- reduce the power  $P_{tx}(p)$  by  $\delta$ ;
- repeat Step (c) so long as that results in a globally feasible solution;
- go back to Step (b) and consider a different PoA.

The above procedure can be repeated, in a manner similar to Newton's bisection algorithm [32], setting a smaller  $\delta$  and starting afresh from Step (a). By doing so, we can get even more fine-grained (hence, higher quality) decisions, at the cost of a longer running time. Concretely, we shall start from a large value of  $\delta$ , and then reduce it (e.g., by halving it) until the required precision has been reached. Importantly, the choice of  $\delta$  does not affect feasibility, since feasibility is checked globally and, no matter the value of  $\delta$ , step 3 never moves from a feasible to an infeasible solution. Also notice that, by considering PoAs in a sequential manner and changing one beam at a time, CtM reduces the impact of suboptimal decisions: indeed, if a suboptimal decision is made, it will not affect PoAs other than the current one.

### 5.1. Algorithm analysis

We now formally prove that CtM makes high-quality decisions *swiftly*, by having a low computational complexity regardless the size of the problem to solve.

**Property 2.** *The CtM strategy has polynomial worst-case computational complexity.*

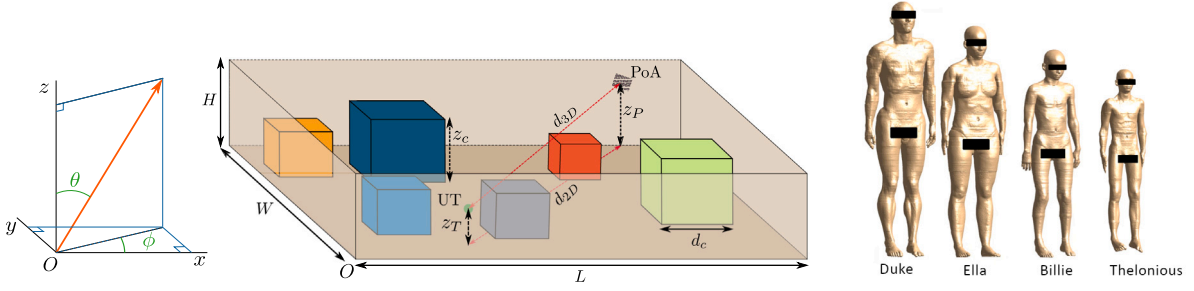
**Proof.** We prove the property by separately considering the three steps depicted in Fig. 3. Concerning Step 1, virtually all popular clustering techniques have polynomial complexity; taking  $k$ -means in particular, recent implementations can even achieve linear complexity [33]. The matching in Step 2 can be performed optimally through multiple approaches; considering the Hungarian algorithm, its complexity is cubic in the number of items to match (in our case, beams), as shown in [34]. In Step 3, width decisions are made only once per beam, hence, the added complexity is linear in the number of beams. Power levels are instead reduced at most  $|P| \left\lceil \frac{P_{tx}^{\max}}{\delta} \right\rceil$  times, hence, the added complexity is linear in the number of PoAs.

As each of the above steps has polynomial time complexity, and we perform them sequentially, it follows that the total complexity of the full CtM procedure is polynomial as well.  $\square$

The concrete meaning of Property 2 is that the computational complexity, hence, the time required by CtM to make its decisions, remains manageable as the size of the problem instance (e.g., the number of end users and PoAs) grows. This, in turn, implies that CtM decisions are made in time to be useful (e.g., to be acted upon) even in large and/or complex, real-world scenarios.

## 6. Reference scenario

For our reference scenario, we turn to ETSI technical report TR 138.901 [7]. The report collects the outcomes of study item RP-151606, covering different activities (e.g., spectrum allocation, scenarios of interest, measurements) related to frequencies between 6 and 100 GHz. The study item had identified different scenarios of interest and, for each scenario, a suitable channel model; the technical report TR 138.901 unifies those models into a single one, working for all scenarios and all bands between 6 and 100 GHz. From TR 138.901, we draw:



**Fig. 4.** Global coordinate system (GCS) (left), hall model for the InF-DH indoor scenario with colored objects representing clutters (center), and the detailed anatomical models used for electromagnetic field exposure assessment (right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

- two reference scenarios, to wit, an indoor factory (InF-DH) and an outdoor urban canyon (UMI-SC);
- the channel model that use throughout all our experiments.

Each scenario includes a set of PoAs,  $\mathcal{P}$ , of end users,  $\mathcal{D}$ , and of humans  $\mathcal{H}$ . Furthermore, indoor scenarios may include walls, ceiling, and clutter, i.e., metallic machinery or other irregular objects, as exemplified in Fig. 4(center). The signal propagation in these environments depends upon several factors, including the frequency used for the communication, the orientation of antennas and beams, and the position of PoAs and users. Indoor scenarios also account for the hall size and the clutter density.

The environment is associated with a spherical or Cartesian global coordinate system (GCS), as depicted in Fig. 4 (left) for the indoor scenario. In the GCS, the azimuth angle,  $\phi \in [-180^\circ, 180^\circ]$ , is measured counterclockwise from the  $x$  axis, whereas the zenith angle,  $\theta \in [0^\circ, 180^\circ]$ , is measured from the  $z$  axis. Thus, directions on the horizon have  $\theta = 90^\circ$ . To characterize the propagation channel, we first discuss the antenna model and then we detail the channel impulse response.

### 6.1. Antenna model and beamforming

The ETSI specifications [7] assume that PoAs and receivers are equipped with rectangular arrays (antenna panels) of  $M \times N$  elements organized in a regular grid. To each panel is associated a local coordinate system (LCS), with the  $x$  axis being along the panel broadside direction. The mechanical orientation of the panel thus completely defines the relation between its LCS and the GCS. To simplify the model description, we here assume that all panels have tilt angle equal to  $90^\circ$  and slant angle  $0^\circ$ , whereas the azimuth angle,  $\alpha$ , can be arbitrarily chosen. The radiation power pattern of a panel element can be described in the LCS by the function  $A_{dB}(\theta, \phi)$ ,<sup>1</sup> given by  $A_{dB}(\theta, \phi) = 0$  the ideal isotropic model, and by

$$A_{dB}(\theta, \phi) = 8 - \min\{A_{1,dB}(\theta) + A_{2,dB}(\phi), 30\}$$

with

$$A_{1,dB}(\theta) = \min\left\{12\left(\frac{\theta - 90^\circ}{65}\right)^2, 30\right\}$$

$$A_{2,dB}(\phi) = \min\left\{12\left(\frac{\phi}{65}\right)^2, 30\right\},$$

according to the 3GPP model with 8 dBi gain.

An antenna panel can generate a beam  $b$  toward a desired direction, characterized by the angles  $(\tilde{\theta}, \tilde{\phi})$ , measured w.r.t. the LCS. Assuming that all antenna elements are vertically polarized, the electromagnetic

field sent by the panel along the direction  $(\theta, \phi)$  (w.r.t. the LCS) is given by

$$F(\theta, \phi, \tilde{\theta}, \tilde{\phi}, M, N) = \tilde{F}(\theta, \phi) \frac{\text{sinc}(M g_1)}{\text{sinc}(g_1)} \frac{\text{sinc}(N g_2)}{\text{sinc}(g_2)} \sqrt{NM} e^{j\pi[(M-1)g_1 + (N-1)g_2]}$$

where  $\tilde{F}(\theta, \phi) = 10^{A_{dB}(\theta, \phi)/20}$  is the field generated by an array element, and  $g_1 = d_v(\cos \theta - \cos \tilde{\theta})$  and  $g_2 = d_h(\sin \phi \sin \theta - \sin \tilde{\phi} \sin \tilde{\theta})$ . Note that such an antenna model can describe both the field associated with antennas used for communication as well as a human body. In the latter case, a portion of a body with surface  $\lambda^2/4\pi \text{ m}^2$  can be associated with a single isotropic antenna element with gain given by  $A_{dB}(\theta, \phi) = 0$  and effective area  $\lambda^2/4\pi$ , where  $\lambda$  is the signal wavelength.

### 6.2. Channel impulse response

We now focus on a single wireless link connecting a generic PoA  $p \in \mathcal{P}$  to a generic target  $t \in \mathcal{T}$  where  $\mathcal{T} = \mathcal{D} \cup \mathcal{H}$  is the set comprising end users and humans. For simplicity, we omit the dependence on PoA  $p$  and target  $t$ . The distance between the PoA and the target is denoted by  $d_{3D}$ , and its projection on the ground is referred to as  $d_{2D}$  (see Fig. 4). The InF-DH and UMI-SC model provide expressions for the channel impulse response,  $\tilde{h}(\tau)$ , which accounts for both line of sight (LoS) and non-LoS (nLoS) propagation. The probability of the PoA and the target being in LoS depends on several factors, including the distance  $d_{2D}$ , on the PoA and target height and (for the InF-DH scenario) on the clutter density [7]. The channel impulse response has, in general, the following expression:  $\tilde{h}(\tau) = 10^{(\gamma_{dB} + P_{x,dB})/20 + \zeta} h(\tau)$  where  $P_{x,dB}$  is the power emitted by the PoA  $p$  (in dBm),  $\gamma_{dB}$  is the pathloss, and  $\zeta$  is a random variable modeling shadow fading. The function  $h(\tau)$  describes the wireless link as a set of rays grouped in a (scenario-specific) number  $N_c$  of clusters. Cluster  $j$  is described by the normalized power  $P_j$  and collects the contribution of  $N_r = 20$  rays, which are characterized by a common delay  $\tau_j$ . Ray  $\ell$  of cluster  $j$  has azimuth and zenith angles of departure (AoD)  $(\theta_{p,j,\ell}, \phi_{p,j,\ell})$ , measured w.r.t. to the GCS and, similarly, azimuth and zenith angles of arrival (AoA)  $(\theta_{t,j,\ell}, \phi_{t,j,\ell})$ , when observed from the target. If the PoA and the target are in LoS, an additional path, with LoS AoD  $(\theta_{p,0}, \phi_{p,0})$  and LoS AoA  $(\theta_{t,0}, \phi_{t,0})$  (in the GCS), has to be considered. Thus, in the delay domain, we have:

$$h(\tau) = \begin{cases} \sum_{j=1}^{N_c} \sqrt{\frac{P_j}{N_r}} \sum_{\ell=1}^{N_r} F_{T,j,\ell} F_{P,j,\ell} e^{j\varphi_{j,\ell}} \delta(\tau - \tau_j) & \text{if nLoS propagation apply} \\ \sqrt{\frac{1}{1+K}} h^{\text{nLoS}}(\tau) + \sqrt{\frac{K}{1+K}} h^{\text{LoS}} \delta(\tau - \tau_1) & \text{if LoS propagation apply} \end{cases} \quad (4)$$

where  $\delta(\tau)$  is the Dirac delta function,  $K$  is the Rician fading coefficient,  $\varphi_{j,\ell}$  is a random phase and

$$F_{T,j,\ell} \triangleq F(\theta_{T,j,\ell}, \phi_{T,j,\ell} - \alpha_T, \tilde{\theta}_T, \tilde{\phi}_T, M_T, N_T)$$

<sup>1</sup> A variable name with value expressed in dB shows the subscript “dB”.

$$F_{p,j,\ell} \triangleq F(\theta_{p,j,\ell}, \phi_{p,j,\ell} - \alpha_p, \tilde{\theta}_p, \tilde{\phi}_p, M_p, N_p)$$

are the PoA and target fields computed according to (4). Specifically, for computing  $F_{T,j,\ell}$ , we substituted in (4) the antenna parameters of the target (e.g.,  $M_T, N_T$ ), and beam direction  $\tilde{\phi}_T, \tilde{\theta}_T$  while, for computing  $F_{p,j,\ell}$ , we employed those of the PoA, namely,  $M_p, N_p, \tilde{\phi}_p$  and  $\tilde{\theta}_p$ . Moreover, in such expressions, the azimuth angles  $\alpha_p$  and  $\alpha_T$  represent the azimuth mechanical orientation of the PoA and target antennas, respectively. Finally,  $h^{\text{LoS}} = F_{T,0} F_{p,0} \exp(-j\frac{2\pi}{\lambda} d_{3D})$  with

$$\begin{aligned} F_{T,0} &\triangleq F(\theta_{T,0}, \phi_{T,0} - \alpha_T, \tilde{\phi}_T, \tilde{\theta}_T, M_T, N_T) \\ F_{p,0} &\triangleq F(\theta_{p,0}, \phi_{p,0} - \alpha_p, \tilde{\phi}_p, \tilde{\theta}_p, M_p, N_p). \end{aligned} \quad (5)$$

Note that, according to [7], the values of the Rician fading coefficient  $K$ , of the AOA and AOD of each ray, of the cluster powers  $P_j$  and of the delays  $\tau_j$  take random values and are obtained through a long and complex procedure whose description would require a large amount of space. The interested reader is referred to [7] for details. If in the scenario there are several PoAs and targets, the wireless links connecting any pair of them are to be considered mutually independent.

### 6.3. SINR, rate, and power density at the target

Given the set of simultaneously transmitting PoAs ( $\mathcal{P}$ ) and the set of targets ( $\mathcal{T}$ ), we recall that with each PoA  $p$  is associated a position  $[x_p(p), y_p(p), z_p(p)]$ , a carrier frequency  $f(p)$ , a panel with  $M_p(p) \times N_p(p)$  elements, and a mechanical orientation angle  $\alpha_p(p)$ , a transmit power  $P_{\text{tx},dB}(p)$ , and a beam  $b$ , whose direction, in the panel LCS, is specified by the angles  $\tilde{\phi}_p(p)$  and  $\tilde{\theta}_p(p)$ . Similarly, the target has position  $[x_T(t), y_T(t), z_T(t)]$ . For simplicity, we assume that the target has an antenna panel with  $M_T(t) = N_T(t) = 1$  element and its power pattern is isotropic. Thus, the target mechanical orientation angle becomes meaningless as well as the notion of beam, i.e., the variables  $\alpha_T(t)$ ,  $\tilde{\phi}_T(t)$ , and  $\tilde{\theta}_T(t)$  can be arbitrarily chosen since they have no effect. Then, for every  $p$  and  $t$ , we define the impulse response of the channel connecting them as  $\tilde{h}_{p,t}(\tau)$ , computed as specified in Section 6.2, with power  $\mathcal{E}_{p,t} = \int_{-\infty}^{+\infty} |\tilde{h}_{p,t}(\tau)|^2 d\tau$ . Let  $\mathcal{F}$  be the set of frequencies used by the PoAs and  $\mathcal{P}(f) \subseteq \mathcal{P}$  the subset of PoAs transmitting at frequency  $f \in \mathcal{F}$ , where  $\cup_{f \in \mathcal{F}} \mathcal{P}(f) = \mathcal{P}$  and  $\mathcal{P}(f) \cap \mathcal{P}(f') = \emptyset \forall f, f' \in \mathcal{F}$ .

Thus, when  $p$  transmits and the target  $t$  is an end user  $d$ , the instantaneous measured SINR is

$$\text{SINR}_{p,d} = \frac{\mathcal{E}_{p,t}}{N_0 W_p + \int_{-\infty}^{+\infty} \left| \sum_{p' \in \mathcal{P}(f(p)), p' \neq p} \tilde{h}_{p',d}(\tau) \right|^2 d\tau},$$

where  $W_p$  is the signal bandwidth used by  $p$  and  $N_0 = -174$  dBm/Hz is the thermal noise power spectral density. The corresponding achievable rate is then given by

$$R_{p,d} = W_p \log_2(1 + \text{SINR}_{p,d}).$$

Instead, when the target  $t$  is a human  $h$ , the power received from all transmitting PoAs at frequency  $f$  is:

$$P_{\text{rx},h,f} = \int_{-\infty}^{+\infty} \left| \sum_{p \in \mathcal{P}(f)} \tilde{h}_{p,h}(\tau) \right|^2 d\tau$$

and the corresponding power density over the human body at  $h$  is obtained by dividing the received power in a certain point by the effective area  $\lambda^2/(4\pi)$ , i.e.,

$$S_{h,f} = 4\pi f^2 / c^2 P_{\text{rx},h,f} \quad \text{W/m}^2.$$

### 6.4. Exposure model

Given the power density over the human body  $S_{h,f}$ , the EMF exposure is assessed by estimating the specific energy absorption rate (SAR) over the whole body,  $\text{SAR}_{\text{wb}}$ , which is the power of the EMF absorbed over the entire body mass. We first derive the incident electric field

$E_{\text{inc}} = \sqrt{S_{h,f} Z_0}$  V/m where  $Z_0 = 377 \Omega$  is the free-space impedance. Considering a human with body mass index  $\text{BMI}_h$  kg/m<sup>2</sup>, the  $\text{SAR}_{\text{wb}}$  can be estimated as

$$\text{SAR}_{\text{wb}} = \left( \frac{E_{\text{inc}}}{E_{\text{ref}}} \right)^2 \cdot \frac{\text{BMI}_h}{\text{BMI}_{\text{ref}}} \text{SAR}_{\text{ref}} \quad (6)$$

where  $\text{SAR}_{\text{ref}}$ , measured in W/kg is the whole-body SAR induced by a reference incident field  $E_{\text{ref}}$  in a reference human body of mass index  $\text{BMI}_{\text{ref}}$ . Eq. (6) assumes far-field conditions, which is a reasonable assumption for the exposure scenarios and the frequency considered in this study.

The whole-body specific absorption rate ( $\text{SAR}_{\text{ref}}$ ) of the reference human body is defined as [35]:

$$\text{SAR}_{\text{ref}} = \frac{P_{\text{wb,ref}}}{M_{\text{ref}}} = \frac{1}{M_{\text{ref}}} \int_{\text{wb}} \sigma(r) E_{\text{RMS}}^2(r) dV \quad (7)$$

where  $P_{\text{wb,ref}}$  (in W) and  $M_{\text{ref}}$  (in kg) are, respectively, the whole-body absorbed power and the mass of the reference human body. As for the reference human body: (i)  $\sigma$  (in S/m) is its electrical conductivity; (ii)  $E_{\text{RMS}}$  (in V/m) is the root mean square value of the induced electric field, and (iii)  $V$  is its volume. The values of  $\text{SAR}_{\text{ref}}$  considered in this study were calculated by [36] through electromagnetic computational techniques in anatomical human models and determined in far-field conditions, when the reference field  $E_{\text{ref}}$  was set equal to 2.45 V/m. To account for the anatomical variability of human bodies, we estimated the  $\text{SAR}_{\text{wb}}$  for four computational whole-body anatomical human models, two adults and two children: Ella (female, 26 years), Duke (male, 34 years), Thelonious (male, 6 years) and Billie (female, 11 years) (see Fig. 4(right)). These models belong to the ‘‘Virtual Population (ViP)’’, a collection of anatomical models obtained from high-resolution Magnetic Resonance Imaging (MRI) data widely used for electromagnetic field exposure assessment [37], representing the gold standard for *in silico* biophysical modeling applications. Each model resembled the anatomical characteristics of a real individual and included up to more than 80 different tissues across the whole body.

The  $\text{SAR}_{\text{ref}}$  values provided by [36] were calculated at a set of frequencies slightly different from those considered in this study, i.e., 3 and 5 GHz. Since the dielectric properties of human tissues vary with frequency (thereby affecting the absorption of the power by human tissues), we needed to identify which value of  $\text{SAR}_{\text{ref}}$  could be considered reliable and accurate enough to approximate those obtained precisely at 3 and 5 GHz. To do so, we calculated the ratio between the dielectric properties across all tissues, at 3 and 5 GHz and those at the nearest frequencies reported in [36], that is, 2.45, 3.5 and 5.2 GHz. All obtained values were well below 2, the threshold for which, variations in dielectric properties do not substantially influence the whole-body SAR [38]. We can therefore state that this approximation does not affect the reliability the estimation of the exposure. For 5 GHz, the dielectric properties of human tissues [39,40] were very similar to those reported at 5.2 GHz, resulting in average ratios for both conductivity and permittivity equal to about 1. As to the 3 GHz frequency, the dielectric properties were compared with those reported for both 2.45 and 3.5 GHz, to identify which approximation could be the best. Although for both frequencies the ratios obtained as average across all the tissues were well below 2, the best approximation was found to be the 2.45 GHz, with ratios equal to 0.8 and 1.2 for conductivity and permittivity, respectively. From the  $\text{SAR}_{\text{wb}}$  of each frequency, the whole exposure level was obtained by summing the single frequency contribution, as defined in ICNIRP guidelines for multiple frequency exposure [35]. Evaluating SAR values allows assessing the compliance of the dose of EMF absorbed with the basic restrictions limits recommended by the guidelines set by ICNIRP [35]. For the whole body SAR the basic restrictions for general public are equal to 0.08 W/kg for both the considered frequencies [35].

### 6.5. Reference topologies

We consider two scenarios, both drawn from ETSI technical report [7]:

- An indoor factory scenario (Indoor Factory with Dense clutter and High base station height in [7], or InF-DH), considering a  $20 \times 80 \text{ m}^2$  factory hall where  $\rho_c = 40\%$  of the floor space is occupied by clutter, whose typical height is 2 m;
- An urban canyon scenario (UMI-SC in [7]), describing a  $40 \times 800 \text{ m}^2$  road stretch.

Indoor factory and urban canyon represent two of the most relevant scenarios for beyond-5G networks, hence, studying how CtM fares in such scenarios is of great interest. Furthermore, they are very different from one another, with challenges coming from different sources – clutter in the factory, buildings in the canyon: this allows us to explore to which extent CtM can adapt to diverse environment conditions.

For EMF exposure estimation, we consider the four human anatomical models described above, consisting of two adults, one adolescent and one child. Using models of people of different age and body structure allows us to enhance the human-centric aspect of our work, by more accurately estimating the effect of networks over the entire population. Notice that we only consider models of adults (Ella and Duke) for the indoor factory scenarios, as the youngest individuals are unlikely to be found there; on the other hand, we use all models for the urban canyon scenario. Also notice how, as per [7], the urban canyon requires a minimum of 10 meters distance between users and the PoAs serving them; humans, on the other hand, can be placed anywhere in the topology.

### 7. Performance evaluation

In the following, we first introduce the benchmarks we consider to assess the performance of our CtM solution (Section 7.1), and the simulator we developed for performance evaluation (Section 7.2). Subsequently, we compare our CtM strategy against the benchmark called MaxRate, which uses simulated annealing, (Section 7.3), and then against two ML-based benchmarks (Section 7.4). The main parameters of our reference scenarios are summarized in Table 1.

Throughout our performance evaluation, we consider the following metrics:

- the *data rate* experienced by the users, as our performance related metric;
- the *EMF exposure* experienced by all humans and the total *energy consumption*, as human-centric metrics.

We will use energy as our objective, and data rate and exposure as constraints.

#### 7.1. MaxRate and ML-based benchmarks

We compare the performance of the proposed CtM solution against three different benchmarks.

The first we consider, called MaxRate, uses the simulated annealing general-purpose optimization algorithm [41], combined with our system model as presented in Section 3, to optimize the following objective:

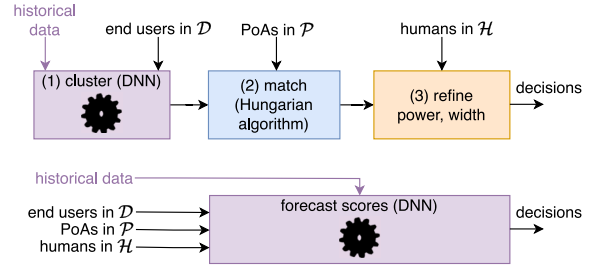
$$\max_{y, P_{\text{tx}}, \phi, \theta, \omega, d} \min_{\text{Rate}(y, P_{\text{tx}}, \phi, \theta, \omega, d),$$

i.e., to find a solution that maximizes the minimum among the data rates achieved by devices (Section 7.3).

By doing so, MaxRate can mimic the traditional approach of improving the overall network performance while guaranteeing fairness among devices. Examples of works following this approach in the field of next-generation network management abound, focusing on such

**Table 1**  
Simulation parameters.

Indoor factory scenario	
Hall size	$80 \times 20 \times 8 \text{ m}^3$
PoA height	7 m (3 GHz), 6 m (5 GHz)
User height	1.5 m
Clutter height	2 m
Clutter density	40%
No. of PoAs	8
Frequency	3 GHz (PoAs 1–2) 5 GHz (PoAs 3–8)
Bandwidth	20 MHz
No. of end users	100
No. of humans	200
Required rate	100 Mb/s
SAR <sub>wb</sub> limit	80 mW/kg [35]
Urban canyon scenario	
Scenario size	$800 \times 40 \text{ m}^2$
PoA height	10 m
User height	0.9 m or 1.5 m
No. of PoAs	8
Frequency	3.5 GHz (PoAs 1–2) 5.2 GHz (PoAs 3–8)
Bandwidth	20 MHz
No. of end users	100
No. of humans	200
Required rate	100 Mb/s
SAR <sub>wb</sub> limit	80 mW/kg [35]



**Fig. 5.** The ML-based benchmarks we compare against: DtM (top), replacing stage 1 of CtM with DNN-based clustering [45], and ML-only (bottom), using a DNN to make all decisions. Additions and changes with respect to CtM (Fig. 3) are marked in purple.

disparate aspects as cognitive networking [42], reflective intelligent surfaces (RISs) [43], and multi-cell coordination [44].

Furthermore, we compare the performance of CtM against two ML-based benchmarks, whose structure is summarized in Fig. 5. The first benchmark, depicted in Fig. 5(top), is called DNN-then-Match, or DtM for short. It essentially replaces the first stage of CtM, i.e., clustering, with a feed-forward DNN based upon [45]. The DNN is trained on historical data, i.e., past decisions and their outcome, and returns a user-to-cluster assignment that should result in the best performance.

The second benchmark, presented in Fig. 5(bottom), is called ML-only, and reflects the recent trends towards *model-free* approaches where decisions are made with little contribution from domain-specific knowledge or algorithms. ML-only reproduces the recent trend towards using ML techniques – most commonly, reinforcement learning – to solve network management problems, including beamforming [46–48], power control [46,47], and user assignment [46,48].

Importantly, both ML-based benchmarks require additional information, i.e., the historical data marked in purple in Fig. 5, and must be trained before they can be used. Depending upon the concrete case at hand, historical data may or may not be available, and devoting the necessary resources and time to training may or may not be practical.

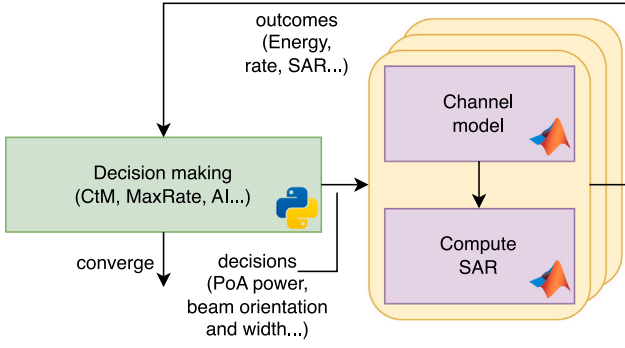


Fig. 6. Scheme of the simulator used for performance evaluation.

## 7.2. System simulator

For our performance evaluation we leverage an ad-hoc simulator, whose architecture is summarized in Fig. 6 and includes three main components:

- a decision-making module, coded in Python and implementing the CtM and MaxRate strategies;
- an implementation of the channel model described in Section 6, coded in Matlab;
- a module computing EMF exposure values as described in Section 6.4, coded in Matlab.

The architecture is highly decoupled: modules run as separate processes, can be – indeed, are – coded in different languages, and communicate by exploiting local or remote inter-process communication (IPC) mechanisms. Such decoupling has the distinct advantage of naturally exploiting the multi-core capabilities of modern hardware; indeed, since each module runs as a separate process, they can efficiently exploit different processors. Even more important, multiple instances of the channel model and SAR computation values can be ran in parallel, so as to evaluate multiple decisions at the same time. This is especially useful for techniques based upon simulated annealing or Markov decision processes, where *groups* of alternatives (“population”, “arms”...) must be evaluated concurrently.

Notably, our decoupled architecture makes it trivial to replace any module with a different implementation thereof; as an example, we might have: (a) a decision-making module based on Markov decision processes or reinforcement learning; (b) a different channel model; (c) a different way to quantify EMF exposure, using alternative metrics and/or models. On the negative side, the communication and coordination between different processes inevitably results in an overhead. However, compared to the running time of individual modules (especially the ones implementing channel models), such overhead is negligible. It can also be further reduced by using modern IPC techniques like message passing.

We take advantage of our decoupled architecture when implementing the ML-only benchmark, and replace the whole CtM block with a feed-forward DNN taking as an input the decisions (beam orientation, width, and power) and the resulting performance. For DtM and ML-only, we train the DNN over a dataset including 10,000 examples of randomly-generated decisions, along their outcome, i.e., the resulting rate, exposure, and energy consumption. The DNN is then used to obtain the set of decisions deemed to yield the best performance.

## 7.3. CtM performance against MaxRate

We now study the performance of the basic CtM algorithm described in Section 5 and compare it to the MaxRate benchmark in the two scenarios described in Section 6.5.

**Indoor factory scenario.** We begin by investigating the most basic aspect of CtM’s and MaxRate’s performance, i.e., the incurred energy consumption. Concerning this critical aspect, Fig. 7(left) shows that CtM outperforms MaxRate, yielding a power consumption that is almost an order of magnitude lower. Groups of bars in the plot correspond to different PoAs, and confirm that CtM can significantly reduce the transmission power of all PoAs compared to MaxRate. Also notice how, for both strategies, PoAs 1 and 2, operating at 3 GHz and typically serving farther-away end users, get assigned higher values of transmission power than the other PoAs that operate at 5 GHz.

Striking as Fig. 7(left) is, one might rightfully wonder whether such a reduced energy consumption comes at a cost in terms of data rate and service quality. The answer, as shown in Fig. 7(center), is both positive and negative. Indeed, under the CtM strategy (blue line in the plot) end users get lower data rates than under MaxRate (red line in the plot); however, such data rates are always (and sometimes significantly) above the required level (yellow line in the plot). This is consistent with the way data rates are included in our problem formulation in Section 4 and, more specifically, with constraint (2): so long as all end users are guaranteed  $\text{Rate}^{\min}(d)$ , there is no reason to further increase data rates. At a more general level, Fig. 7(left) and Fig. 7(center) suggest how adopting a human-centric approach, hence, using data rate performance as a constraint and not an objective, can bring substantial energy savings without jeopardizing service requirements.

In Fig. 7(right), we move to the other human-centric metric we consider, i.e., EMF exposure as quantified through  $\text{SAR}_{\text{wb}}$ . Once can immediately see that CtM yields a much lower  $\text{SAR}_{\text{wb}}$  than MaxRate; further,  $\text{SAR}_{\text{wb}}$  levels for either strategies are significantly below the limit values recommended in [35], to wit, 80 mW/kg. It is perhaps even more interesting to remark how there are *two* lines in the plot for each strategy, one dashed and one dotted, corresponding (resp.) to the Ella and Duke models introduced in Section 6.4. This further confirms that our approach can account for individual characteristics when assessing  $\text{SAR}_{\text{wb}}$ .

We now check whether there is any clear space pattern in the distribution of rate by plotting, in Fig. 8, the location of each end user and the rate they get under the MaxRate and CtM strategies. Consistently with Fig. 7(center), MaxRate ensures to virtually all end users a data rate that is much higher than required, hence, all markers in the left-hand side map are deep blue. In the right map, instead, we can see many lighter markers, though no red ones – highlighting the fact that CtM yields rate values that are closer to, but above, the required one. As one might expect, nodes with lower rate tend to be farther away from the PoA serving them, hence, experience higher attenuation and/or interference.

Similarly, Fig. 9 presents the  $\text{SAR}_{\text{wb}}$  levels experienced by humans in  $\mathcal{H}$ , under the MaxRate and CtM solutions. As we can expect from Fig. 7(right), CtM results in uniformly low  $\text{SAR}_{\text{wb}}$  values, hence, dark green markers. MaxRate, on the other hand, results in higher  $\text{SAR}_{\text{wb}}$  levels, hence, slightly lighter markers. As per Fig. 7(right), however, all  $\text{SAR}_{\text{wb}}$  values are significantly lower than the limit, thus, there are no purple markers on the plot. Also notice how the shape of the marker identifies the model (Ella or Duke) used to determine the  $\text{SAR}_{\text{wb}}$  of that specific human.

In Fig. 10, we look a little deeper into the relationship between the transmission power selected by PoAs and the  $\text{SAR}_{\text{wb}}$  level experienced by humans. Each marker in the plot corresponds to a solution (including infeasible ones) considered by any of the strategies. Its positions along the  $x$ - and  $y$ -axes correspond, respectively, to the total power level and 95th percentile of  $\text{SAR}_{\text{wb}}$  (computed over all humans). We

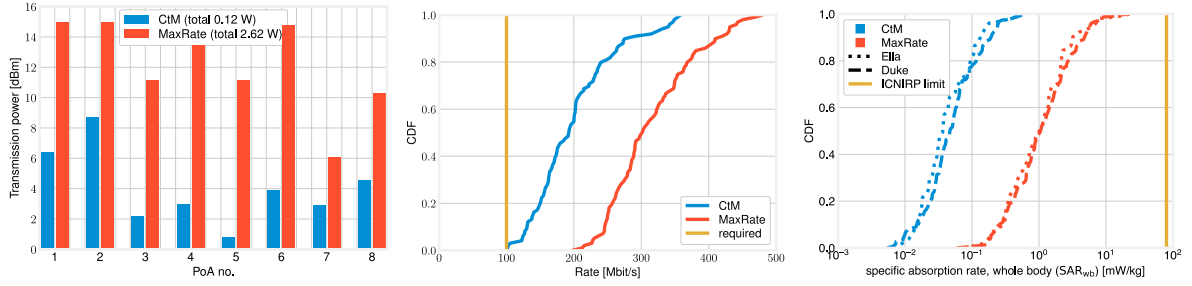


Fig. 7. Indoor factory scenario, CtM and the MaxRate benchmark: transmitted power (left), distribution of data rates (center), and  $SAR_{wb}$  values (right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

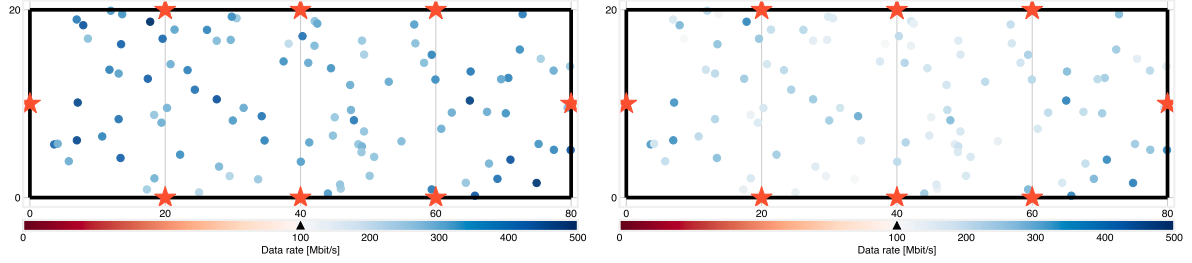


Fig. 8. Indoor factory scenario: data rate experienced by different end users under the MaxRate (left) and CtM (right) strategies. The black marker on the color bar corresponds to the required rate. Red stars represent PoAs. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

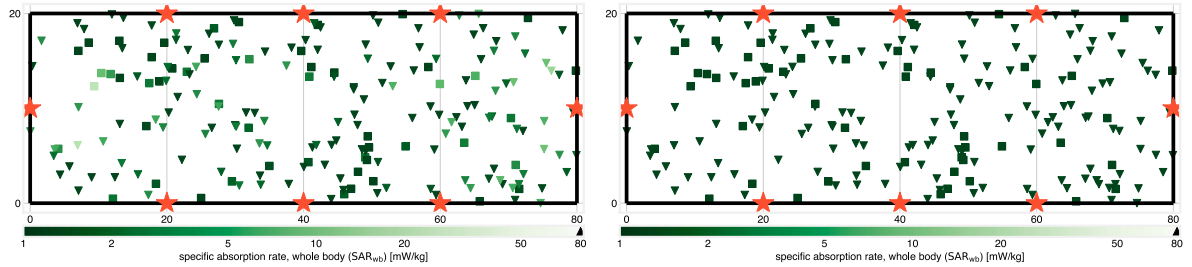


Fig. 9. Indoor factory scenario:  $SAR_{wb}$  experienced by different humans under the MaxRate (left) and CtM (right) strategies. The black marker on the color bar corresponds to the ICNIRP limit. Square and triangle markers correspond (resp.) to humans associated with the Duke and Ella models. Red stars represent PoAs. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

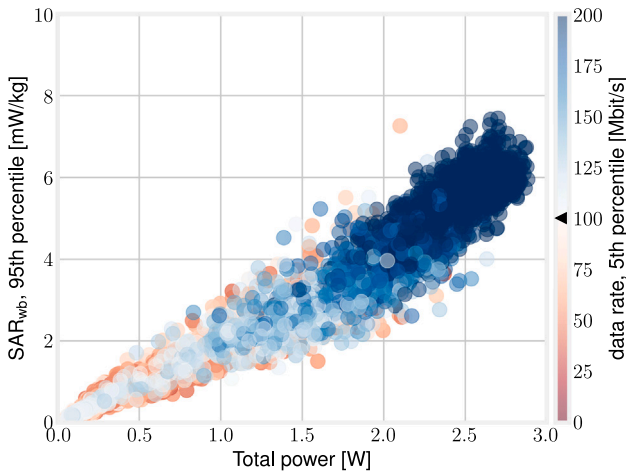


Fig. 10. Indoor factory scenario, all solutions considered by all strategies: relationship between total emitted power and 95th percentile of  $SAR_{wb}$ , with the color of each marker corresponding to the rate experienced by the 5th percentile of end users. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

can observe a very strong correlation between these two quantities. Even more interestingly, from the color of the markers, one can see that solutions with high power are almost invariably feasible. On the other hand, low power levels are associated with many infeasible solutions but also – critically – some feasible ones. Indeed, CtM can explore that part of the solution space, and find solutions that are consistent with the data requirements while incurring little energy consumption and EMF exposure.

The latter two quantities, although distinct, are linked: indeed, energy consumption is a *global* metric describing the PoAs and how much power they *emit*; SAR, on the other hand, depends upon the power level humans *receive*, and is a quintessentially *local* metric. In other words, there could be corner cases where a globally good solution, i.e., one with small energy consumption, has a locally unacceptable SAR for some specific human.

Finally, Fig. 11 sheds additional light on how PoAs operating at different frequencies are employed by the MaxRate and CtM strategies. Each marker in the plots corresponds to an end user; its position along the x- and y-axes corresponds, respectively, to the achieved data rate and the incurred interfering power; marker colors represent the frequency of the PoA serving the associated end user. We can observe that the higher frequency (5 GHz, purple markers) is associated with higher interfering power and lower data rates. Nonetheless, CtM (right

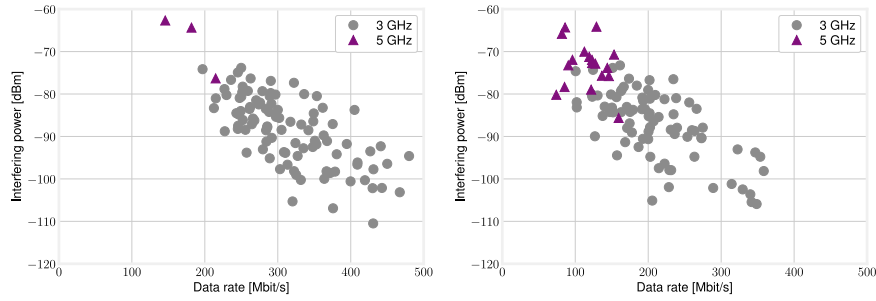


Fig. 11. Indoor factory scenario: data rate (x-axis) and interfering power (y-axis) experienced by all end users under the MaxRate (left) and CtM (right) strategies. The color of each marker denotes the frequency of the PoA serving the corresponding end user. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

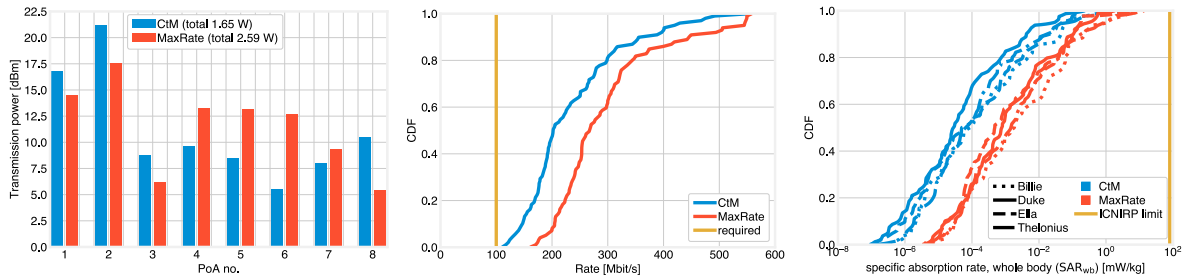


Fig. 12. Urban canyon scenario, CtM and the MaxRate benchmark: transmitted power (left), distribution of data rates (center), and SAR<sub>wb</sub> values (right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

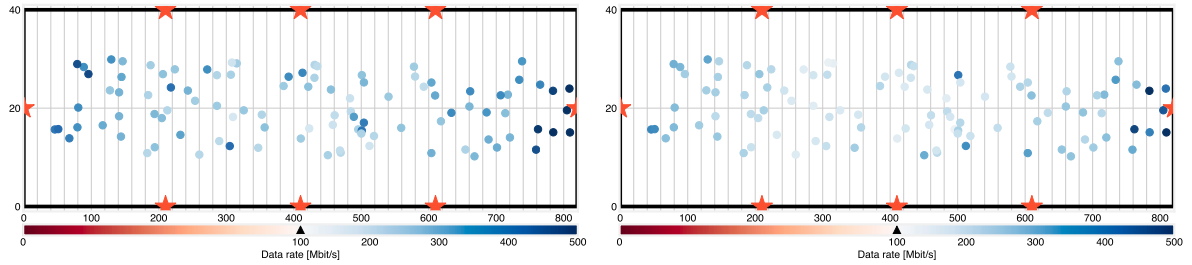


Fig. 13. Urban canyon scenario: data rate experienced by different end users under the MaxRate (left) and CtM (right) strategies. The black marker on the color bar corresponds to the required rate. Red stars represent PoAs. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

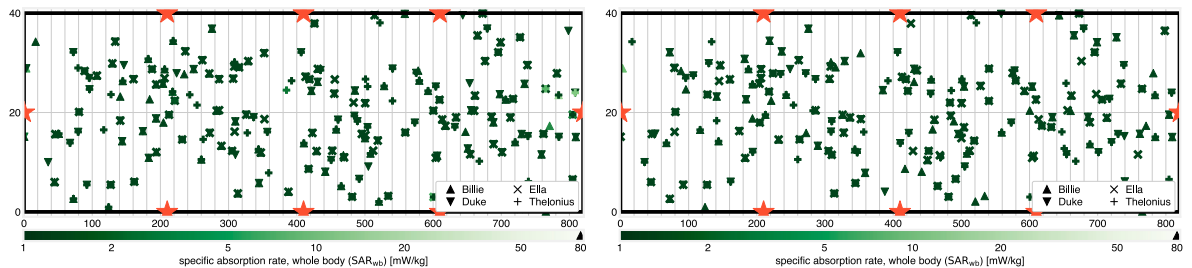


Fig. 14. Urban canyon scenario: SAR<sub>wb</sub> experienced by different humans under the MaxRate (left) and CtM (right) strategies. The black marker on the color bar corresponds to the ICNIRP limit. Square and triangle markers correspond (resp.) to humans associated with the Duke and Ella models. Red stars represent PoAs. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

plot) can leverage such a frequency to serve a much larger number of end users, while MaxRate (left plot) leverages almost exclusively the lower, 3 GHz frequency (gray markers in the plots). As also highlighted in Fig. 7(left), however, intensively using that frequency necessitates higher transmission power, which results in higher energy consumption and EMF exposure.

**Urban canyon scenario.** We now move to the urban canyon scenario, and assess whether the behavior of CtM and MaxRate follows the same behavior as in the indoor factory.

Fig. 12(left) confirms the behavior we observed in Fig. 7(left), with CtM resulting in a significantly lower power consumption than MaxRate. However, we can notice that the power used by CtM is higher than in the indoor factory scenario (Fig. 7(left)); the main reason for this difference is the larger size of the urban canyon scenario, which results in longer distances between PoAs and users, hence, higher attenuation. By looking at the columns reporting the power used by individual PoAs, we can also observe how CtM tends to use PoAs 1 and 2, placed on the two ends of the road stretch, to a larger extent than MaxRate. This is due to the elongated shape of the urban canyon

scenario: serving many users through PoAs 3–8 (placed along the long sides of the area under study) would require using wider beams, which would result in more energy (and EMF exposure); on the other hand, PoAs 1 and 2 (placed on the short sides of the area) can serve many users with narrower beams.

As one might expect, CtM's lower power and reliance on the lower-frequency PoAs result in a significantly lower rate than MaxRate, as shown in Fig. 12(center). However, all rate values are consistent with the 100 Mb/s requirement (denoted by a yellow line in the plot), hence, CtM's solutions are feasible. Importantly, this is essentially the same behavior as in Fig. 7(center); in both scenarios, CtM can provide users with the required data rate while minimizing power consumption, configuring the PoAs in the way that best suits the scenario at hand.

Consistently with Fig. 12(center), and similarly to Fig. 7 (right), Fig. 12 (right) shows that CtM yields a much lower exposure than MaxRate. This is true for all phantoms (notice that all four are used for the urban canyon scenario), even though individual phantoms do have different exposure distributions. Finally, similarly to Fig. 7(right), the EMF exposure is always significantly smaller than the ICNIRP reference level, marked with a yellow line in the plot.

Fig. 13, depicting the data rate experienced by each network user, further highlights the different way in which MaxRate (left map) and CtM (right map) use the PoAs. The difference is especially clear by considering the users around the middle of the topology: as CtM mostly uses PoAs 1–2, those users consistently have a lower rate. MaxRate is instead able to give some (not many) users therein higher rates by using the other PoAs, at the cost of a higher energy consumption.

Last, Fig. 14 confirms that, consistently with Fig. 12(right) and similarly to Fig. 9, SAR<sub>wb</sub> levels experienced by all humans in the topology are comfortably within the ICNIRP reference values. This is true for all phantoms (represented with different markers in the maps), and even for humans in the coverage areas of multiple PoAs, e.g., around the top-left part of Fig. 14(left).

**Runtime.** For our performance evaluation, we ran CtM on a server equipped with IBM POWER9 processors and a total of 1 TiB of RAM, running the Linux 4.18.0-305.el8.ppc64le kernel. With reference to Fig. 6, the decision-making module ran on Python 3.9, and the rate and exposure computation ones ran on GNU Octave 9.1 (this is because there is no MATLAB version for the Power ISA architecture used by POWER9 systems). Our reference implementation is extremely simple and, most significantly, lacks multi-processing capabilities. The running time are as follows:

- Octave (SAR computation): 4 s;
- Octave (rate computation): 39 s;
- Python (decision-making): 3 s;
- **total:** 47 s.

Notice that the individual components sum up to 46 seconds, but the total time includes additional contributions due to inter-process communication, context switching, and other types of overhead. Such a time is acceptable, since we are not reconfiguring the network in a real-time manner, but rather reacting to longer-scale mobility and demand patterns, e.g., due to traffic. Accordingly, we can expect to run CtM several times per day.

Furthermore, it is important to stress that there are vast margins to further optimize our running time, most importantly, by exploiting the multi-processing capabilities that are commonplace in modern hardware, as well as enabling GPU acceleration (which is supported in MATLAB but not in Octave).

#### 7.4. CtM performance against ML-based benchmarks

Back to the indoor factory scenario, we now compare the performance of CtM against the ML-based benchmarks presented in Section 7.1: DtM, which replaces the clustering stage in CtM with a

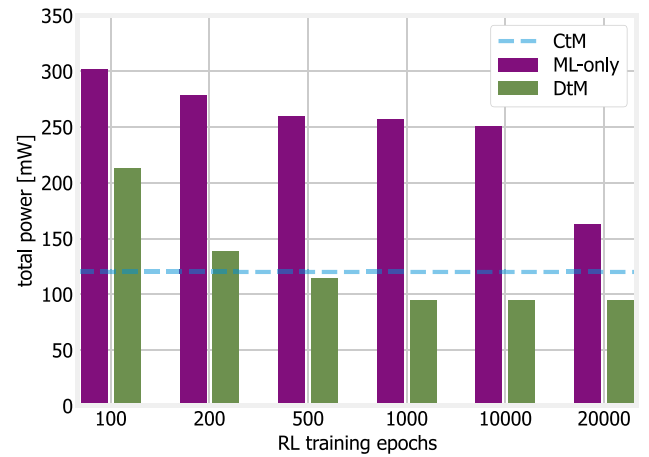


Fig. 15. Total power consumption under CtM, DtM and ML-only approaches as a function of the number of epochs allowed for DNN. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

DNN-based block, and ML-only, seeking to replace the whole CtM with a DNN.

Fig. 15 presents the total power consumption yielded by DtM (green bars) and ML-only (purple bars), with the power consumption of the original CtM reported as a blue, dashed line as a reference. We can immediately see that the green bars are consistently lower than the alternatives; indeed, DtM outperforms the original CtM provided that a sufficient number of epochs is allowed for training. As for ML-only, it results in a power consumption that is much higher than the alternatives. The situation improves as more training epochs are performed; however, even after tens of thousands of epochs, ML-only is unable to match CtM. The advantage of CtM and DtM lies in the fact that they exploit domain-specific knowledge and approaches; this allows them a much better efficiency compared to fully-automated, model-agnostic approaches.

The results in Fig. 15 confirm the intuition behind the design of DtM: ML is best used to *complement* domain-specific knowledge and approaches, as opposed to replacing them. One may conjecture that additional training (i.e., going further to the right in Fig. 15) might lead to ML-only eventually outperforming CtM; however, the resources and time needed for such training would render that approach infeasible or impractical in many real-world cases. The comparison between CtM and DtM also suggests that *k*-means is a reasonable choice as our clustering algorithms. Indeed, DtM's performance does exceed CtM's – which suggests better user-to-PoA association decisions –, but doing so requires a state-of-the-art DNN, thousands of training samples, and hundreds of training epochs: we can conjecture that a comparable complexity would be required from any alternative clustering algorithm to exceed *k*-means' performance.

Last, let us focus on the two ML-based benchmarks, and seek to understand the effect of combining ML with a domain-specific approach, as opposed to replacing the latter with the former. Specifically, Fig. 16(left) shows that DtM uses substantially lower power levels than ML-only, for all PoAs. This further highlights the effectiveness of retaining the blocks of CtM in charge of making power and width decisions, as per Fig. 5(top). For the same reasons, DtM yields a lower exposure than ML-only (Fig. 16(right)) while honoring the minimum-rate constraint (Fig. 16(center)).

**Training time.** The training time, using the same hardware we leveraged for the CtM tests, is about 24 s per epoch. It follows that it takes about three hours of training for DtM to match CtM's performance, and ML-only is unable to do so even after five days of training. Although the same remarks we made for CtM hold here, i.e., our implementation is

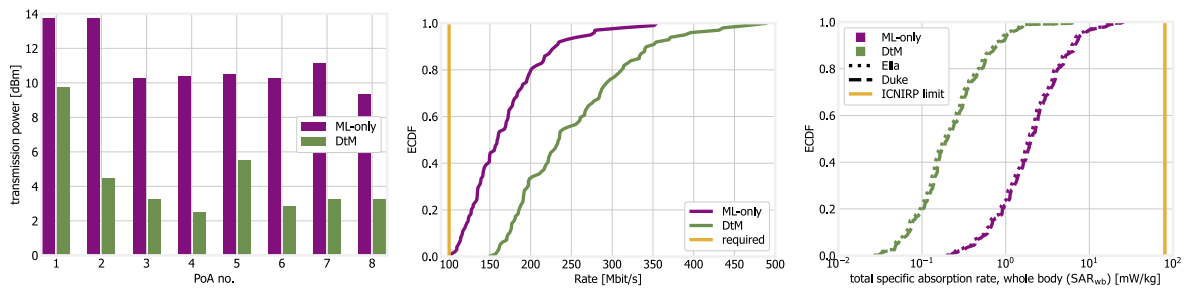


Fig. 16. DtM and ML-only for 1000 training epochs: transmitted power (left), distribution of data rates (center), and SAR<sub>wb</sub> values (right).

susceptible to significant improvements, this major difference suggests that the long training times required by ML-based approaches render them suitable for some, but not all, concrete scenarios.

## 8. Conclusions

We have addressed the scenario of cell-less networks, where there is no fixed, *a priori* association between users and PoAs, and management decisions must be made in a human-centric fashion, accounting for network performance as well as power consumption and EMF exposure. In order to cope with the problem complexity, we proposed a solution concept called CtM, for cluster-then-match. As its name suggests, CtM performs two main steps: first clustering end users, and then associating PoAs with clusters. CtM is able to minimize energy consumption; at the same time, CtM meets the rate requirements of all users and honors EMF exposure limits for all humans on the topology, including those who are not network users. Our performance evaluation, which leverages detailed models for EMF exposure estimation and standard-specified signal propagation models, compares CtM's performance against that of traditional and ML-based benchmarks, showing that CtM can reduce the power consumption by over 80%.

Future work will seek to improve CtM along two main directions: performance and user-to-PoA assignment. Concerning the former, we will optimize our implementation of CtM in order to further reduce the runtime and associated overhead; we may even consider departing from our current, modular architecture if the performance gain we obtain warrants it. Concerning the latter, we will explore approaches alternative to clustering, aiming to match the effectiveness of DtM's association decisions.

## CRediT authorship contribution statement

**Emma Chiamarello:** Writing – review & editing, Writing – original draft, Software, Methodology, Investigation. **Carla Fabiana Chiasserini:** Writing – review & editing, Supervision, Investigation, Conceptualization. **Francesco Malandrino:** Writing – review & editing, Writing – original draft, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Alessandro Nordio:** Writing – review & editing, Writing – original draft, Software, Data curation. **Marta Parazzini:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Investigation. **Alvaro Valcarce:** Writing – review & editing, Conceptualization.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgment

This work was supported by the SNS-JU-2022 project CENTRIC under the EU's Horizon Europe research and innovation programme, Grant Agreement No. 101096379.

## Data availability

Scenarios will be available upon request. Phantoms are commercial and cannot be shared (but are available from their seller.)

## References

- [1] T. Nakamura, S. Nagata, A. Benjebbour, Y. Kishiyama, T. Hai, S. Xiaodong, Y. Ning, L. Nan, Trends in small cell enhancements in LTE advanced, *IEEE Commun. Mag.* (2013).
- [2] B. Bangerter, S. Talwar, R. Arefi, K. Stewart, Networks and devices for the 5G era, *IEEE Commun. Mag.* (2014).
- [3] M. Ismail, W. Zhuang, E. Serpedin, K. Qaraqe, A survey on green mobile networking: From the perspectives of network operators and mobile users, *IEEE Comm. S&T* (2014).
- [4] E. Chiamarello, S. Flocchi, M. Parazzini, P. Ravazzani, J. Wiart, Stochastic dosimetry for radio-frequency exposure assessment in realistic scenarios, *Uncertain. Model. Eng. App.* (2019).
- [5] S.S. Zhekov, M. Yao, C. Di Paola, B. Xu, S. Zhang, A study on EMF exposure assessments with different metrics for user equipment antennas at 6 and 10 GHz, *IEEE Trans. Electromagn. Compat.* (2023).
- [6] I. Ahmed, et al., A survey on hybrid beamforming techniques in 5G: Architecture and system model perspectives, *IEEE Commun. Surv. Tutor.* (2018).
- [7] ETSI, 5G; Study on Channel Model for Frequencies from 0.5 to 100 GHz (3GPP TR 38.901 Release 16), Tech. Rep., 2020.
- [8] M. Maier, 6G as if people mattered: From industry 4.0 toward society 5.0, in: *IEEE ICCCN*, 2021.
- [9] M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, M. Zorzi, Toward 6G networks: Use cases and technologies, *IEEE Commun. Mag.* (2020).
- [10] K.F. Ystgaard, K. De Moor, Bring the human to the network: 5G and beyond, in: *IEEE ICT*, 2021.
- [11] S. Wang, M.A. Qureshi, L. Miralles-Pechuán, T. Huynh-The, T.R. Gadekallu, M. Liyanage, Applications of explainable AI for 6G: Technical aspects, use cases, and research challenges, 2021, arXiv preprint arXiv:2112.04698.
- [12] W. Saad, M. Bennis, M. Chen, A vision of 6G wireless systems: Applications, trends, technologies, and open research problems, *IEEE Netw.* (2019).
- [13] S. Verma, S. Kaur, M.A. Khan, P.S. Sehdev, Toward green communication in 6G-enabled massive internet of things, *IEEE Internet Things J.* (2020).
- [14] X. Yu, D. Xu, D.W.K. Ng, R. Schober, IRS-assisted green communication systems: Provable convergence and robust optimization, *IEEE Trans. Commun.* (2021).
- [15] G. Vallero, D. Renga, M. Meo, M. Ajmone Marsan, Processing ANN traffic predictions for RAN energy efficiency, in: *ACM MSWiM*, 2020.
- [16] B. Mao, F. Tang, Y. Kawamoto, N. Kato, AI models for green communications towards 6G, *IEEE Commun. Surv. Tutor.* (2021).
- [17] B. Yang, X. Cao, C. Huang, Y.L. Guan, C. Yuen, M. Di Renzo, D. Niyato, M. Debbah, L. Hanzo, Spectrum-learning-aided reconfigurable intelligent surfaces for “green” 6G networks, *IEEE Netw.*
- [18] R. Schwartz, J. Dodge, N.A. Smith, O. Etzioni, Green AI, *Commun. ACM* (2020).
- [19] Z. Du, Y. Deng, W. Guo, A. Nallanathan, Q. Wu, Green deep reinforcement learning for radio resource management: Architecture, algorithm compression, and challenges, *IEEE Veh. Tech. Mag.* (2020).
- [20] M. Xue, et al., EosDNN: An efficient offloading scheme for DNN inference acceleration in local-edge-cloud collaborative environments, *IEEE Trans. Green Commun. Netw.* (2021).
- [21] L. Chiaraviglio, S. Turco, G. Bianchi, N. Blefari-Melazzi, “Cellular network densification increases radio-frequency pollution”: True or false? *IEEE Trans. Wirel. Commun.* (2021).
- [22] S. Liesegang, S. Buzzi, EMF-aware power control for massive MIMO: Cell-free versus cellular networks, 2023, arXiv:2311.08989.
- [23] E. Chiamarello, C.F. Chiasserini, F. Malandrino, A. Nordio, M. Parazzini, A. Valcarce, Cluster-then-Match: efficient management of human-centric, cell-less 6G networks, in: *IEEE WoWMoM*, 2024.

- [24] D. Colombi, B. Thors, C. Törnevik, Implications of EMF exposure limits on output power levels for 5G devices above 6 GHz, *IEEE Ant. Wir. Prop. Lett.* (2015).
- [25] B.M. Hochwald, et al., Incorporating specific absorption rate constraints into wireless signal design, *IEEE Commun. Mag.* (2014).
- [26] A.B. Sediq, R.H. Gohary, R. Schoenen, H. Yanikomeroglu, Optimal tradeoff between sum-rate efficiency and Jain's fairness index in resource allocation, *IEEE Trans. Wir. Comm.* (2013).
- [27] D.G. Cattrysse, L.N. Van Wassenhove, A survey of algorithms for the generalized assignment problem, *Eur. J. Oper. Res.* (1992).
- [28] J.A. Hartigan, M.A. Wong, Algorithm AS 136: A k-means clustering algorithm, *J. R. Stat. Soc. Ser. C (Appl. Statist.)* (1979).
- [29] E. Schubert, J. Sander, M. Ester, H.P. Kriegel, X. Xu, DBSCAN revisited, revisited: why and how you should (still) use DBSCAN, *ACM Trans. Database Syst.* (2017).
- [30] S. Sur, I. Pefkianakis, X. Zhang, K.-H. Kim, Practical MU-MIMO user selection on 802.11 ac commodity networks, in: *ACM MobiCom*, 2016.
- [31] R. Jonker, T. Volgenant, Improving the hungarian assignment algorithm, *Oper. Res. Lett.* (1986).
- [32] R.B. Kearfott, Some tests of generalized bisection, *ACM Trans. Math. Software* (1987).
- [33] M.K. Pakhira, A linear time-complexity k-means algorithm using cluster shifting, in: *IEEE CICC*, 2014.
- [34] J. Edmonds, R.M. Karp, Theoretical improvements in algorithmic efficiency for network flow problems, *J. ACM* (1972).
- [35] I.C. on Non-Ionizing Radiation Protection, Guidelines for limiting exposure to electromagnetic fields (100 kHz to 300 GHz), *Health Phys.* (2020).
- [36] I. Liorni, M. Capstick, L. van Wel, J. Wiart, W. Joseph, E. Cardis, M. Guxens, R. Vermeulen, A. Thielens, Evaluation of specific absorption rate in the far-field, near-to-far field and near-field regions for integrative radiofrequency exposure assessment, *Radiat. Prot. Dosim.* (2020).
- [37] A. Christ, et al., The virtual family—development of surface-based anatomical models of two adults and two children for dosimetric simulations, *Phys. Med. Biol.* (2009).
- [38] P. Gajsek, W. Hurt, J. Ziriak, P. Mason, Parametric dependence of SAR on permittivity values in a man model, *IEEE Trans. Biomed. Eng.* (2001).
- [39] C. Gabriel, S. Gabriel, E. Corthout, The dielectric properties of biological tissues: I. Literature survey, *Phys. Med. Biol.* (1996).
- [40] S. Gabriel, R.W. Lau, C. Gabriel, The dielectric properties of biological tissues: II. Measurements in the frequency range 10 Hz to 20 GHz, *Phys. Med. Biol.* (1996).
- [41] J. Schneider, S. Kirkpatrick, *Stochastic optimization*, 2007.
- [42] L. Zheng, C.W. Tan, Maximizing sum rates in cognitive radio networks: Convex relaxation and global optimization algorithms, *IEEE J. Sel. Areas Commun.* (2013).
- [43] Y. Ju, S. Gong, H. Liu, C. Xing, J. An, Y. Li, Beamforming optimization for hybrid active-passive RIS assisted wireless communications: A rate-maximization perspective, *IEEE Trans. Commun.* (2024).
- [44] J. Qiu, J. Yu, A. Dong, K. Yu, Joint beamforming for IRS-aided multi-cell MISO system: Sum rate maximization and SINR balancing, *IEEE Trans. Wirel. Commun.* (2022).
- [45] A.Q. Ohi, M.F. Mridha, F.B. Safir, M.A. Hamid, M.M. Monowar, Autoembedder: A semi-supervised DNN embedding system for clustering, *Knowl.-Based Syst.* (2020).
- [46] F.B. Mismar, B.L. Evans, A. Alkhateeb, Deep reinforcement learning for 5G networks: Joint beamforming, power control, and interference coordination, *IEEE Trans. Commun.* (2019).
- [47] F. Fredj, Y. Al-Eryani, S. Maghsudi, M. Akrou, E. Hossain, Distributed beamforming techniques for cell-free wireless networks using deep reinforcement learning, *IEEE Trans. Cogn. Commun. Netw.* (2022).
- [48] M. Fozi, A.R. Sharafat, M. Bennis, Fast MIMO beamforming via deep reinforcement learning for high mobility mmwave connectivity, *IEEE J. Sel. Areas Commun.* (2021).