

On Cost-effective, Reliable Coverage for LoS Communications in Urban Areas

Gabriele Gemmi, *Student Member, IEEE*, Renato Lo Cigno, *Senior Member, IEEE*,
Leonardo Maccari, *Member, IEEE*

Abstract—The use of ultra high frequencies in 5G and future networks to improve transmission speeds and capacity requires that users’ equipment remain in Line of Sight with the access antennas most of the service time. This requirement implies a change in perspective to plan the coverage: Antennas cannot be placed on roofs or remote antenna sites, and a robust coverage is based on multi-antenna visibility from any point. This paper tackles the problem of public street coverage in urban areas with a data-driven methodology. Starting from 3D digital maps, we formalize the problem of antenna placement as a set coverage problem and leverage powerful heuristics to implement a general algorithm that allows the exploration of different policies, returning the detailed coverage, the antenna placement, and the cost of the coverage. Results on 15 areas in 3 Italian cities show the properties of different policies and confirm for the first time on large scale real data the feasibility of Line of Sight communications with a sustainable number of antennas per km².

Index Terms—Wireless networks and cellular networks, Coverage planning, Economic aspects, Internet connectivity and Internet access services, Access networks, Mathematical optimization

I. INTRODUCTION

The performance boost of next generation access networks, 5G and 6G included, is based mostly on the densification of base stations and the use of ultra high frequencies (mmWave, TeraHertz, up to optical frequencies) to provide increased capacity to users’ equipment. The consequence is that communications require Line-of-Sight (LoS) (or near LoS) links, since any obstacle (wall, vehicles, but also humans) at these frequencies is an almost impassable barrier. In these conditions, network reliability requires that each point in a public space is covered by multiple antennas to compensate dynamically for the shadowing of LoS links, which exacerbates the well-known requirement of a higher density of base stations compared to 4G. Manufacturers estimate that 5G will require more than 100 base stations per km², up from the 10 per km² in LTE [1], and even more in 6G. In the literature there are works suggesting densities with more than a hundred next Generation NodeB (gNB)/km² [2].

The placement of antennas (gNB from now on) for LoS communications in urban areas is a very complex task given the complexity of the 3D topology defined by the building

layout, yet computing the coverage, its redundancy for reliability, and the cost to achieve it is of paramount importance for the success of next generation mobile networks. gNBs should be moved from antenna sites and roof-based positions to building facades, light poles, or other suitable positions that grant LoS on streets and public spaces, yet there is a lack of scientific evidence on the feasibility and cost-effectiveness of achieving redundant LoS coverage in urban areas. Traditional coverage studies and analysis are based on Non Line-of-Sight (NLoS) propagation models devised for lower frequencies, where diffraction around corners and reflections have very different characteristics from mmWave. Modeling of 3D characteristics and cities is thus very different; moreover antennas at lower frequencies allow over-the-roof placement with good coverage on the ground, which is not possible at mmWave frequencies, completely changing the rules of network coverage and planning.

This paper uses a novel 3D approach to estimate the w -coverage, i.e., the coverage with multiplicity w , and the cost of its deployment. We study and adapt to the problem at stake heuristics developed for set coverage and minimum partial set multi-cover problems (both of them known to NP-complete). The solution to these problems returns the optimal placement of gNBs for LoS communications in urban areas, and we propose different metrics (or score functions) that allow adapting the heuristic solution to target different goals of the coverage.

The knowledge base are LiDAR-derived open datasets that describe the 3D shapes of buildings. This paper analyzes parts of 3 Italian cities (Trento, Firenze, and Napoli) where we selected 5 central areas with an average size of 0.7 km², for a total of 15 different locations. These three cities, even if all in Italy, have different characteristics, and allow a good insight and general conclusions stemming from the results obtained. Extension to other cities and urban areas is straightforward once the LiDAR-based data are available. In an Open Science effort, we publish both the code¹ and datasets [3] in public repositories.

We consider placing antennas on the building facades, and we use state-of-the-art ray-tracing techniques based on GPU computation to estimate the LoS coverage with its multiplicity on the ground area identified by public streets. gNBs are supposed to be able to cover the entire area seen by their placement site, possibly with more than a single antenna if required. In other words, a gNB on the wall of a building

R. Lo Cigno is with the Department of Information Engineering, University of Brescia, Italy (e-mail: renato.locigno@unibs.it)

G. Gemmi and L. Maccari are with the Department of Environmental Sciences, Informatics and Statistics, Ca’ Foscari University of Venice, Italy. (e-mail: gabriele.gemmi@unive.it; leonardo.maccari@unive.it)

¹<https://github.com/UniVe-NeDS-Lab/TrueBS/>

has a 180° visibility, while if mounted on the corner it has a 270° one. We devise a general algorithm that accepts different score metrics of the candidate gNB positions, thus enabling the exploration of different policies in pursuing a given w -coverage, where $w = 1, 2, \dots$ expresses the level of redundancy requested for the reliability. For instance, $w = 3$ instructs the algorithm to pursue a coverage where each point in the streets has visibility toward 3 gNBs.

Given the score function describing the coverage policy selected, we evaluate the coverage (the actual outcome of the algorithm), and the cost necessary to achieve it, offering significant insight on the problem of robust coverage for LoS communications, ranging from the effective number of gNBs required to achieve a given coverage (e.g., 90 or 95%) to building selection strategies to reduce the coverage cost.

This paper's contribution addresses two different aspects of coverage planning related to mmWave and in general ultra-high frequency technologies that require LoS for good performance.

The first one relates to the algorithms available to solve the multi-coverage problem to increase robustness and resilience of the network. We consider state-of-the-art algorithms derived from visibility analysis problems and adapt them to tackle communication networks' planning, showing how to modify them and how they affect the network design.

The second contribution is, to the best of our knowledge, the first quantitative analysis on real urban data on the required density of gNBs to achieve a robust and reliable coverage with mmWave links. The results we present show that in general between 60 and 120 gNBs per km^2 are sufficient for LoS coverage in public streets when the maximum distance between the user equipment and the gNB for LoS communications is set to 300 m [4]. This is consistent with the estimate of 100 gNBs per km^2 often mentioned by operators [1] and academics [5].

The remaining part of the paper is organized as follows. First, Sect. II reviews the state of the art and the work related to our contribution. Next, Sect. III formalizes the problem, introduces the notation we use, and reduces it to a visibility analysis problem, allowing the use of powerful tools and heuristics to solve it. Sect. IV describes the optimization objectives of coverage problems together with the constraints that limit the solution space. The following Sect. V and VI introduce heuristics at the state of the art that we adapt to the specific problem at stake, also extending the heuristics to reduce the cost of coverage. Sect. VII and VIII define the metrics describing the fitness and robustness of the solutions found and present numerical results for the three Italian cities we have considered. Sect. IX closes the paper with a final discussion.

II. STATE OF THE ART

The background of this contribution is rooted in different but related research areas: cellular networks, land and population coverage, foundational results on set coverage, and sensors and camera placement problems. In fact, the LoS requirement changes significantly the definition of the network coverage problem itself, which essentially becomes a problem of covering extremely complex topologies (the 3D maps of cities) that

are represented digitally, with a resolution that can be as good as one square meter or even better. With this representation the task of covering urban areas can be naturally mapped onto (very large) set covering problems, thus allowing the exploitation of powerful heuristics developed for this class of NP-complete problems. At the same time, the wealth of results and methodologies coming from the cellular networks' design must not be disregarded, even if NLoS coverage and LoS coverage are subject to different modeling solution techniques. Similar problems have also been addressed in the study of sensor and camera placement for monitoring purposes.

A. Cellular Networks Coverage

The placement of base transmitting stations has been widely studied since the deployment of the first mobile networks [6]. The densification of gNBs required by next generation networks has brought the topic back in the limelight, but the specific topic of LoS coverage of public areas has not received vast attention yet; even very recent works as [5] base their analysis on antenna directionality, but disregard the 3D geometry of the problem. Anjinappa et al. [7] investigate the optimal placement for gNBs and passive reflectors in two urban areas; however, they do not consider any cost model or try to account for the Capital Expenditure (CapEx) to deploy such a network. Two other works, [8] and [9] focus on the placement of wall-mounted gNBs and approach the problem using computational geometry in two dimensions, thus disregarding the impact of buildings and the height of the antennas. Zhang et al. [10] tackle the problem of gNBs placement trying to minimize the outage probability by studying a regular, Manhattan style urban topology, while Haile et al. [2] use an approach similar to ours, for NLoS communications. Another branch of research focuses on the placement of Unmanned Aerial Vehicles (UAVs) base stations [11] but has different requirements compared to our problem.

B. Set Covering Algorithms

The family of set covering problems are classical ones in computer science. Given a set Λ and a collection $\Omega = \{\sigma^j \mid j \in [0, N]\}$ where each σ^j is a subset of elements of Λ , the typical problem is to find a minimal subset $\Omega' \subset \Omega$ that covers the entire Λ . This problem is known to be NP-complete, but polynomial-time heuristics exist with bounded errors [12], [13]. In our case, Λ is the set of points on a public street and σ^j is the set of points covered by a gNB placed at a certain position. There are two differences compared to the classical problem. The first is that to achieve reliability we want to cover each element of Λ at least w times, and the second is that we do not impose total coverage, but only the coverage of a percentage of points. This is called a partial set multi-cover problem and is less studied than the previous one. Ran et al. address the difficulty of the problem [14], while recent works [15], [16] provide heuristics with bounded error that still require solving an integer linear problem at each step, which is computationally impossible in our context.

C. Sensors and Camera Placement

The research on sensor networks has addressed a similar problem. Sensors are assumed to be able to perceive events in a certain area around them, and one problem is to find the minimal deployment of sensors that cover each point at least w times. This problem is not new [17] and received a lot of attention [18], but the constraints and the solutions are very different from network coverage: The visibility regions of urban public spaces are not unit circles nor are they stochastic regions; furthermore, there is no reason to keep points connected because they are not communicating one another, and our points are contiguous and not randomly distributed on a 2D space. All the available solutions are based on geometrical considerations that do not apply to our case.

A last area of research we mention as part of set covering is the problem of placing the minimal set of cameras to monitor a certain area, or a certain number of objects in space [19]. The problem is only partly similar because cameras can be oriented, they can pan and zoom, and deployments are in the order of tens of cameras, while we have tens of thousands of potential locations. However, given the similarity to network coverage, we adapt the heuristic proposed in [20] that tries to achieve a fair coverage of selected targets from multiple cameras.

III. PROBLEM STATEMENT

The optimal placement of gNBs for LoS communications requires an efficient solution to determine the visibility between any two points. The solution to the problem is enabled by two innovations. First, public administrations publish 3D shapes of cities obtained using Laser Imaging Detection and Ranging (LiDAR) technology as open data (generically called Digital Elevation Models (DEMs)). These models can reach a horizontal precision down to 30 cm and a vertical precision in the order of 10 cm. Second, GPU computing allows the efficient computation of LoS existence between two points in a 3D space for thousands of potential gNBs on a high precision DEM at reasonable costs.

In our research, we take advantage of visibility analysis, a methodology that, given a point in space, provides the points on the DEM that are in LoS. The outcome of the visibility analysis is usually called *viewshed*, and the algorithms are called *viewshed algorithms*. We have implemented a viewshed algorithm [21] using Numba [22], a Python compiler that allows producing object code for Nvidia GPUs. We have introduced this technique in preparatory works to study the connections between gNBs [23] while here we describe it applied to ground coverage.

The DEM \mathcal{D} we consider is a set of points in space expressed by their 3D coordinates, with a sampling precision of 1 point per squared meter (for the main notation used in the paper see Tab. I). The coordinates are relative to the area, so a point p in the DEM is given by its coordinates (x, y, z) where x and y are integer values and z is a real number. \mathcal{D} is represented by a matrix \mathbf{E} of dimension $m_x \times m_y$ where $E_{x,y}$ is the element of the matrix of indices x and y , that is, the elevation of the point with planar coordinates (x, y) . \mathbf{E} is a

Symbol	Type	Definition
\mathcal{D}	Set	Digital Elevation Model (DEM)
\mathbf{E}	Matrix	Matrix Representation of DEM
p_i, p_j	Vector	Generic points in a 3D space
$Y(p_i, p_j)$	Function	LoS function between p_i and p_j
$d(p_i, p_j)$	Function	distance between p_i and p_j
\mathbf{V}^i	Matrix	Viewshed from p_i
s^r	Set	The points of the perimeter of building r
$A(s^r)$	Function	Function that returns the points inside s^r
R	Scalar	Number of buildings
\mathbf{S}^r	Matrix	Rasterized building shape
$\delta()$	Function	Dilation
\mathbf{P}^r	Matrix	Rasterized perimeter of the building r
C^r	Set	Set of potential coordinates on the border of building r
Λ	Set	(x, y) coordinates of street points to cover
σ^i	Matrix	Ground visibility from the point p_i
Ω	Set	Ground visibility matrices
$\Gamma()$	Function	Maximum coverage heuristic

TABLE I: Symbols and notation used in the paper. In general scalars and vectors are standard math characters and matrices' names are bold. Sets are calligraphic letters, elements of a matrix use the same name of the matrix and appropriate pedices, e.g., $\sigma_{x,y}^i$ is an element of matrix σ^i .

bi-dimensional matrix, and every point p in \mathcal{D} corresponds to a triplet $(x, y, E_{x,y})$. In some cases, we may refer to a point p elevated from the 3D surface ($z > E_{x,y}$), the context will make it clear. Since the space of (x, y) coordinates is sampled we can enumerate the points, and we often refer to points p_j or p_i .

Given \mathcal{D} let Y be a function that takes in input two points p_i and p_j and returns 1 if there is LoS between p_i and p_j or 0 otherwise. Computing $Y(p_i, p_j)$ is not a trivial task, as it requires calculating the segment between the two points and checking if it intercepts the surface identified by \mathcal{D} , zero times (thus, there is LoS) or an even number of times as p_i and p_j are on or above the surface identified by \mathcal{D} (then there is no LoS).

Given a point p_i , we can define the viewshed from p_i as the binary matrix \mathbf{V}^i of size $m_x \times m_y$ so that for every point $p_j \in \mathcal{D}$ of coordinates $(x, y, E_{x,y})$ it holds:

$$V_{x,y}^i = Y(p_i, p_j), \forall p_j \in \mathcal{D} \quad (1)$$

If p_i is the position of a gNB then \mathbf{V}^i represents the points in \mathcal{D} that have direct line of sight with the gNB. We say that point p_j is covered by gNB i if $V_{x,y}^i = 1$ and we say that p_j has w -coverage if, given a certain placement of gNBs, it is covered by at least w gNBs. The goal of our contribution is to study the deployment of gNBs that maximizes the w -coverage of users' equipment placed at 1.5 meters from the ground given a specific density of gNBs per km^2 .

In this initial work we only consider buildings, disregarding the effect of trees, lamps, vehicles and any other static or moving object that is not found in the public datasets. This choice leads to an optimistic evaluation of the LoS probability, which is instead influenced also by these objects as studied recently in [24]. Adding a probabilistic LoS model to our study is not difficult, but, unless the model can be tuned with realistic parameters for the area, it can also be misleading. The

results we obtain and present in VIII are thus related to the *coverage* and not to the (complement of) *outage probability*: They represent the upper-bound of the probability that a user in the area can receive the service through a mmWave communication interface.

IV. PROBLEM CONSTRAINTS AND OPTIMIZATION GOALS

As LoS communications prevent placing gNBs on roofs, we assume their deployment on the facades of buildings, a location that is already used by telecommunication companies to install fiber-to-the-home splitters and 4G small cells.

Let $A(s^r)$ be the area contained in a polygon s^r that represents the 2D perimeter of a building, which we obtain from open datasets, and let R be the total number of buildings in the area. Given a building r whose 2D perimeter is represented by s^r , we implement a function that transforms the polygon into a binary matrix \mathbf{S}^r of the same dimensions of \mathbf{E} , so that:

$$S_{x,y}^r = 1 \iff (x, y) \in A(s^r) \quad (2)$$

In order to isolate the border points of the building, we use a morphological operation $\delta(\cdot)$ called dilation, which dilates the binary matrix of the polygon by one unit in every direction [25] and sets to 1 the corresponding coordinates. Then given the matrix \mathbf{S}^r we define its perimeter using a matrix \mathbf{P}^r as:

$$\mathbf{P}^r = \delta(\mathbf{S}^r) - \mathbf{S}^r \quad (3)$$

Figure 1 shows the graphical representation of the matrixes \mathbf{S}^r and \mathbf{P}^r .

We assign to every point on the perimeter a height from the ground z_r that is the minimum between the height of the building minus one meter and 10m from the street level, and finally we define the set of coordinates C^r as:

$$C^r = \{(x, y, z_r) \mid P_{x,y}^r == 1\} \quad (4)$$

Eq. (4) identifies a set of potential positions for gNBs on the facade of building r . Each point lies on the border of the buildings and it is placed at the height of the building, or at 10m from the ground if the building is taller than 10m (as the ITU recommendations for micro cells in urban areas suggest [26]). In every area we have roughly 500 to 700 buildings, the set of points $p_i \in \{\cup_r C^r\}$ is the set of potential locations where to place our gNBs. As the average number of points per building perimeter is roughly 100 we have a total of 50,000 to 70,000 potential gNB locations.

A. Defining the Ground Visibility Matrix

We use street maps provided by Openstreetmap, where each road is identified by a mono-dimensional line that represents the center of the street. We expand the line to make it a 2D surface and we call Λ the set of all (x, y) coordinates of the points on the street, obviously $(x, y) \in \Lambda \rightarrow (x, y) \notin A(s^r) \forall r < R$. Let us now consider all points p_j of coordinates $(x, y, E_{x,y} + 1.5)$ where $(x, y) \in \Lambda$. The point p_j is a point in the street elevated by 1.5m from the ground, a common assumption of the position of a mobile terminal. Assume p_i




Fig. 1: Dilation of a building r with its perimeter points \mathbf{P}^r in blue (each pixel represents a point) and the rasterized building shape \mathbf{S}^r in black.

is the position of a gNB, we call σ^i a $m_x \times m_y$ matrix whose elements are defined as:

$$\sigma_{x,y}^i = \begin{cases} 0 & (x, y) \notin \Lambda \\ 0 & d(p_i, p_j) \geq d_{\max} \\ \Upsilon(p_i, p_j) & \text{otherwise} \end{cases} \quad (5)$$

where $d(p_i, p_j)$ is the euclidean distance between two points and d_{\max} is a technology-dependent, arbitrary maximum communication distance. σ^i is a binary matrix that represents all the points in the street elevated of 1.5m that are in line of sight with the gNB placed at p_i , and that are at a maximum distance d_{\max} . We generally refer to σ^i as the ground visibility matrix from point p_i . Figure 2 shows the graphical rendering of a matrix σ^i .

Given all the potential locations of gNBs, i.e., the union of all the sets C^r , let Ω be the collection of all potential visibility matrices:

$$\Omega = \{\sigma^i \mid p_i \in \{\cup_r C^r\}\} \quad (6)$$

and let k be the number of gNB we want to install in that specific area. The problem can be expressed as the search of a subset $\Omega' \subseteq \Omega$ that maximizes w -coverage, with $|\Omega'| \leq k$.

The algorithm used to compute the viewshed from p_i [27] has a worst case complexity $O(|E|^2)$ as it has to check all the points in the DEM. However, we limit the maximum distance of each ray to d_{\max} . This leads to a complexity $O(d_{\max}^4)$ for each viewshed. Since we compute a different viewshed for each potential location in the set C^r the overall complexity for the computation of the set Ω is $O(|C^r| \cdot d_{\max}^4)$. We start by defining 1-coverage, which is intuitive, and we then extend the approach to arbitrary values of w .

B. Maximizing 1-coverage

Let $|\cdot|_0$ be the L_0 norm of a matrix (the number of non zero elements); the problem of maximizing 1-coverage can be formalized as the search of a subset $\Omega' \subseteq \Omega$, such that

$$|\Omega'| \leq k \quad \text{and} \quad \left| \sum \sigma^i \in \Omega' \right|_0 \quad \text{is maximised.} \quad (7)$$

This is a classical maximum coverage problem, in which we have a collection Ω of sets, each set has elements in Λ and given k , we need to find the k sets that cover the largest number of elements of Λ . Being a set covering problems, it is




Fig. 2: Ground coverage σ^0 (light green) from a gNB in point p_0 (black) in a portion of the city of Florence. Figure shows also the buildings' shapes s^r (red), the set of points on the street Λ (grey) and other gNB locations (purple).

NP-complete and can not be solved exactly when $|\Omega|$ and $|\Lambda|$ are in the order of tens of thousands. We leverage a polynomial heuristic with bounded error [28] that allows finding a semi-optimal solution in a reasonable time, extending it to fit our problem. Since the maximum coverage problem can be seen as a special case of maximization of submodular functions with a cardinality constraint, we can state that the error of this polynomial heuristic is bounded and it matches the theoretical bound [12] [29].

C. Maximizing w -coverage

The interpretation of Eq. (7) is straightforward: given a number k of gNBs that can be deployed to cover a certain area, we want to have visibility from at least 1 gNB to the largest possible number of points in the public streets.

The 1-coverage is sufficient to claim to be covering an area, but does not provide reliability because there can be obstacles that obstruct the LoS with the gNB, the most obvious one being the person that holds the terminal, but also other people, cars, trucks, and so forth. The goal of a good coverage plan must be to provide 1-coverage to the largest area, but also a reliable service to mobile users, thus the problem is better formulated as a w -coverage, with a suitable w .

Maximizing w -coverage can be formulated as finding Ω' such that:

$$|\Omega'| \leq k \quad \text{and, given} \quad \Phi = \sum \sigma^i \in \Omega' \quad (8)$$

$$|\{\Phi_{x,y} \geq w\}| \quad \text{is maximised,}$$

where Φ is a matrix populated with the multiplicity of the coverage for each point.

This is an extension of the set covering problem known as *minimum partial set multi-cover problem*, it is NP-complete [16], and no efficient heuristics are yet known. To solve this problem, we start from a heuristic proposed for multi-camera visibility [20] and, after analyzing the specific problem and goal, we propose a simpler one with comparable performance.

V. A GENERIC HEURISTIC FOR w -COVERAGE

The goal of w -coverage is to guarantee a robust and efficient coverage, but in communications it has many subtleties that need to be considered, in particular, the problem that a

Algorithm 1 Computation of approximate solution to the maximum coverage and minimum partial set multi-cover problems

Require: Ω (Set of viewsheds), k (number of gNB), w (w -coverage)

Ensure: Ω^* (Set of the viewsheds from optimal locations)

```

1: procedure  $\Gamma(\Omega, k, w)$ 
2:    $\mathbf{C} = \mathbf{0}$ 
3:    $\Omega^* = \{\}$ 
4:   for  $i \leftarrow 0$  to  $k$  do
5:      $h^* = -\infty$ 
6:     for  $\sigma^j \in \Omega$  do
7:        $\mathbf{C}^* = \mathbf{C} +_w \sigma^j$ 
8:        $h_j = \rho(\mathbf{C}^*, w)$  ▷ Calculate score
9:       if  $h_j > h^*$  and  $\sigma^j \notin \Omega^*$  then
10:         $\sigma^* = \sigma^j$ ;  $h^* = h_j$ 
11:      end if
12:    end for
13:     $\mathbf{C} = \mathbf{C} +_w \sigma^*$  ▷ Update covered elements
14:     $\Omega^* = \Omega^* \cup \{\sigma^*\}$ 
15:  end for
16:  return  $\Omega^*$ 
17: end procedure

```

complete w -coverage in real cities is probably impossible or simply too costly to pursue, so one has to take into account also the points where the coverage is smaller than w , solving dilemmas as, for instance, *is it better to add an antenna that improves the 1-coverage in n points or one that improves the w coverage in m other points?*

To achieve this goal we wish to have a single approximation algorithm where different metrics (or score functions) can be applied to explore different balances of the coverage. Algorithm 1 describes the heuristic we propose to tackle the w -coverage problem, extending the heuristic algorithm for set coverage proposed in [28].

Let us first introduce the notation used in the algorithm. Let $\mathbf{0}$ be the matrix of dimension $m_x \times m_y$ of all zeros. We also introduce the *sum-capped-to- w* operator with symbol $+_w$ so that given matrices \mathbf{B}, \mathbf{D} , then $\mathbf{A} = \mathbf{B} +_w \mathbf{D}$ is a matrix whose elements are: $\mathbf{A}_{i,j} = \min(w, \mathbf{B}_{i,j} + \mathbf{D}_{i,j})$.

Algorithm 1 works as follows. It takes in input a set of ground visibility matrices Ω , each one corresponding to a point p_j where a gNB can be placed, a number k of gNBs to be placed, and the desired value for w . It returns a subset $\Omega^* \subset \Omega$ corresponding to the set of the viewsheds (1-t-1 mapped to gNB positions) that provide a close-to-optimal w -coverage. Line 2 initializes the matrix \mathbf{C} at $\mathbf{0}$. The points on public streets are considered by the algorithm, the others remain untouched at zero. Line 3 initializes the set Ω^* to an empty set. Every iteration of the loop beginning at Line 4 chooses a new viewshed σ^j (corresponding to a possible gNB location) to be added to Ω^* . To select the viewshed above, every iteration of the loop beginning at Line 6 evaluates the additional coverage offered by every possible viewshed, i.e., every possible candidate position for a gNB. \mathbf{C}^* is defined at Line 7 as the matrix of the points covered if the candidate

viewshed σ^j is added to the set. Every element of \mathbf{C}^* is thus the result of the element by element sum capped to w of the already existing coverage plus the viewshed σ^j . The score function $\rho(\mathbf{C}^*, k)$ at Line 8 is the core of the algorithm that actually defines the metric by which candidate gNB locations are ranked. Sect. V-A discusses three different score functions we propose and evaluate.

The score functions are defined so that a better coverage obtained by adding σ^j corresponds to an algebraically higher value of h_j . Line 9 checks the score and, if the score for candidate σ^j is greater than the largest found so far, the candidate σ^* and its corresponding score h^* are updated. Finally, Lines 13 and 14 update the coverage \mathbf{C} and the list of selected viewsheds Ω^* ; at the end of the algorithm, Ω^* is returned, also yielding the selected gNB positions.

A. Score Functions

The final outcome of Algorithm 1 depends on the score function $\rho()$, which specifies, step-by-step, the metric for the best σ^j selection. These functions express how the algorithm will push towards 1-, 2-, ... or w -coverage at each iteration. We introduce here three functions that assign a different importance to different levels of coverage; Sect. VIII analyzes their impact on the final outcome.

1) *w-coverage Maximization (w-CM)*: The first score we consider is a simple 1-norm (sum of all the elements) on the candidate coverage \mathbf{C}^* computed at Line 7 in Algorithm 1:

$$\rho(\mathbf{C}^*, w) = |\mathbf{C}^*|_1 \quad (9)$$

For $w = 1$ applying Eq. (9) in Algorithm 1 solves the canonical maximum set coverage problem. For $w > 1$ it returns the global weight of \mathbf{C}^* elements, thus Algorithm 1 ends up in maximizing the weighted coverage, without any attempt to prioritize the points that have zero coverage over the ones that have coverage between 1 and $w - 1$. Notice that \mathbf{C}^* are capped to w , so that points whose coverage is larger than w do not influence the score. Besides its simplicity, this naive function completely disregards the difference between a point that is not served at all and one that already has coverage, but whose coverage is improved by adding the considered viewshed.

2) *w-coverage Fairness (w-CF)*: Reducing the difference between different levels of coverage resembles a problem of fairness, thus we recall the well-known fairness index proposed by Ray Jain [30], that applied to our problem can be written in terms of matrix norms as

$$\text{FI}(\mathbf{C}^*) = \frac{(|\mathbf{C}^*|_1)^2}{|\Lambda|(|\mathbf{C}^*|_2)^2} \quad (10)$$

the 2-norm is the square root of the sum of the squared elements of the matrix.

The fairness pushes all elements of \mathbf{C}^* to a similar value, but alone does not provide a useful score, since we are interested in a wide, fair coverage and not only in a fair one (all 0-coverage is perfectly fair). This problem has been analyzed in the context of visual coverage in sensor networks in [20], and we adopt the same cost function proposed there, which,

adapted to the problem and with the notation of this paper can be formalized as

$$\rho(\mathbf{C}^*, w) = \text{FI}(\mathbf{C}^*) \times \frac{|\mathbf{C}^*|_1}{|\Lambda|w} = \frac{(|\mathbf{C}^*|_1)^2}{|\Lambda|(|\mathbf{C}^*|_2)^2} \times \frac{|\mathbf{C}^*|_1}{|\Lambda|w} \quad (11)$$

$\text{FI}(\mathbf{C}^*)$ is weighted by the ratio between the 1-norm of the 1-coverage and the target w -coverage ($|\Lambda|w$), so that the first factor of Eq. (11) tries to balance the coverage and the second tries to extend it.

3) *w-coverage Gap (w-CG)*: The score defined by Eq. (11) is composite (a multiplication of two factors aiming at different goals), and its interpretation not always straightforward, thus we propose a third, simple metric whose goal is to weight the gap between the actual coverage and the target coverage with a quadratic function. Let $\mathbb{1}$ be the $m_x \times m_y$ matrix with 1 in all the positions that need to be covered and 0 otherwise. The target coverage can be expressed simply as $\mathbb{1} \cdot w$, and the score becomes

$$\rho(\mathbf{C}^*, w) = -(|\mathbb{1} \cdot w - \mathbf{C}^*|_2)^2 \quad (12)$$

We use the squared value of $|\cdot|_2$ for computational efficiency. Since this score measures a gap from the intended coverage, we use its negative value so that Algorithm 1 remains consistent in selecting the algebraic maximum.

When $w = 1$ Eq. (9) and Eq. (12) yield exactly the same result, while Eq. (11) returns a different numeric result (values are divided by $|\Lambda|^2$) but the obtained ranking is the same of the other two. For this reason, Sect. VIII presents only 1-CM results, while for other values of w all metrics are reported.

B. Complexity

Algorithm 1 is composed of two nested loops, the first one iterates over the number k of gNBs that are to be deployed and the second one iterates over the number of potential locations $|\Omega|$. At each inner loop the algorithm computes a score h_j by calling the function ρ . All score functions loop iteratively over all the values of $|\mathbf{C}^*|$, so they have time complexity $O(|\mathbf{C}^*|)$. This leads to a worst case time complexity of $O(k|\Omega||\mathbf{C}^*|)$.

The memory required to execute the algorithm is bounded by the number of ground points (Λ), as the algorithm allocates a 2-dimensional matrix of size $(|\Lambda|, |\Omega|)$ (Line 2). At Line 7 we allocate a new temporary matrix \mathbf{C}^* of the same size. Since we only need to store a boolean value for each cell of the matrix, we can use the smallest datatype available, which is `uint8`². Thus the memory footprint is $2|\Lambda||\Omega|$ bytes. The memory footprint of ρ is constant.

VI. THREE-STEP HEURISTIC

The heuristic presented in the previous section finds what we call the semi-optimal placement for the gNBs. Given a certain number k of gNBs that the operator can afford to place, it provides the best locations where these gNBs should be placed. Iterating on values of k , we can answer the question: *what is the minimum number of gNBs and their positions to obtain w-coverage of a certain fraction of points on the street, e.g. 95%?*

²Theoretically a single bit could be used, but the Numba architecture does not support it.

However, Algorithm 1 treats every σ^j in Ω the same way, irrespectively of the building it corresponds to. This means that if convenient, the algorithm can place one gNB per building, or even all gNBs in the same building. Albeit theoretically possible, these configurations would be practically impossible or not cost-effective. To produce realistic results we need to introduce two practical constraints based on the cost estimation of a gNB deployment.

In a study done in 2020, Oughton et al. calculated the cost of a 5G small cell to be 20100\$, of which only 3380\$ were accountable to the radio³ while the rest includes several fixed costs for constructions, hardware, backhaul fiber network, etc. that are needed to realize the gNB [31]. We assume that a multi-radio small cell can be realized keeping the fixed cost constant, and increasing the number of radios up to a maximum of 5 per gNB. This is a free parameter of our algorithm and can be adjusted to any other value.

We designed a heuristic that applies Algorithm 1 in three steps to limit the number of buildings on which to place gNBs and the number of radios per gNB: First, for every building, it finds the best 5 points that maximize the building visibility; Second, it finds the top g buildings ordered by their cumulative visibility provided by the best 5 points; Third, among the best 5 points of the top g buildings, it finds the overall best k locations for the gNBs. The choice of g is key for the final result, in fact, if the target of the operator is to provide w -coverage to a certain percentage of the points on the ground, if g is too small such objective may not be reachable. On the other hand, the larger is g , the higher is the cost of the deployment. The optimal g is of course scenario-dependent and can be found iterating on g with our heuristic. Since reducing the number of buildings (and thus the cardinality of Ω) makes the 3-step heuristic faster than the original one, the optimal g for an area of roughly 0.7km² with 135 gNB/km² can be computed in around 1 minute.

Note that the formulation of the $\Gamma()$ function in Algorithm 1 is extremely generic as it takes a set Ω of viewsheds as an input, and it returns another set of viewsheds. We can thus change the input Ω and k and apply the algorithm three times obtaining semantically different results, without the need to modify its internals.

A. First step: Point selection

At the first step we limit Ω to the set of all points from a single building. Let Ω^r be the set of all the visibility matrices from points on the (dilated) perimeter of builing r , then:

$$\Omega^r = \{\sigma^j \forall p_j \in C^r\} \quad (13)$$

We apply $\Gamma()$ using Ω^r as the first argument, and limiting the choice to 5 points. We obtain a list of viewsheds from the best 5 points in the building.

$$\Omega_5^r = \Gamma(\Omega^r, 5, w) \quad (14)$$

We repeat this step for every building r .

³These costs have been converted to Dollars (\$) for the sake of readability, in the original research they were expressed in British Pounds (£).

B. Second step: Building Selection

Let Φ^r be the ground visibility matrix of building r , i.e., the visibility matrix from the best 5 points on the facade of building r . Φ^r is defined as the logical OR (\vee) of all the binary visibility matrices:

$$\Phi^r = \bigvee_{\sigma^j \in \Omega_5^r} \sigma^j \quad (15)$$

and let $\Omega^R = \{\Phi^r \forall r < R\}$ be the set of all the visibility matrices from all buildings. Invoking the function $\Gamma()$ using Ω^R and g as an input, we obtain the cumulative viewsheds from the best g buildings. Since the areas in which we run the experiments are different and contain a different number R of buildings, in practical terms it is more convenient to refer to the percentage X of buildings to be used, where $g = \lceil \frac{X}{100} * R \rceil$. We then define:

$$\Omega_X^R = \Gamma\left(\Omega^R, \left\lceil \frac{X}{100} * R \right\rceil, w\right) \quad (16)$$

that is the set of the best cumulative viewsheds from a percentage X of the buildings of the area.

C. Third step: gNB Selection

Step two returns a set of cumulative viewsheds obtained considering the best 5 points on $X\%$ of the buildings, and we use it to select the set B_X of the indices of the best buildings:

$$B_X = \{r \forall r \mid \Phi^r \in \Omega_X^R\} \quad (17)$$

and thus we can define the input to the third and last step of the heuristic:

$$\Omega_X = \{\sigma^j \in \Omega_5^r \forall r \in B_X\} \quad (18)$$

Ω_X contains all the possible viewsheds from each of the best 5 points (selected in step 1), from $X\%$ of the buildings that have the best cumulative viewshed (selected at step 2). As an example consider an area with 600 buildings each of which has an average of 100 points on its border. If $X = 10$ then $g = 60$, and the third step of our heuristic would explore at most $|\Omega_X| = 60 \times 5 = 300$ locations out of the 60.000 available.

Finally, we apply $\Gamma()$ again and obtain:

$$\Omega_{X,k}^* = \Gamma(\Omega_X, k, w) \quad (19)$$

Ω_X^* contains exactly k viewsheds that correspond to k points in space where to place gNBs. These places represent the semi-optimal choice among Ω_X , with a maximum of 5 radios per building and $X\%$ of buildings.

Similarly to what we did passing from g to X , it is more practical to target a desired density λ of gNBs per squared km, rather than a number k . So, given a certain area and λ we have:

$$\Omega_{X,\lambda}^* = \Gamma(\Omega_X, \lambda * \text{area}, w) \quad (20)$$

We are now able to compare the coverage obtained in heterogeneous areas with a target density of gNBs per squared km, using X as a tuning parameter to obtain the most favorable cost/performance trade-off.



Fig. 3: Graphical rendering of the area chosen for the experiment (in dark grey), its buffer of radius $d_{\max}/2$ (in light grey), and the set of points on the street Λ in black.

Note that while the original heuristic has a bounded error, our three steps heuristic loses this theoretical property. However, when used with $X = 100\%$ the effect is to pick the best 5 points for each building, and then apply the bounded heuristic to this restricted set of points. This reduces the dimension of the problem, but also corresponds to a practical constraint in real deployments, since the number of gNBs per building cannot grow arbitrarily. The application of the three-step heuristic with $X = 100\%$ thus maintains the bounded error, and Sect. VIII shows that results with $X = 4\%$ are very close with a strong reduction of the cost of the infrastructure. Finally, since Algorithm 1 has polynomial complexity, its three-step application is again of polynomial complexity.

VII. EXPERIMENTS SET UP AND METRICS

As already mentioned we evaluate 5 areas for each of the three densely populated Italian cities for a total of 15 scenarios. Trento, Firenze, and Napoli, however, have quite different urban textures, so they represent a valid sample of different urban areas. The average size of the areas is roughly 0.7 km^2 and is limited by the computational power we had access to, as obtaining $\Omega_{X,\lambda}^*$ is a computationally intensive task that requires handling tens of thousands of large integer matrices. To speed-up the algebra on large matrices we implemented the algorithms on an NVIDIA V100 GPU equipped with 5120 CUDA cores, which enables parallelization of matrix operations. The limit is given by the 16 GB of RAM of the GPU which limits computation to areas of up to 1 km^2 . Results can be extended to larger areas straightforwardly with more resources, and with some effort by improving the memory space optimization.

Placing the gNBs in a constrained area causes a border effect: Points on the border of the area can be covered only by the gNBs placed inside the area, but not from the ones outside the area, which would be present in the real world coverage. For this reason, we enlarge the area where gNBs can be placed with a guard band of width $d_{\max}/2$ m as shown in Fig. 3 (light gray area), but we measure coverage only in the inner area (dark gray area).



Fig. 4: Graphical rendering of the coverage $\Phi_{4,45}$ (shades of green), and the set of optimal points (black dots) in a portion of one of the considered areas, in Napoli.

A. Evaluation Metrics

To evaluate the performance of different algorithms and their parameters we calculate the effective coverage, that is the sum of all the viewsheds $\sigma^i \in \Omega_{X,\lambda}^*$:

$$\Phi_{X,\lambda} = \sum_{\sigma^j \in \Omega_{X,\lambda}^*} \sigma^j \quad (21)$$

This matrix, with the same dimensions of \mathbf{E} , contains values ranging from 0 (no LoS with any gNB) to $k = \lambda * \text{area}$ (LoS with all gNBs). Fig. 4 shows the effective coverage on a sample area of Napoli.

Every experiment is repeated varying λ , X and w as reported in Tab. II; in graphs we report the confidence interval on the 15 areas with confidence level $\alpha = 0.95$ as error bars.

Parameter	Value	Description
λ	{15, 45, 75, 105, 135, 165, 195 }	gNB density
X	{2, 4, 100}%	Percentage of buildings
d_{\max}	300 m	Maximum distance for the links [4]
c_{gNB}	16720 \$	Cost of a gNB without any radio [31]
c_{radio}	3380 \$	Cost of each radio of a gNB [31]
w	{1,3}	Target coverage

TABLE II: Parameters used for the numerical results presented in the paper.

1) *Coverage*: The coverage metric, analyzed as a function of X , counts the number of points with w -coverage normalized by the total number of points on the public street;

$$c_w(\lambda) = \frac{\gamma_w(\Phi_{X,\lambda})}{|\Lambda|} \quad (22)$$

where $\gamma_w(\Phi_{X,\lambda})$ is a function that counts the number of elements of $\Phi_{X,\lambda}$ that are greater or equal than w . In the numerical results we show $c_1(\lambda)$, which expresses the portion of streets covered by at least one gNB, and $c_3(\lambda)$ the portion of streets covered with higher reliability (at least 3 gNBs). We also show the coverage distribution on the points of Λ , i.e., the Empirical Probability Density Function (e-pdf) of the coverage values.




Fig. 5: Ground projection of two examples of 3-coverage. Black dots are gNBs.

2) *Resistance to Obstruction*: Fig. 5 depicts the ground projection of the rays that connect a point p to three gNBs in two different cases. In both scenarios, an obstacle that obstructs an angle $\alpha_2 + \alpha_3$ may totally shadow p , but in the right figure this is way more likely than in the left one, so w -coverage alone does not necessarily imply resistance to obstruction. We define $\alpha_c(p) = 360^\circ - \max(\alpha_1, \alpha_2, \alpha_3)$ to be a measure of robustness against shadowing. It is easy to show that one single object that obstructs an angle smaller than $\alpha_c(p)$ cannot totally prevent LoS with some gNB, so the larger is $\alpha_c(p)$ the better it is. Let Λ_2 be the set of points with coverage larger or equal than 2, then

$$O_R = \frac{|\{p \in \Lambda_2 \mid \alpha_c(p) > 45^\circ\}|}{|\Lambda_2|} \quad (23)$$

provides the fraction of points for which one obstacle that obstructs 45° does not prevent LoS. In the absence of finer metrics, O_R provides a heuristic measure of how well the coverage can resist to shadowing.

3) *Total Cost*: The cost function we use to evaluate the CapEx for the deployment of the network follows a cost model conceived for 5G networks [31]. From the model we extract a fixed part c_{gNB} for the deployment of the gNB in a building, and a variable part for each radio interface of the node (c_{radio}):

$$\text{cost}(\lambda) = c_{\text{gNB}} * N_{\text{gNB}} + c_{\text{radio}} * k \quad (24)$$

where N_{gNB} counts the number of buildings on which at least one gNB has been deployed and $k = \lambda * \text{area}$. The $\text{cost}()$ value depends on the desired density of gNBs per km^2 , so it is expressed as a function of λ , but it is also influenced by X . We also define upper and lower bounds of the cost function, as:

$$\begin{aligned} \text{UB} &= (c_{\text{radio}} + c_{\text{gNB}}) \cdot k \\ \text{LB} &= c_{\text{radio}} \cdot k + c_{\text{gNB}} \end{aligned} \quad (25)$$

The upper bound UB is given by the worst case where every coverage point is a single antenna gNB on a different building, while the lower bound LB is given by the (practically impossible) network where all radios are placed on a single building.

4) *Marginal Cost*: To estimate the Return of Investment (RoI) of deploying a robust network, we also evaluate the cost-effectiveness of adding new gNBs to further improve the coverage. We evaluate the cost of a deployment with growing densities λ_i . Since the heuristic is deterministic, if $\lambda_{i-1} < \lambda_i$ then $\Omega_{X, \lambda_{i-1}}^* \subset \Omega_{X, \lambda_i}^*$, and the difference $\text{cost}(\lambda_i) - \text{cost}(\lambda_{i-1})$ is exactly the marginal cost of the added gNBs. Similarly,

$c_w(\lambda_i) - c_w(\lambda_{i-1})$ is the marginal increase of relative coverage. Thus we can define the incremental cost metric m_c :

$$m_c(\lambda_i) = \frac{\text{cost}(\lambda_i) - \text{cost}(\lambda_{i-1})}{c_1(\lambda_i) - c_1(\lambda_{i-1})} \quad (26)$$

that provides an efficiency metric ($\$/\text{m}^2$) to estimate how cost-effective it is to increment the density of gNBs to improve coverage.

VIII. RESULTS

We present the results following the same order we introduced for the metrics: coverage, total cost, and marginal cost.

A. 1-Coverage

Fig. 6 reports $c_w(\lambda)$ for $X = 2, 4, 100\%$ and $w = 1, 3$. Each point is the average of all 15 areas we consider; vertical bars are the 95% confidence intervals. Since areas are never overlapping and thus they are independent, we can safely consider their average coverage as the outcome of an i.i.d. random variable, thus the confidence interval can be computed as a Student-t distribution with 15 degrees of freedom. The confidence intervals are reasonably compact, thus we deem that the results are reliable.

If we focus on the first row (1-coverage) the plots show one key point of our analysis, that it is possible to obtain a very large c_1 in urban areas with λ below, or very close to 100. This confirms that in urban areas a very high LoS coverage is achievable with a number of gNBs that is close to what is expected (roughly ten times what is used for 4G, so $\lambda \simeq 100$). We see in fact that with $X = 100$ and with $X = 4$ all the curves reach $c_1 = 80\%$ ($\lambda = 45$), $c_1 = 90\%$ ($\lambda = 75$), and 1-CM and 3-CG reach $c_1 = 95\%$ ($\lambda = 105$). We stress that these results were obtained with realistic assumptions: a limited, realistic number of devices per building and a precise ray-tracing model based on real data. All curves follow a trend expected in a set covering problem, with a steep initial rise followed by a saturation phase.

1-CM provides the highest 1-coverage, which is consistent with its design, while it is important to note that 3-CG, when evaluated on c_1 , dominates the other 3-coverage strategies on values of $c_1 > 80\%$. The difference is not very large, but it is consistent in the saturation phase, in which it is extremely expensive to gain even a single percentage point, thus this very simple score function serves well both coverage and robustness. We also see that with $X = 2\%$ all strategies stay below $c_1 = 90\%$, while as said, with $X = 4\%$ we can achieve results that are very close to $X = 100$.

1) *n-Coverage and Reliability*: Let us now focus on the 3-coverage row, which confirms that $X = 2\%$ is outperformed by the other two configurations (and thus, we do not comment on results with $X = 2\%$ anymore). It is evident that a remarkable difference between 1-CM and the 3-coverage strategies exists. Among the latter ones, 3-CM is the one that performs slightly better than others, as it focuses only on 3-coverage and disregards other possible goals. In this case, the difference in the saturation point between $X = 4$ and $X = 100$ is more evident, as with $X = 4$ the highest 3-coverage is 84%.

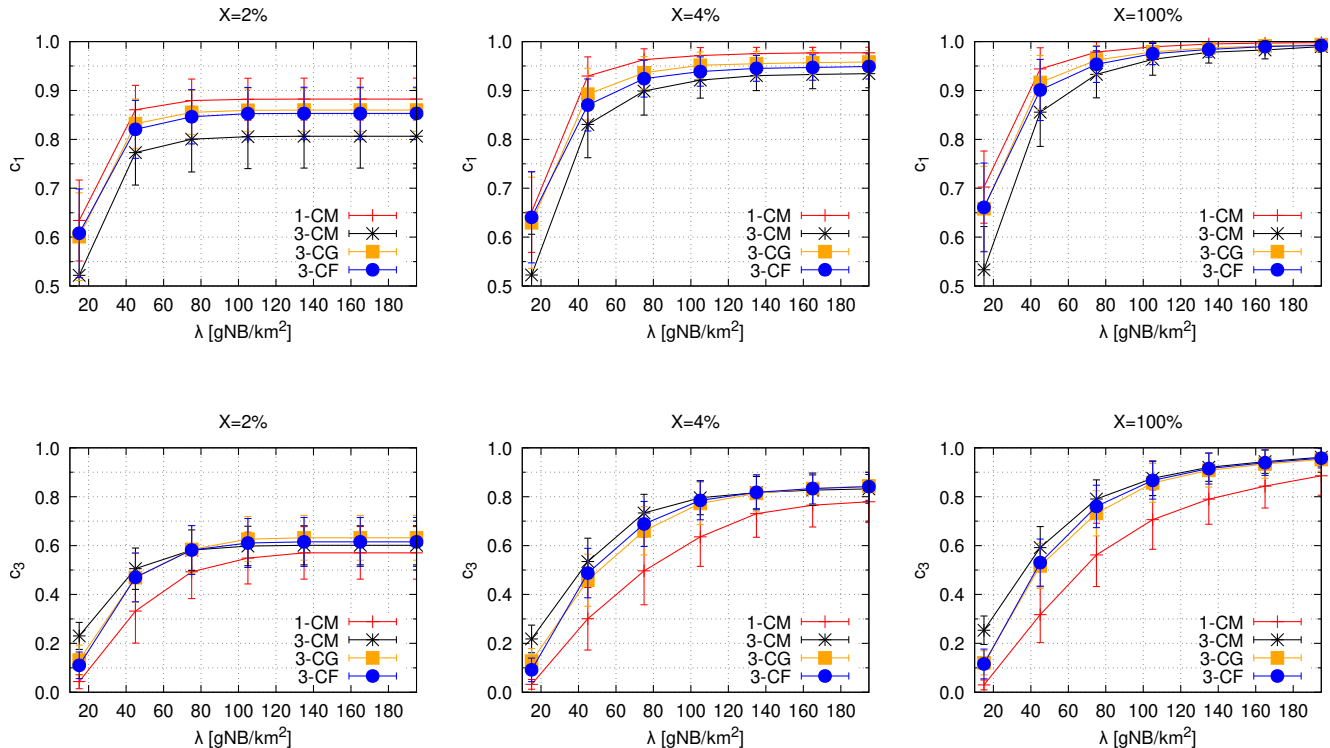


Fig. 6: Coverage metrics c_1 and c_3 for different values of X and different score functions evaluated on all the 15 areas considered.




Fig. 7: Empirical pdf (e-pdf) PDF of the coverage for $X = 4$ and different λ .

The take-away from this set of results is that if an operator wants just to provide 1-coverage, it should use 1-CM, and can reach 95% coverage with $\lambda = 75$ and $X = 4\%$. This is the most cost-effective solution, as shown later, but obviously is also the most fragile. If the operator is interested only in 3-coverage, because its target customers need extremely high reliability, then the best choice is 3-CM, calibrating X on the desired coverage. Most likely the operator is interested primarily in 1-coverage (because it enables the service to people) and with lower priority in 3-coverage (for users that need high reliability). In this case, 3-CF and 3-CG are the best choices, as they perform very close to 1-CM in terms of 1-coverage while being very close to 3-CM, outperforming 1-CM, in terms of 3-coverage.

To better understand the behavior of the coverage as a function of the gNB density, and why it tends to saturate before reaching 100% coverage, Fig. 7 shows the e-pdf of the coverage multiplicity for $X = 4\%$ and three values of λ . First, it is clear that to reach a good 1-coverage some areas are covered many times, with the tail of the coverage distribution growing significantly. Second, 3-CM penalizes the 1-coverage to the point that the probability of 0-coverage is higher than the probability of strictly 1-coverage. The mass of the distribution is different between 1-CM and the 3- score functions, but the tails remain very close to one another suggesting that there are geometric properties that force a very high degree of coverage in some areas when trying to cover the areas that have remained uncovered.




Fig. 8: O_R for all points with coverage larger or equal 2, $X = 4\%$.




Fig. 9: ECDF of the length of the link to the nearest BS. The dashed lines marks the 95th percentile.

Fig. 8 shows the robustness to obstruction O_R for all points $p \in \Lambda_2$ for different λ values. We report also the bar for 1-CM for completeness, but previous results show that 2-coverage for 1-CM is bad (for each reported λ slightly more than 50% of 2-coverage compared to the worst of the other strategies), thus O_R should be weighted by its lower 2-coverage and it is not really meaningful. The other strategies perform similarly, with a fairly high fraction of points (between 63% and 77%) that resists to shadowing from an obstacle that obstructs a 45 degree angle, and 3-CG better than the others in three cases. On the other hand, increasing the density of gNBs by a factor of 3 improves O_R only about 10%, which indicates that to improve this result a different heuristic that takes into account this specific objective should be considered.

Finally, Fig. 9 presents the E-CDF of the length of the link to the closest gNB for every covered point. We report this metric to show that with any value of λ the largest majority of the best links are below 100 meters of length, and the 95% of them are below 214 meters (less than 200 meters for higher density), so our choice of setting the maximum link length to 300 meters only marginally impacts the results.

B. Total Cost

Fig. 10 shows the cost of each deployment as a function of λ for $X = 4\%$; and $X = 100\%$. Changing X gives rise to very different trends, with the cost for $X = 4\%$ that grows





Fig. 10: Deployment cost of the network with the four score functions and for $X = 4$ and 100%.

	coverage	1-CM	3-CG
c_1	0.84	19%	24%
	0.90	27%	24%
	0.95	30%	36%
c_3	0.84	–	70%

TABLE III: Cost increase passing from $X = 4\%$ to $X = 100\%$.

remarkably more slowly than the cost for $X = 100\%$. This is the beneficial effect of limiting the number of buildings. Fig. 11 reports the trends of the cost using as a free variable the coverage, with vertical lines to highlight key coverage values (90% and 95% in c_1 and 84% in c_3). The cost for $X = 100\%$ is always higher than the cost for $X = 4\%$ at the same coverage: our heuristic with $X = 4\%$ may need a higher λ compared to $X = 100\%$ (as Fig. 6 shows), however, the total cost of for $X = 100\%$ is higher because it uses more buildings. This results highlight two facts. First, the cost to ensure nearly complete w -coverage increases very fast with a sort of asymptotic behavior independently from w . Second, and rather obvious, including all buildings in the selection allows a better coverage again independently from w .

We quantify this increase in Tab. III reporting the relative cost difference between $X = 4\%$ and $X = 100\%$ with the 1-CM and the 3-CG strategies at the same coverage value (the intercept of the vertical lines and the curves in Fig. 11). Not limiting the number of buildings produces an increased cost of 19% and 36% for c_1 and a probably unacceptable 70% for c_3 .




Fig. 11: Deployment cost of the network as a function of coverage for 1-CM and 3-CG.

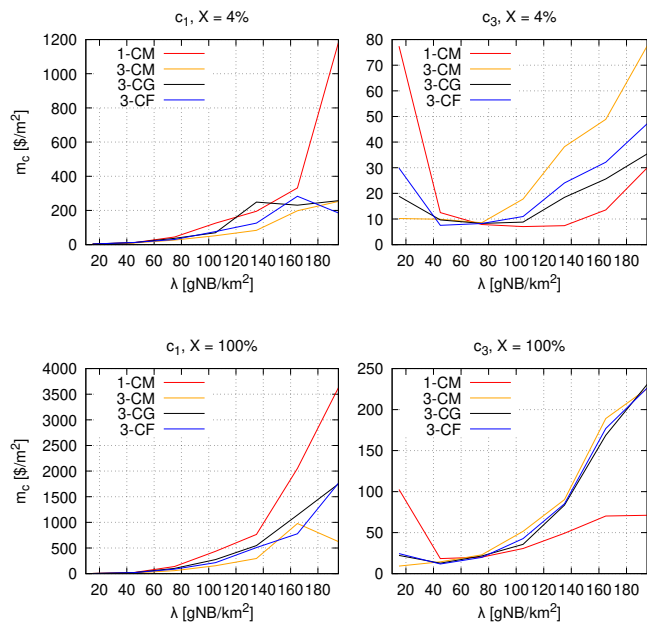


Fig. 12: Marginal cost per covered m^2 for different w -coverage and strategies.

C. Marginal Cost

Fig. 12 reports the marginal cost per m^2 $m_c(\lambda)$. Each point is the ratio between the increased cost and the increased covered area (Eq. (26)). The figure shows that c_1 (left plots) monotonically increases for the 1-CM strategy, with very high

costs when the gNB density becomes very high ($\lambda = 160, 180$). This is consistent with the design of the metric as it is trying to cover every single uncovered point, at the cost of placing new gNBs that cover just a few squared meters. The other metrics instead grow up to a certain coverage and then oscillate with a more noisy trend. This is an artifact of the metric, as the incremental cost is monotonously growing as in Fig. 10, but the incremental covered area is not monotonously growing for strategies that use $w = 3$ that do not simply try to maximize c_1 .

Looking at the right column, we see that the marginal cost per m^2 of c_3 has a different trend. Initially, the cost for all strategies decreases, as the 3-coverage is very low (see Fig. 6) and it is easy to improve it. When the coverage curves start to saturate, then m_c increases. The 1-CM strategy has a lower m_c simply because it has a lower c_3 and thus, it is not saturating even at high values of λ .

Our approach makes it possible to obtain these data that are vital for operators, because they help to identify a realistic maximum value for coverage, beyond which it is not cost-effective to add more gNBs.

IX. CONCLUSIONS

The increased demand for bit-rate and the growth of the number of devices pushed the designers of 5G and 6G to use high frequencies (mmWave or THz), a direction that is surely nonreversible. At high frequencies, the communication takes place primarily in LoS, and we know that this implies a densification of base stations. Yet, we do not know how to plan a robust access network that uses LoS communications, we do not have evidence to state if this approach is actually feasible at large scale, and we do not know if it is cost-effective.

This paper provides some initial answers to these open issues using a data-based approach to plan a LoS network in urban areas. We exploited the availability of open data and fast GPU-based computation and we implemented several algorithms to study the realization of robust coverage of open areas, showing that the problem can be solved with a reasonable density of gNB per km^2 . We have shown that the problem of maximizing 1-coverage or 3-coverage are competing ones, and there are no algorithms in the literature that we can use to solve them at the needed scale.

The new heuristic proposed prove to be the best trade-off between 1-coverage (that is the minimal requirement for an operator) and 3-coverage, and finally, we have evaluated the cost of our solution using state-of-the-art estimations.

We believe the approach we propose can open the way to new research, improving our own, and introducing several other challenges. As an example, we mention the use of Integrated Access and Backhaul (IAB) to further reduce the cost of each gNB, or the design of efficient solutions for rural areas, where the cost per user is way larger than in urban areas. Data-driven design can tackle those challenges and, to support the research community in making these steps, we publish all our code and data, together with the paper.

REFERENCES

- [1] Small Cell Forum, “Hyperdense HetNets: Definition, drivers and barriers,” Tech. Rep., Feb 2017. [Online]. Available: https://scf.io/en/documents/180_Hyperdense_HetNets_Definition_drivers_and_barriers.php
- [2] B. B. Haile, E. Mutafungwa, and J. Hämäläinen, “A Data-Driven Multiobjective Optimization Framework for Hyperdense 5G Network Planning,” *IEEE Access*, vol. 8, pp. 169 423–169 443, 2020.
- [3] G. Gemmi, R. Lo Cigno, and L. Maccari, “Dataset: On Cost-effective, Reliable Coverage for LoS Communications in Urban Areas,” Jul. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6813283>
- [4] K. Haneda and et al., “5G 3GPP-Like Channel Models for Outdoor Urban Microcellular and Macrocellular Environments,” in *2016 IEEE 83rd Vehicular Technology Conference (VTC Spring)*, May 2016, pp. 1–7.
- [5] M. M. Ahamed and S. Faruque, “5G Network Coverage Planning and Analysis of the Deployment Challenges,” *MDPI Sensors*, vol. 21, no. 19, pp. 6698(1–26), Oct. 2021.
- [6] M. Wright, “Optimization methods for base station placement in wireless applications,” in *IEEE Vehicular Technology Conference*, vol. 1, 1998, pp. 387–391.
- [7] C. K. Anjinappa, F. Erden, and I. Güvenç, “Base Station and Passive Reflectors Placement for Urban mmWave Networks,” *IEEE Transactions on Vehicular Technology*, vol. 70, no. 4, pp. 3525–3539, 2021.
- [8] S. S. Szyszkowicz, A. Lou, and H. Yanikomeroglu, “Automated Placement of Individual Millimeter-Wave Wall-Mounted Base Stations for Line-of-Sight Coverage of Outdoor Urban Areas,” *IEEE Wireless Communications Letters*, vol. 5, no. 3, pp. 316–319, 2016.
- [9] N. Palizban, S. Szyszkowicz, and H. Yanikomeroglu, “Automation of millimeter wave network planning for outdoor coverage in dense urban areas using wall-mounted base stations,” *IEEE Wireless Communications Letters*, vol. 6, no. 2, pp. 206–209, 2017.
- [10] Y. Zhang, L. Dai, and E. W. M. Wong, “Optimal BS Deployment and User Association for 5G Millimeter Wave Communication Networks,” *IEEE Trans. on Wireless Communications*, vol. 20, no. 5, pp. 2776–2791, 2021.
- [11] N. Adam, C. Tapparello, W. Heinzelman, and H. Yanikomeroglu, “Placement Optimization of Multiple UAV Base Stations,” in *2021 IEEE Wireless Communications and Networking Conference (WCNC)*, 2021, pp. 1–7.
- [12] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, “An analysis of approximations for maximizing submodular set functions—i,” *Mathematical programming*, vol. 14, no. 1, pp. 265–294, 1978.
- [13] P. Slavík, “A tight analysis of the greedy algorithm for set cover,” in *Symposium on Theory of Computing (STOC)*. Association for Computing Machinery (ACM), 1996, pp. 435–441.
- [14] Y. Ran, Y. Shi, and Z. Zhang, “Local ratio method on partial set multi-cover,” *Journal of Combinatorial Optimization*, vol. 34, no. 1, pp. 302–313, Jul. 2017.
- [15] Y. Shi, Y. Ran, Z. Zhang, J. Willson, G. Tong, and D.-Z. Du, “Approximation algorithm for the partial set multi-cover problem,” *Journal of Global Optimization*, vol. 75, no. 4, pp. 1133–1146, Dec. 2019.
- [16] Y. Ran, Y. Shi, C. Tang, and Z. Zhang, “A primal-dual algorithm for the minimum partial set multi-cover problem,” *Journal of Combinatorial Optimization*, vol. 39, no. 3, pp. 725–746, Apr. 2020.
- [17] S. Yang, F. Dai, M. Cardei, and J. Wu, “On multiple point coverage in wireless sensor networks,” in *IEEE International Conference on Mobile Adhoc and Sensor Systems Conference*, Nov. 2005.
- [18] M. Farsi, M. A. Elhosseini, M. Badawy, H. Arafat Ali, and H. Zain Eldin, “Deployment techniques in wireless sensor networks, coverage and connectivity: A survey,” *IEEE Access*, vol. 7, pp. 28 940–28 954, 2019.
- [19] V. P. Munishwar and N. B. Abu-Ghazaleh, “Coverage algorithms for visual sensor networks,” *ACM Transactions on Sensor Networks*, vol. 9, no. 4, jul 2013.
- [20] S. M. B. Malek, M. M. Sadik, and A. Rahman, “On balanced k-coverage in visual sensor networks,” *Elsevier Jou. of Network and Computer Applications*, vol. 72, pp. 72–86, 9 2016.
- [21] A. Osterman, L. Benedičič, and P. Ritoša, “An IO-efficient parallel implementation of an R2 viewshed algorithm for large terrain maps on a CUDA GPU,” *Int. Jou. of Geographical Information Science*, vol. 28, no. 11, pp. 2304–2327, 2014.
- [22] S. K. Lam, A. Pitrou, and S. Seibert, “Numba: A LLVM-Based Python JIT Compiler,” in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, 2015, pp. 1–6.
- [23] G. Gemmi, R. Lo Cigno, and L. Maccari, “Wip: Analysis of feasible topologies for backhaul mesh networks,” in *2021 IEEE 22nd International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 2021.
- [24] P. Koivumäki, A. F. Molisch, and K. Haneda, “Line-of-Sight Probability in Cluttered Urban Microcells: Analyses Using Poisson Point Process and Point Cloud,” *IEEE Transactions on Antennas and Propagation*, vol. 70, no. 3, pp. 2161–2173, 2022.
- [25] H. Heijmans and C. Ronse, “The algebraic basis of mathematical morphology I. Dilations and erosions,” *Computer Vision, Graphics, and Image Processing*, vol. 50, no. 3, pp. 245–295, 1990.
- [26] ITU, “Guidelines for evaluation of radio interface technologies for IMT-2020,” Tech. Rep. M.2412-0, Oct 2017. [Online]. Available: <https://www.itu.int/pub/R-REP-M.2412>
- [27] W. R. Franklin and C. K. Ray, “Higher isn’t necessarily better: Visibility algorithms and experiments,” in *Advances in GIS Research: Sixth International Symposium on Spatial Data Handling*. Taylor & Francis, 1994, pp. 751–770.
- [28] D. S. Hochbaum, “Approximating covering and packing problems: set cover, vertex cover, independent set, and related problems,” in *Approximation algorithms for NP-hard problems*, 1996, pp. 94–143.
- [29] C. Lund and M. Yannakakis, “On the hardness of approximating minimization problems,” *J. ACM*, vol. 41, no. 5, pp. 960–981, sep 1994.
- [30] R. Jain, D. Chiu, and W. Hawe, “A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems,” DEC Research Report TR-301, Tech. Rep., 9 1984.
- [31] E. J. Oughton and T. Russell, “The importance of spatio-temporal infrastructure assessment: Evidence for 5G from the Oxford–Cambridge Arc,” *Computers, Environment and Urban Systems*, vol. 83, Jun. 2020.



Gabriele Gemmi received an MSc in computer science from the University of Trento in 2020. He is now pursuing a double Ph.D. in computer science at the Ca’ Foscari University of Venice (IT) and in computer architecture at Polytechnic University of Catalonia (ES). His research interests are in wireless communication networks, distributed systems, and network economics.



Renato Lo Cigno is full professor at the University of Brescia, Italy. He received a degree in electronic engineering with a specialization in Telecommunications from Politecnico di Torino; in 1998/9, he was visiting scholar at UCLA; from 2002 to 2019 he was with the University of Trento, Italy. He is currently Associate Editor for IEEE/ACM TNET and Editor for Elsevier Computer Communications. His current research interests are in performance evaluation of wired and wireless networks, modeling and simulation techniques, vehicular networks and networks beyond 5G, and networked systems in general. Renato Lo Cigno is senior member of IEEE and ACM and has co-authored around 250 papers in peer reviewed journals and conferences.



Leonardo Maccari is an associate professor at the Ca’ Foscari University of Venice, Italy. He received a Ph.D in the Faculty of Computer Science Engineering in the University of Florence in 2011 and he was a Marie Curie fellow and an assistant professor at the University of Trento (Italy) till 2019. He is an IEEE and ACM member, he co-authored more than 60 publications in refereed conferences, journals and book chapters. He is also among the authors of European and US patents. He was part of the TPC of relevant conferences such as IEEE INFOCOM, IEEE

ICDCS, IFIP Networking, IEEE Globecom, IEEE ICC. Among his research interest there are network protocols for large-scale wireless mesh networks, Community Networks, interdisciplinary studies and network security.