

User/Operator Utility-Based Infrastructure Deployment Strategies for Vehicular Networks

Pasquale Cataldi* and Jérôme Härri*[†]

Mobile Communications Department

EURECOM[‡]

06904 Sophia-Antipolis, France

Email: {cataldi, haerri}@eurecom.fr

Abstract—In this paper, we propose a framework for user/operator utility-based infrastructure deployment in vehicular networks. We extend the well known maximum coverage problem to consider the following four complementary aspects: (i) a non-circular polygon-based coverage representing measured heterogeneous any-directional communication conditions (ii) a generic intensity-based utility function reflecting variations of the intensity of a desired metric over the polygonal area (iii) comparable metrics either by normalizing the utility over the coverage, or by an intensity threshold maintaining a minimal utility regardless of coverage, (iv) a joint user/operator utility optimization minimizing the infrastructure deployment. We finally illustrate the benefits of our framework over related coverage-based approaches in an exemplary study.

I. INTRODUCTION

Heterogeneous Vehicular Networks (HVN), which Intelligent Transportation Systems (ITS) is probably the most popular application area, are intended to provide connectivity to vehicles on the road benefiting from heterogeneous access technologies (WLAN, 802.11p, LTE) or vehicle-to-vehicle (V2V) or vehicle-to-infrastructure (V2I) communication patterns over a single hop or multiple hops. Despite the dedicated V2V communication capabilities, a key component of HVN is a sufficient coverage provided by Infrastructure Nodes (INs) for V2I networking. Coverage could be provided by cellular networks, but a particularity of HVN is also to consider benefit from *customized* intermittent coverage for localized services, networking, or localized connectivity enhancements provided by INs. The common aspect is that a continuous coverage is not required but INs deployed for such customized benefit should still be optimized due to monetary or logistic reasons. In this situation, we not only need to optimize the locations where to deploy these INs as to maximize the benefit to users, but also to minimize the number of INs required to satisfy operators. The question we pose in this paper is as follows: Assuming an arbitrarily road topology considering realistic radio conditions, mobility and INs of heterogeneous capabilities, can we find a joint user/operator

utility maximizing the vehicular *customized* connectivity and minimizing the number of INs?

Beside the complexity of optimizing the number and locations of INs, we have to address the following key aspects currently not sufficiently considered in available solutions. First, the coverage area of INs cannot be modeled as a circular shape characterized by a fixed theoretical transmission range, but is of variable shape in space depending on the intensity of the vehicular communication conditions. Second, coverage does not reflect the quality of the experienced connectivity, as the intensity of the connectivity may significantly fluctuate within the coverage area (i.e. being connected does not mean being capable of communicating). Finally, comparing the benefits from INs of variable capabilities or technologies requires comparable metrics to avoid bias.

In order to provide an answer to the given question, we present in this paper an optimization framework defining different deployment strategies based of specific user and operator requirements. We define a general *Maximum Benefit Problem (MBP)* and provide two specific formulations considering either a normalized benefit, B^n , or a minimal guaranteed benefit, B^{mgn} . The benefit functions take into account the following aspects: (i) we do not consider theoretical benefit functions (such as a theoretical radio range), but instead analyze benefits from calibrated simulated data. (ii) the benefit area is modeled as a generic polygon rather than a circle reflecting spatial and directional variations. (iii) we define the benefit functions as a continuous 3D function, where the third dimension represents the intensity of the benefit over the 2D area. (iv) non-compatible meanings of metrics (RSSI, range, degree, link-duration) being one matter of bias in comparing benefits from INs of heterogeneous capabilities, the framework is based on a generic benefit function with the only constraint being to be normalized over the size of the value space (e.g., 2D area in this paper).

As a case study of our framework, we addressed the placement of INs in the Bologna area using the iTETRIS platform [1], which allowed us to consider both mobility and wireless communication with realistic level of details. The outcome of our study illustrated how the proposed framework provides a higher joined user/operator utility that leads to better design of the deployment of INs.

The rest of this work is organized as follows: Section II

*P. Cataldi and J. Härri acknowledge the support of the European Commission through FP7 ICT Project iTETRIS: An Integrated Wireless and Traffic Platform for Real-Time Road Traffic Management Solutions (No. FP7 224644)

[†]J. Härri further acknowledges the support of the French DGCI/OSEO for the FOT Project SCOREF.

[‡]EURECOM industrial members: BMW Group, Cisco, Monaco Telecom, Orange, SAP, SFR, STEricsson, Swisscom, Symantec, Thales.

provides related studies in the domain of infrastructure deployment, while Section III covers the modeling of the intensity-based benefit function. In Section IV, we introduce the infrastructure deployment framework, then describe in Section V its implementation and illustrate its benefits. We summarize this work and give directions of future work in Section VI.

II. RELATED WORK

Infrastructure deployment is investigated in either cellular contexts or in WLAN or sensor network contexts. Considering cellular networks, contractual and financial reasons made operators optimize a continuous connectivity at a maximum received signal or best access with a minimum number of INs, such as in [2], [3]. In the other context, full connectivity is mostly not required but rather customized given particular connectivity or networking contexts (i.e. for localized services). On the one hand, early approaches, such as [4]–[6], formulated the objective as a Maximum Coverage Problem, where k INs should maximize the total radio coverage. On the other hand, networking or application-based benefits from INs could be optimized instead of coverage. For instance, in works described in [7]–[9], authors opted to optimize the benefit from k INs as function of dissemination capacity, data gathering or link duration. A common limitation in most of the approaches in this context is an accurate modeling of the true communication capacities of INs. Indeed, the coverage or networking range of INs is modeled as a circle with constant connectivity or networking benefits. In a recent study, Yu *et al.* [10] attempted to address such shortcomings by segmenting a coverage circle in zones of variable ranges, but modeled the benefit function with a random or stochastic value rather than a real connectivity or networking capabilities. Finally, similarly to the cellular world, deploying INs also requires considering the cost of INs. Accordingly, not only the location of INs should be optimized but also their numbers. These are the aspects we address in this paper.

III. INTENSITY-BASED INFRASTRUCTURE BENEFIT

Before formulating the infrastructure deployment problem in Section IV, we need to characterize the benefit associated to an IN. We here define the *benefit* of an IN as 3D function¹ that quantitatively describes the advantage of placing the node at a given position. We describe an accurate benefit function in space (i.e. as a function of the positions around the considered INs) in Section III-A, its variable intensity (i.e. as a function of defined metrics) in Section III-B and, finally, normalize it for comparison between INs in Section III-C.

A. Non-Convex Polygon-Based Benefit Function

The benefit area of INs should not be modeled as a circular shape considering that radio propagation is not omnidirectional due to several factors, such as fading, channel interference and obstacles. In addition, in vehicular scenarios positions of the communicating nodes are constrained by

¹In the rest of the paper we will refer to a 2D spacial area within which the benefit function is defined as the *benefit area*.

streets. Areas where nodes cannot be found should therefore not be included. Finally, assuming a circular area at a fixed benefit does not provide detailed information about communication or service quality and may lead to poor design choices. In Fig. 1 we present the benefit of an IN as a function of the position of the covered nodes. As we can observe, the circular area only provides rough information about the actual benefit. A better approach is to consider the benefit area as a polygon. The convex polygon depicted in dash-dotted line could be one option, but it still lacks the level of details needed for an optimal deployment of INs, as nodes cannot be found in a large part of its area.

In this paper we propose to describe the benefit area as a generic non-convex polygon, as shown in Fig. 1. This approach allows to depict the benefit area shape with maximum details.

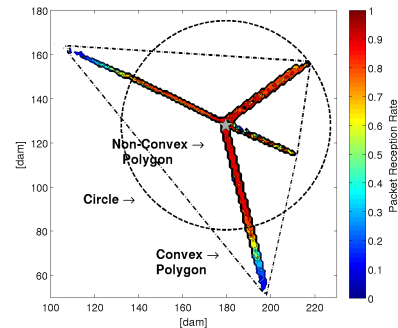


Fig. 1: Polygon coverage, where the heat map represents the intensity of a given benefit function associated to an IN located at the intersection of four roads.

B. Non-Homogeneous Benefit

Related studies mostly model a constant benefit function over a circular area although communication-based utility functions are more expected to vary in space as illustrated by the legend in Fig. 1. The traditional constant benefit significantly overestimates a more realistic benefit varying as function of distance as depicted in Fig. 2. The framework introduced in this paper considers a continuous function modeling the variation of the intensity of the benefit function over non-convex areas.




Fig. 2: Intensity of an infrastructure node in the coverage area.

C. Normalized Benefit Metric

We define two intensity-based benefit metrics: a *normalized benefit* B^n and the *minimum guaranteed normalized benefit*

B^{mgn} . The objective of the first metric is to provide a comparable representation of the IN's expected benefit over distance, while the second aims at better reflecting a comparable desired benefit required from the INs. We illustrate in Fig. 3 how these two metrics are obtained. B^n is normalized by integrating the continuous benefit function over the area (see Fig. 3(a)). B^{mgn} is instead obtained by limiting the range as function of a user requested minimum guaranteed benefit and again normalized over the area (see Fig. 3(b)).




Fig. 3: Two benefit metrics for infrastructure deployment.

IV. INFRASTRUCTURE DEPLOYMENT FRAMEWORK

Given a set of possible locations for the placement of INs, we are interested in placing at **most** k of them such that the joint user/operator benefit function is maximized.

The problem can be formulated as a Maximum Benefit Problem (MBP), which can be defined for a generic benefit function w as follows.

Theorem 1 (Maximum Benefit Problem): Given a collection of N sets $S = \{s_1, s_2, \dots, s_N\}$, where each set is a subset of a given set of ground elements $E = \{e_1, e_2, \dots, e_M\}$, each of those associated with a benefit value $w(e_j)$, select the k sets that maximize the benefit of the union set U .

If the benefit functions are unitary steps, the problem is reduced to the basic Maximum Coverage Problem (MCP), where the objective is to find the k sets that maximize the cardinality of the union set U .

As the proposed problem is a generalization of the MCP, it is also classified as NP-hard, as an exhaustive search of the locations for the k INs requires investigating all the $\binom{N}{k}$ combinations.

A. Greedy Formulation

We describe here an extension of the greedy algorithm in [7] for a generic benefit function that provides a very good approximation of the optimal solution of Theorem 1 in $O(N^2)$.

Assuming that at most k sets can be selected, the greedy algorithm iterates on the search set X of sets that have not been selected yet to find the one that maximizes the total current benefit of the selected ones. The set with maximum benefit is added to the union set U and removed from the search set X . The output of the algorithm is given by the union set U and its associated benefit. In particular, the order of insertion of the sets in U indicates the rank of every set. In this way, if only a number $k' < k$ has to be selected, the first k' of U will be the optimal ones.

Algorithm 1 Greedy Maximum Benefit Problem

Require: sets S_i, k

- 1: $U = \emptyset$
- 2: $B_{tot} = 0$
- 3: Add all S_i to the search-set X .
- 4: **for** $step = 1$ to k **do**
- 5: **if** $X \neq \emptyset$ **then**
- 6: initialize max benefit: $B^* = 0$
- 7: initialize best set: $S^* = \emptyset$
- 8: **for all** $S_i \subset X$ **do**
- 9: compute benefit: $B = computeBenefit(S_i \cup U)$
- 10: **if** $B > B^*$ **then**
- 11: $B^* = B$
- 12: $S^* = S_i$
- 13: **else if** $B_{tot} \geq B_{tot} + B$ **then**
- 14: $X = X - S_i$
- 15: **end if**
- 16: **end for**
- 17: $U = U + S^*$
- 18: $X = X - S^*$
- 19: total benefit $B_{tot} = B_{tot} + B^*$
- 20: **else**
- 21: **return** U, B_{tot}
- 22: **end if**
- 23: **end for**
- 24: **return** U, B_{tot}

It is also important to observe that, for a particular objective function, less than k sets can achieve the same results as N sets. Therefore, the algorithm is able to detect the sets (i.e., the candidate locations of the INs) that do not contribute to the overall benefit, thus minimizing the number of chosen sets and decreasing the number of operations to be performed.

B. Set Definition

To better understand the correspondence between the MBP and our problem, we consider that each set S is associated to the benefit function of an IN. In particular, the matrix $S_{R \times C}^2$ covers the whole map and the value $S_{i,j}$ of the generic element is defined **only** inside the benefit area, expressing the intensity of the benefit function. By considering the two metrics B^n and B^{mgn} presented in Section III-C, we can express the benefit of an IN as in Eq. (1) and Eq. (2), respectively:

$$B^{nc} = \frac{\sum_i^R \sum_j^C \begin{cases} S_{i,j} & \text{if } S_{i,j} \text{ defined;} \\ 0 & \text{otherwise.} \end{cases}}{R \cdot C} \quad (1)$$

$$B^{mgn} = \frac{\sum_i^R \sum_j^C \begin{cases} S_{i,j} & \text{if } S_{i,j} \text{ defined and } S_{i,j} \geq th; \\ 0 & \text{otherwise.} \end{cases}}{R \cdot C} \quad (2)$$

where th represents the minimum benefit threshold.

The union set U obtained by placing k INs represents the cumulative benefit area of the infrastructure in the map. In particular, the element $U_{i,j}$ of the union matrix has an intensity value defined as

$$U_{i,j} = \max(S_{i,j}^1, S_{i,j}^2, \dots, S_{i,j}^k) \quad (3)$$

Equation (3) indicates that the benefit intensity at a position (i, j) of the map is equal to the maximum benefit in that position achieved by the k considered INs. This means that given a certain position, its associated benefit intensity depends only on the IN whose benefit intensity at that point is the highest, i.e. all other INs are not useful to *cover* that position.

² R and C are respectively the number of rows and columns of the matrix.

C. Utility-Based Optimal RSU Deployment

The problem formulated above and the definition of the sets allows a ranking of the candidate positions of the INs as a function of the benefit function. Given these definitions, it is now possible to define the user utility \widehat{U}_u as the measure of the advantage for the user of having a certain number of INs optimally placed. In this paper we modeled the user utility as proportional to the average message reception rate that the user can experience in any point of the union benefit area, as described by Eq. (4):

$$\widehat{U}_u = \frac{\sum_i^R \sum_j^C \begin{cases} S_{i,j} & \text{if } S_{i,j} \text{ defined;} \\ 0 & \text{otherwise.} \end{cases}}{\sum_i^R \sum_j^C \begin{cases} 1 & \text{if } S_{i,j} \text{ defined;} \\ 0 & \text{otherwise.} \end{cases}} \quad (4)$$

As we can observe by the formulation of \widehat{U}_u , the user utility depends on the definition of the benefit area. As a consequence, inaccurate descriptions will result in low \widehat{U}_u , as the over-estimated areas is filled by '0's.

D. Joined User/Operator Utility

The function \widehat{U}_u is monotonically increasing with the number of infrastructure nodes. This means that the user's advantage will not decrease when increasing the number of INs. Unfortunately, realistic infrastructure deployments must consider other parameters than just user utility that represent deployment cost, available budget, deployment time, company strategy, target service, etc. Many utility functions can describe these parameters and an accurate analysis is out of the scope of the paper (and would widely depend on a business model). However, in general we can assume that, if represented as a function of the number of infrastructure nodes to be placed, the operator utility function \widehat{U}_o is monotonically decreasing.

The optimal selection of the locations for the selection of the INs to be deployed can now be obtained by multiplying \widehat{U}_u by \widehat{U}_o , as in Eq. (5), and by choosing the number n of INs that corresponds to the maximum of the resulting joined user/operator function \widehat{U}_j . The locations of the n INs will be provided by the ranking list obtained as output of the MBP.

$$\widehat{U}_j = \widehat{U}_u \cdot \widehat{U}_o \quad (5)$$

V. PERFORMANCE EVALUATION

In this section we present the simulation results of the proposed framework and compared it over the benchmark MCP. In order to limit the optimization process time, we selected 55 candidate positions for the placement of the INs, as shown in Fig. 4. The choice of placing INs at intersections is motivated by the fact that the connectivity potential is sensibly higher than when the INs are placed along road segments, as presented by Trullós *et al.* [7]. However, the proposed framework is not limited to this assumption, nor to the number of potential locations.




Fig. 4: Candidate positions of the INs.

A. Simulation Environment

In the proposed framework, we used matrices covering a $2200 \times 2200 m^2$ area, with each element covering $10 \times 10 m^2$. Without loss of generality about the use of the framework, we present in this paper the optimal placement of 802.11p infrastructure nodes (RSUs) of the same type and characteristics.

The intensity values of the benefit functions are obtained through the reception rate of beacon messages. In particular, we defined the two benefit functions B^n and B^{mgn} 0.9 as the beacon reception rate at the RSU side with respect to the position, and the minimum benefit area that guarantees a reception rate higher than 90% threshold, respectively. In order to compute the reception rates, we used the iTETRIS [1] platform, which considers realistic urban fading models and vehicular mobility patterns. We developed an application that allows RSUs to collect the received beacons generated by each node in the map and accordingly measure their reception rate. We reported the main simulation settings in Table I.

TABLE I: iTETRIS Simulation parameters.

Map [in SUMO]	Urban (Costa-Pasubio in Bologna)
Simulation Time	300 s
beacon frequency	2 Hz
beacon size	132 bytes
RSU transmission power	20 dBm
RSU height	6 m
Fading model [in ns3]	WINNER II B1 [11]
MAC [in ns3]	IEEE 802.11p CCH

B. Optimization Results

In Fig. 5 we present the evolution of the benefit as a function of the number of selected locations. For the computation of the benefit, we considered the two metrics B^n and B^{mgn} 0.9 and, as benchmark, two common two dimensional approaches, i.e. maximum circular coverage *Circle* (i.e. the coverage of each node is given by the maximum range where the signal can be received) and the case where the transmission range is defined a priori and set to $400m$, *Circle 400m*. In order to compare the benefits of the different approaches, we normalized their values. In this way, a benefit $B = 1$ represents the maximum benefit obtainable by any metric.

As we can observe, the maximum benefit is reached for all the considered metrics with a number of RSUs lower than 55.

This means that above a certain number of RSUs, there is no increased benefit in placing more of them. In addition, we can observe that the circular MCP approaches converge faster to the maximum benefit with respect to the non-convex MBP ones, due to the over-estimation of the benefit area and the modeling of the benefit as step function.




Fig. 5: Benefit as a function of the number of RSUs.

Figure 6(a) depicts the user utility \widehat{U}_u as a function of the number of RSUs deployed in the considered map. As we can observe, the metric $B^{mgn} 0.9$ provides the highest level of utility. In fact, although the benefit areas considered during the MBP optimization are smaller, the quality of the communication when a user is within the area is higher. This balances the need to have more infrastructure nodes to cover the whole area. In addition, unlike in Fig. 5, we can observe that the average user utility for circular coverage areas is very low, as they provide benefit where the user does not need it.

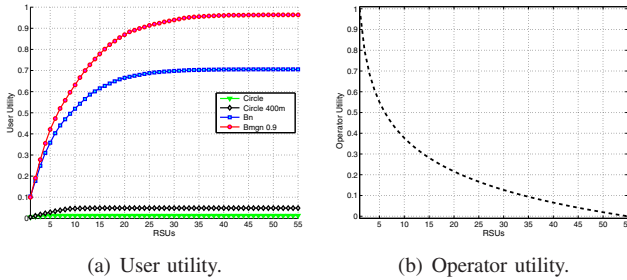


Fig. 6: Utilities as a function of the number of RSUs.

Finally, we modeled the operator utility as a Cobb-Douglas function [12], as presented in Fig. 6(b). In particular we described \widehat{U}_o as a function of the decreasing returns to scale amount of labor, expressed as number of INs deployed.

Given the user and the operator utility functions, it is now possible to compute the joined-utility, as in Eq. 5, which will give direction about the optimal deployment strategy.

In Fig. 7 we present the joined-utility results according to the different strategies that are used in the MBC optimization. As we can observe, the highest utility is reached for a number of RSUs between 7–8, significantly lower than the k candidate locations. Also, it can be observed that $B^{mgn} 0.9$ provides a higher satisfaction than all other metrics at the same cost. The

actual locations of the RSUs will now depend on the rank that the optimization process generated for each metric.

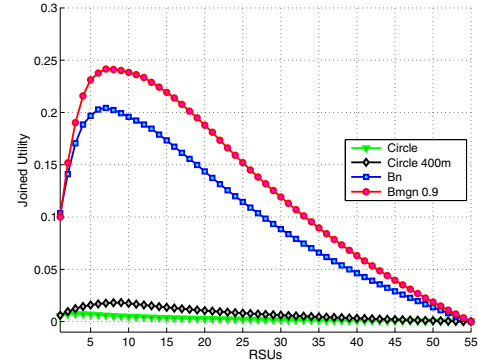


Fig. 7: Joint user/operator utility as a function of the number of RSUs.

VI. SUMMARY AND FUTURE WORK

We presented in this paper a framework for the deployment of infrastructure nodes (INs) in vehicular networks maximizing a joint user/operator utility function. We illustrated how jointly considering non-circular shape areas, non-constant connectivity-based benefits within areas and operators' deployment costs could significantly improve the design decisions on deployment strategies by providing a higher user/operator satisfaction with less required INs. This framework is also a first step towards complementary deployment strategies of heterogeneous infrastructure nodes for future HVNs. The generic formulation of the intensity benefit functions notably allows comparisons between INs of heterogeneous capabilities. In this work, we however only considered the number of INs in the operator utility, not their type or location. In future work, we plan to also consider the impact of these aspects on the cost, as it expected to significantly impact the design strategies.

REFERENCES

- [1] iTETRIS - An Integrated Wireless and Traffic Platform for Real-Time Road Traffic Management Solutions, <http://www.ict-itetris.eu/>.
- [2] R.M. Whitaker, L. Raisanen and S. Hurley *et al.*, "The infrastructure efficiency of cellular wireless networks", in *Computer Networks* Volume 48, Issue 6, pp. 941–959. 2005.
- [3] M.H. Wright, "Optimization methods for base station placement in wireless applications", in *Proc. of the IEEE Vehicular Technology Conference (VTC)*, Ottawa, Canada, 1998.
- [4] K. Tutschku, "Demand-based radio network planning of cellular mobile communication systems", in *Proc of the IEEE INFOCOM*, San Francisco, USA, 1998.
- [5] E. Amaldi, A. Capone, M. Cesana and F. Malucelli, "Optimizing WLAN radio coverage", in *Proc. of the IEEE International Conference on Communications*, Madison, USA, 2004.
- [6] X. Bai, D. Xuan, Z. Yun, T. H. Lai and W. Jia, "Complete Optimal Deployment Patterns for Full-Coverage and k -Connectivity ($k < 6$) Wireless Sensor Networks", in *Proc. of ACM MobiHoc*, 2008.
- [7] O. Trullols, M. Fiore, C. Casetti, C.F. Chiasserini and J.M. Barcelo Ordinas, "Planning Roadside Infrastructure for Information Dissemination in Intelligent Transportation Systems", in *Computer Communications (Elsevier)*, Volume 33, Issue 4, pp 432–442, 2010.
- [8] C. Lochert, B. Scheuermann, C. Wewetzer, A. Luebke and M. Mauve, "Data Aggregation and Roadside Unit Placement for a Vanet Traffic Information System", in *Proc. of the ACM MobiCom*, S. Francisco, USA, 2008.
- [9] M. Fiore and J.M. Barcelo-Ordinas, "Cooperative download in urban vehicular networks", in *Proc. of the IEEE MASS'09*, Macau, China, October 2009.
- [10] Z. Yu, J. Teng, X. Bai, D. Xuan and W. Jia, "Connected Coverage in Wireless Networks with Directional Antennas", in *Proc. of the IEEE INFOCOM*, Shanghai, China, 2011.
- [11] IST-4-027756 WINNER II D1.1.2 V1.2, "WINNER II Channel Models", 2007.
- [12] C.W. Cobb and P.H. Douglas, "A theory of production", in *The American Economic Review (JSTOR)*, Volume 18, Issue 1, pp 139–165, 1928.