POLITECNICO DI TORINO
Repository ISTITUZIONALE

Persistent Localized Broadcasting in VANETs

Original
Persistent Localized Broadcasting in VANETs / Fiore, Marco; Casetti, CLAUDIO ETTORE; Chiasserini, Carla Fabiana; Borsetti, DIEGO FERRUCCIO. - In: IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS. - ISSN 0733-8716. - ELETTRONICO. - 31:9(2013), pp. 480-490. [10.1109/JSAC.2013.SUP.0513043]

Availability:
This version is available at: 11583/2502544 since:

Publisher:
IEEE

Published
DOI:10.1109/JSAC.2013.SUP.0513043

Terms of use:

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

Publisher copyright

(Article begins on next page)

12 October 2024

Persistent Localized Broadcasting in VANETs

Marco Fiore, Claudio Casetti, Carla-Fabiana Chiasserini, Diego Borsetti

Abstract

We present a communication protocol, called LINGER, for persistent dissemination of delay-tolerant information to vehicular users, within a geographical area of interest. The goal of LINGER is to dispatch and confine information in localized areas of a mobile network with minimal protocol overhead and without requiring knowledge of the vehicles' routes or destinations. LINGER does not require roadside infrastructure support: it selects mobile nodes in a distributed, cooperative way and lets them act as "information bearers", providing uninterrupted information availability within a desired region. We analyze the performance of our dissemination mechanism through extensive simulations, in complex vehicular scenarios with realistic node mobility. The results demonstrate that LINGER represents a viable, appealing alternative to infrastructure-based solutions, as it can successfully drive the information toward a region of interest from a far away source and keep it local with negligible overhead. We show the effectiveness of such an approach in the support of localized broadcasting, in terms of both percentage of informed vehicles and information delivery delay, and we compare its performance to that of a dedicated, state-of-the-art protocol.

Index Terms: Vehicular networks, information dissemination, localized services.

I. Introduction

In the near future, it will not be uncommon for vehicles to be equipped with always-on multi-interface network devices. Therefore, it is conceivable that the node density of vehicular ad hoc networks (VANETs) will be sufficiently high to effectively complement a network infrastructure. In such a scenario, a content provider may want to spread specific delay-tolerant information to localized areas, in the spirit of [1]. Examples abound: drivers may report congested traffic conditions to nearby areas; a parking lot may advertise information on empty parking spots to nearby passing cars; a restaurant may want to advertise its lunch menu to neighboring areas where potential customers are.

This paper addresses the scenario above by investigating a distributed protocol, named LINGER (Localized INformation with GEographical Restrictions), for persistent localized broadcasting. Given a geographical area where a piece of information needs to be periodically distributed, LINGER drives the information from the data source to the area of interest, and, once there, it aims at *keeping the information content within the area* so that passing-by

users can receive it. The main objective of LINGER is therefore different from the one of geographical routing schemes, which aim at delivering an information content at a given location but are not concerned about information persistency within a geographical area.

To achieve the above goals, we let vehicles take turns in shouldering the role of information bearer, broadcasting the content while travelling across the area of interest. The suitability for the role is quantified through an index, which captures the essential mobility characteristics of the vehicle (i.e., current position, movement direction and speed). As a result, unlike previous solutions [2], LINGER does not require knowledge of the network environment or of privacy-sensitive information on mobile users, such as vehicles routes or final destinations, nor the prediction of mobility patterns. Furthermore, it is implemented through a very lightweight communication protocol. All these facts make LINGER a practical solution, which can be implemented in real-world vehicular communication systems.

Arguably, the goals listed above could be achieved by using existing pervasive infrastructure nodes, like cellular telephony base stations and WiFi access points. However, we stress that

(i) operators may rightfully charge for the used bandwidth and information may only be available to customers of a single operator, while our framework relies on an ad hoc communication paradigm;

(ii) a fixed infrastructure constrains a service to the areas where base stations or access points are deployed, while our framework, by using mobile nodes as service providers, allows flexibility and easy reconfigurability.

In the rest of the paper, we first detail the concept of persistent localized broadcasting and discuss related work (Sec. II). Then, in Sec. III we describe the LINGER protocol, while Secs. V and VI present the analysis of the protocol performance and compare it to that of a dedicated state-of-the-art technique. The analysis is carried out through realistic simulations in the scenarios described in Sec. IV. In Sec. VII, we extend the basic LINGER protocol, by introducing the use of multiple copies and defining the local management of these copies at the mobile nodes. Finally, we draw our conclusions in Sec. VIII.

II. Preliminaries and Related Work

The basic operations of a protocol for persistent data dissemination are to drive a copy of a given information content to a specific area (routing phase), and to confine the information

there (persistence phase). In the following we refer to such an area as the *region of interest (RoI)*, and, without loss of generality, we assume it to be a circular region of radius R. Also, we define as *target* the center of the RoI and consider that, when it generates an information content, the source tags it with the coordinates of the target and with the radius R of the corresponding RoI. Note that, if the persistent broadcasting aims at covering wide geographical areas, multiple RoIs, whose union matches the larger area, can be used. If instead the information has to be kept very close to the target, the RoI radius can be set to an arbitrarily small value so as to meet such a requirement.

Several protocols have been proposed that aim at achieving the first of the above goals (see the literature on geographical routing, e.g., [3]–[8], [10]). Fewer works [2], [9], [12], instead, have addressed the second goal, i.e., information persistency within a target area, which is the main objective of LINGER. Below, we review the proposals that are most relevant to our work.

MDDV. Wu *et al*. [7] proposed the Mobility-centric Data Dissemination for Vehicular networks (MDDV) protocol, which allows a source to dispatch some information to a RoI, by leveraging opportunistic vehicle-to-vehicle communication to build a path between the source position and the distant area of broadcasting. More precisely, MDDV exploits knowledge of the road topology, including the per-road lane capacity, to determine the best trajectory to reach the RoI over the street layout. Then, vehicles traveling along the chosen route forward the content to each other so as to reduce its geographical distance from the RoI. However, in MDDV the number of content bearers cannot be controlled since each vehicle autonomously determines its role; this makes the scheme demanding in terms of network resources. Furthermore, MDDV focuses on reaching the RoI only, and does not consider the persistence of information there.

MoVe. Le Brun *et al*. proposed to use vehicles as data mules to transfer data from isolated sensors (i.e., sources) to data gathering sinks (i.e., RoIs) [8]. In the protocol, named Motion Vector (MoVe), the forwarder role is decided at each contact between two vehicles that carry data for the sink. Assuming that both cars maintain their speed and direction, the data are handed over to the vehicle deemed to get close to the sink first. Also MoVe does not deal with data persistence in the RoI.

Election approach. This solution was introduced in [9], with the aim to maintain a content

within a RoI through ad hoc communications only. The mechanism does not address the forwarding phase, rather it employs a generic geocast routing protocol to drive the content from its source to the target, i.e., the RoI center. The first vehicle along the geocast route that happens to be within the RoI assumes the role of information bearer. When the vehicle leaves the RoI, the geocast routing is used again to let the content return towards the target. This basic solution has the drawback that the content is allowed to leave the RoI before it is handed over to a new bearer. The authors then propose to use the scheme in [11], which, however, uses a complex, distributed algorithm to identify the closest node to the target.

DP-IB. In the Data Pouring with Intersection Buffering (DP-IB) scheme [10], the data generated by a gateway node are broadcast along busy roads using vehicles as forwarders: the current forwarder hands its role over to the vehicle among its neighbors that is farther away in the data propagation direction. As the information reaches an intersection, it is buffered by infrastructure nodes, which are then in charge of periodically broadcasting the content to passing-by vehicles. Unlike our scheme, DP-IB requires a fixed network infrastructure to guarantee the information persistence in the RoI.

GVI. A pure VANET is instead assumed in [12], where the Geo-localized Virtual Infrastructure (GVI) protocol is proposed. GVI aims at replacing infrastructure for periodic data broadcast to vehicular users. It assumes that the information is already within the RoI, and that the target corresponds to a road intersection with R equal to the vehicle transmission range. The protocol tries to assign the broadcaster role to the vehicle that is predicted to be the closest to the target at the time the next periodic transmission is scheduled. We point out that GVI must rely on external solutions (e.g., geographic routing) to drive the content to the RoI, and a high traffic density is required for the broadcast not to die out. Finally, the choice of switching the broadcaster role at each transmission yields a significant overhead.

HomeZome. An integrated framework, addressing the routing and persistence phases, has been proposed by Leontiadis *et al*. [2]. The authors consider a scenario where roadside infostations, i.e., the data sources, publish contents addressed to subscribers in specific areas called *HomeZones*, i.e., the RoIs. Once a content has reached its HomeZone, its bearer is in charge of periodically broadcasting the information. The protocol, HomeZone hereinafter, lets each content copy to be carried by one vehicle at a time, selected according to the estimated time to reach the RoI. Indeed, HomeZone assumes a periodic route announcement

by each vehicle; when a bearer receives an announcement by a vehicle deemed to have a lower time-to-RoI, it hands over the content copy. A vehicle time-to-RoI is predicted as the sum of (i) the time that the vehicle will take to reach the point of its future trajectory that is the closest to the RoI, and (ii) the estimated time to go from such a point to the actual RoI (computed using a Dijkstra's algorithm). Thus, unlike LINGER, the HomeZone protocol requires knowledge of the future route at each vehicle, preventing its applicability to scenarios where the future route is not preemptively set or accessible.

III. The Lingering Protocol

As mentioned before, the LINGER protocol covers both the routing and the persistence phase. A single metric referred to as *lingering index* is used to select the most suitable information bearer in either phases. We design LINGER with simplicity in mind; in particular, we do not make use of outside information that may not be readily (if at all) available to mobile nodes. Rather, in order to compute its index, a vehicle requires knowledge only of its own movement to date. This allows us to define a lightweight communication protocol for bearer selection and information transfer between old/new bearers. The lingering index computation algorithm and the communication protocol are detailed below.

A. The Lingering index

The lingering index is computed locally at each node, and is designed so as to require minimum knowledge on the vehicle mobility (i.e., positioning information obtained through GPS) and to generate very low overhead. Furthermore, to effectively prompt the bearer selection, we design the lingering index according to the following general rules: *(i)* the index is positive when a node is well within the RoI or it is approaching it from outside; *(ii)* the index is negative when a node is at the area border or outside the RoI. Thus, a negative index is associated with the need to look for a better bearer.

In order to realize the above requirements, we define the index as

$$i(d, \theta, v) = K\left[f(d) + g(\theta, d) + h(v, \theta, d)\right] \tag{1}$$

where K is a normalization factor computed so that the maximum value taken by index is equal to one, d is the node distance from the target, θ is the node heading (i.e., the angular

distance of the node speed vector with respect to the straight line connecting the node to the target), and v is the node speed module.

The index is built according to *the minimal set of rules* that are needed in order to support localized broadcasting, and the functions $f(d)$, $g(\theta, d)$ and $h(v, \theta, d)$ are the simplest mathematical formulations reflecting such rules. The rules, along with the functions expressions, are given below.

Distance contribution, $f(d)$. The relationship between d and the node's suitability as an information bearer is fashioned after the following rules:

$f.1$ a node in the RoI ($d \leq R$) is considered a better bearer than a node out of the RoI ($d > R$);

$f.2$ the quality as bearer of a node outside the RoI decreases as its distance from the target grows;

$f.3$ to minimize the number of information handovers, a node at distance d, which is neither too close to 0 nor too close to R, should be favored. Indeed, a vehicle moving toward, and already nearing, the target would spend a shorter time within the RoI than another one moving along the same direction but still far away from the target. However, a vehicle that is too close to the border of the RoI is not an ideal choice either: even if the vehicle were currently speeding toward the target, it might abruptly steer in a different direction and move out of the RoI very quickly.

Accordingly, by using simple linear[1] functions, we define $f(d)$ as

$$f(d) = \begin{cases} 2\frac{d}{R}, & \text{if } 0 \leq d \leq \frac{R}{2} \\ 2\left(1 - \frac{d}{R}\right), & \text{if } d > \frac{R}{2} \end{cases} \tag{2}$$

Note that $f(d)$ returns values in $(-\infty, 0)$ for nodes outside the RoI, and in $[0, 1]$ for nodes inside the RoI, which realizes rule $f.1$. Moreover, the negative derivative of the second linear function in (2) guarantees that rule $f.2$ is satisfied, and, jointly with the positive derivative of the first linear function in (2), determines a maximum in $R/2$, which fulfills rule $f.3$.

Heading contribution, $g(\theta, d)$. The heading θ is 0 when the node moves exactly toward the target, while it reaches its maximum value, equal to π, when the node travels exactly

[1]Other functions (e.g., parabolic and trigonometric) gave similar results: thus, we opted for the simplest solution, i.e., linear functions.

away from the target. Nodes travelling toward or away from the RoI must carry different lingering indices, even when at similar distance from the target. In particular,

$g.1$ nodes travelling toward the target are preferable over those moving away from the target;

$g.2$ the impact of the heading becomes of lesser importance when nodes are close to the target itself.

Such considerations lead to the following formulation of the functional $g(\theta, d) = \alpha(d)\nu(\theta)$, where

$$\alpha(d) = \begin{cases} 2\frac{d}{R}, & \text{if } 0 \leq d \leq \frac{R}{2} \\ 1, & \text{if } d > \frac{R}{2} \end{cases} \quad ; \quad \nu(\theta) = -\frac{\theta}{\pi} \tag{3}$$

Function $\nu(\theta)$ is linear with θ with negative derivative, thus fulfilling rule $g.1$. The effect of distance imposed by criterion $g.2$ is modeled by $\alpha(d)$, which multiplies $\nu(\theta)$ in $g(\theta, d)$. From (3), for $d > \frac{R}{2}$, $\alpha(d)$ has value 1, and thus no impact on $g(\theta, d)$. However, for $d \leq \frac{R}{2}$, it is directly proportional to d, tending to 0 for decreasing distances: thus, the closer the node is to the target, the more $\alpha(d)$ mitigates the differences that diverse θ's induce on $g(\theta, d)$.

Speed contribution, $h(v, \theta, d)$. The rationale behind the contribution of speed is as follows:

$h.1$ if a node is approaching the target, its lingering index should be: (i) *increasing with its speed* if the node is outside the RoI, as a fast moving node is expected to reach the area before a slow one; (ii) *independent of speed* if the node lies close to the RoI border, since, in this case, both high and low speeds can lead to potential advantages: a fast node would reach the target first, but a slow one would spend more time within the RoI; (iii) *decreasing as the speed grows* if the node is within the RoI, since at lower speeds the node will linger within a distance R from the target for a longer time;

$h.2$ if a node is moving away from the target, its index should be higher at low speeds than it is at high ones, independently of the distance d from target; indeed, lower speeds allow the information content to linger for a longer time close to the target.

Bearing these remarks in mind, we define the speed functional as

$$h(v, \theta, d) = A(\theta, d)\frac{v}{v_{max}} + B(\theta, d) \tag{4}$$

$$A(\theta, d) = \frac{\min\{d, 2R\}}{R}[\nu(\theta) + 1] - 1 \quad ; \quad B(\theta, d) = -\frac{\min\{d, 2R\}}{2R}[\nu(\theta) + 1] + 1 \tag{5}$$

where v_{max} is the maximum speed that a node can reach, which it is fair to assume nodes do know.

Let us first consider the case of a node moving toward the target ($\theta = 0$): when the distance d from the target is small, $A(\theta, d)$ is negative and $h(v, \theta, d)$ decreases linearly as v increases. However, as d grows, $A(\theta, d)$ tends to zero making $h(v, \theta, d)$ less and less sensitive to speed, until independent of v for $d = R$. For distances d larger than R, $h(v, \theta, d)$ becomes directly proportional to v, up to a value $d = 2R$: for even greater distances, the value of $h(v, \theta, d)$ does not vary. Summarizing, the impact of speed on the index of a node approaching the target is exactly that described in criterion $h.1$.

When instead a node leaves the target ($\theta = \pi$), the inverse proportionality between $h(v, \theta, d)$ and v is unaffected by the distance, as required by $h.2$. Indeed, as θ tends to π, the presence of the factor $\nu(\theta)$ in the expressions of A and B makes them more and more independent of distance d.

The overall lingering index is plotted in Fig. 1 as a function of the distance from the target, for the four boundary combinations of θ and v values (plots presenting the different contributions to the index can be found in [13]). Note that it takes values in the interval $(-\infty, 1]$ and, consistently with our initial requirements, it exhibits the following properties:

- *outer far area ($d > 2R$)*: the heading θ yields the predominant contribution, as a node approaching the target exhibits a higher index than one leaving it. Such a difference is further influenced by speed v, since nodes that quickly carry the information content toward the RoI ($\theta \to 0$, $v \to v_{max}$) are preferable, while nodes leaving the RoI at high speed ($\theta \to \pi$, $v \to v_{max}$) are not good choices;
- *outer near area ($R < d \leq 2R$)*: indices of nodes approaching the RoI ($\theta \to 0$) at different speeds tend to become similar, since, as noted above, different speeds hold complementary advantages. In any case, such indices are higher than those of nodes leaving the RoI ($\theta \to \pi$). Nodes leaving the RoI slowly ($v \to 0$) are obviously better bearers than those leaving it rapidly ($v \to v_{max}$);
- *inside the RoI ($d \leq R$)*: for a given distance d, the most suitable bearers are always nodes moving toward the target at slow speed (θ, $v \to 0$), as they are expected to stay longer in the RoI. More precisely, in the outer rim ($R/2 < d \leq R$), it is desirable that the information content is brought nearer to the target, while in the inner core

($d \leq R/2$), the goal is to move the content as slowly as possible. Also, as shown in Fig. 1, the index becomes negative as nodes approach the border of the RoI.

B. The LINGER message exchange

We now outline the message exchange that allows vehicles to determine the best "candidate" information bearer and to transfer the information content to it. We assume that a contention-based MAC (e.g., belonging to the 802.11 family) is used. A source, either fixed or mobile but located outside the RoI, generates a message including the information content and the following header fields: (i) a target field, specifying the geographical coordinates of the target as well as the RoI radius R, and (ii) an information field used to identify the information content. Then, the source starts a search procedure aimed at identifying a candidate node to carry the content toward the RoI. To this end, it computes its own lingering index and sends a POLL message as a one-hop broadcast message. The POLL carries the following information: (i) the target and information fields defined above, (ii) the current bearer's lingering index i_B, (iii) a sequence number, (iv) the sender's address.

A neighboring candidate node receiving the POLL computes its own lingering index i_C and compares it with the one broadcast in the POLL. If $\Delta = i_C - i_B > 0$, the candidate node replies with a BID message to the current bearer. The BID is sent as a unicast message so that, in case of failure, retransmissions can be performed at the MAC layer and higher reliability is provided. The BID message includes (i) the information field, (ii) the POLL sequence number and (iii) the BID sender's address.

Upon overhearing the transmission of the BID, candidate nodes operating in promiscuous mode will refrain from issuing their scheduled reply, thus curbing the overhead. A flood of BID messages is also avoided by letting candidate nodes reply after a lag time, which decreases with Δ and ranges between 0 and a maximum value. As a consequence, due to the contention-based MAC, the first BID received by the bearer will likely be sent by the best available candidate in the neighborhood. If no BID is received, the procedure is repeated at the next position update until a BID arrives. Upon receiving the first BID, the source replies with one or more unicast messages transferring the information content to the newly elected bearer over one hop. When the new bearer acknowledges the reception, the search procedure ends.

Finally, we point out that the current information bearer keeps track of its position at the update frequency of its positioning system. Upon every position update (e.g., every second according to standard GPS receivers), it computes its own lingering index i_B: if negative (i.e., the bearer is outside the RoI, or it is leaving it), it searches for a better candidate to carry the information content, following the procedure outlined above.

IV. Simulation-based Analysis

We study the performance of LINGER in presence of realistic vehicular environments. Before detailing our approach, we point out that the realism of mobility plays a key role in the correct evaluation of protocols for vehicular networks: as shown for instance in [14], [15], using inappropriate mobility models for the evaluation of vehicular applications can lead to results that are far from reality. Such a consideration is valid in general, but it becomes even more crucial in the case of localized broadcasting schemes, which are *built* over vehicular movements and for which node mobility is a factor of primary importance. The need for realism forces us to avoid analytic approaches that, in order to be tractable, would constrain us to extremely simplified representations of the vehicle mobility.

We thus evaluate the performance of persistent localized broadcasting via simulation, conducted through *ns-2.35* and considering the vehicular mobility in the urban area of Zurich, Switzerland. This area covers 20×20 km^2 and its road topology is as depicted in Fig. 2. Car movements in the region have been synthetically generated via a multi-agent microscopic traffic simulator (MMTS) developed at ETH Zurich [16]. The macroscopic- and microscopic-level models employed in the MMTS have been shown to yield a realistic representation of both large-scale traffic flows and small-scale car-to-car interactions. Although the mobility trace provided by ETH Zurich covers a whole day, we consider an observation window of 30 minutes, between 4.30 pm and 5 pm, encompassing part of the afternoon traffic peak. This choice is due to the computational constraint of the network simulator, as the sub-trace we extract already implies the concurrent simulation of several thousands of vehicles. In the large-scale scenario, we also consider the presence of buildings between roads, with a density that decreases as we move from the center to the more rural zones of the Zurich area. As for the communication settings, vehicles communicate using the IEEE 802.11p technology and signal propagation is represented according to the urban micro-cell radio propagation model derived during the European WINNER project [17] for the 5 GHz band. The model

accounts for the effects of path loss, shadowing and multipath fading, and it differentiates between Line of Sight (LOS) and Non-Line of Sight (NLOS) propagation conditions.

In addition, we support the large-scale Zurich scenario with a small-scale mobility setting, representative of a single intersection. The rationale is that a number of tests we perform require complete control of the vehicular mobility, including the possibility of varying the road traffic intensity or introducing pedestrian users. We employ the VanetMobiSim mobility generator to reproduce controllable traces in the small-scale scenario. Microscopic-level traffic dynamics are regulated by the IDM-LC car-following model, which accounts for car interactions, overtakings, and road signaling. More precisely, the small-scale scenario represents a perpendicular intersection between two single-lane roads, each terminated by other crossways at 500 m from the junction, which is regulated by traffic lights. Desired speeds are randomly chosen between 35 km/h and 50 km/h. In such a scenario, we have tested three different average lane densities, namely 5, 20 and 30 vehicles/km, which range from low to high density, and imply different levels of network connectivity. In the intersection scenario too, the vehicles use IEEE 802.11p, but we use a simplified propagation model. Specifically, we assume that the vehicles have a common radio range $r = 100$ m (unless otherwise specified) and that the distance between roads and the presence of buildings prevent the communication between vehicles on different roads.

As for the message exchange, unless otherwise specified, each content has a size of 1 Kbyte. This implies that, when handed over to another bearer, the content is carried by a single message (which includes an additional header size of 24 bytes). We also set the size of the POLL and BID messages, respectively, to 32 bytes and 16 bytes, consistently with the fields they contain. The RoI radius is equal to $R = 200$ m, unless differently indicated, and the update frequency of the vehicle positioning system is equal to 1 Hz. The maximum lag time is set to 200 ms.

V. Supporting Broadcasting: Driving and Maintaining the Information

We first evaluate the ability of LINGER to drive an information content to a specific RoI. Then, we show how such an infrastructure-less solution can maintain the content within the RoI that has been reached. Preliminary results on these aspects can also be found in an early version of our study in [13].

Reaching the RoI. Since here we discuss the performance in steering the information content to the desired RoI, our metric of interest is the time needed for the content to reach the RoI.

In a first experiment, we selected five intersections in the Zurich scenario, as depicted in Fig. 3, and evaluated the capability of LINGER to bring the information content generated at one of the points toward the RoIs centered at the other four. The results show that LINGER always drives successfully the content within the RoIs, with minimal overhead; the time, averaged over 10 runs, needed for an information content to travel from every source to every other RoI is reported in Table I. From the recorded values, we observe that LINGER needs a few minutes to accomplish the task. This is due partially to the large distances between the data sources and the RoIs (between 3 and 8 km), and partially to the low vehicle density in proximity of some of the RoIs (e.g., target E in Fig. 3). The latter observation is also confirmed by the fact, given a pair of locations, the time to travel between them varies significantly on the two directions, as the different lanes may exhibit very diverse vehicle densities. However, smaller latencies can be achieved by combining LINGER with a dedicated geocasting protocol for the routing phase, such as the one in [2] that exploits future route information, the scheme in [5] that leverages knowledge of road topologies, or the solution in [6] that uses knowledge of real-time vehicular traffic.

Then, in order to evaluate the impact of traffic intensity on the source-to-target latency of LINGER, we employ the small-scale intersection scenario and derived the average time, observed over 10 simulation runs, between the data source and the crossing of the border of the RoI. The obtained values are as follows: 31.82 s for 5 veh./km (sparse traffic), 4.27 s for 20 veh./km (moderate traffic), and 0.706 s for 30 veh./km (heavy traffic). Clearly, under sparse traffic, it takes a longer time to reach the RoI, since the same vehicle may have to carry the information content for a long distance, for want of a better candidate. As the lane density increases from 5 to 20 veh./km, such an occurrence becomes unlikely thanks to multiple handovers, and the RoI is reached in a much shorter time.

Confining the information within the RoI. We now study the performance of our localized broadcasting solution once the information content has reached the RoI, i.e., its ability to confine the content in the RoI. To that end, we focus on the intersection scenario, as it is representative of a typical RoI configuration that has to be faced and, unlike the

Zurich scenario, allows us to vary the road traffic intensity.

The impact of the radio range on the capability of LINGER to maintain the content within the RoI is shown in Fig. 4, in terms of cumulative density function (CDF) of the information distance from the target. Wider radio ranges (i.e., $r = 200$) are associated to bearers spending more time in the RoI. Indeed, widening the radio range increases the number of candidate nodes reached by POLL messages, giving the bearer the opportunity for a more informed selection. Similarly, a better bearer selection, hence performance, is possible with higher vehicle densities. We also point out that the small step at the right end of the CDF plot for the low density case is due to the relative position of intersections and target: very few candidate nodes can be found between 100 m and 500 m from the target, then the likelihood of finding a candidate suddenly surges upward thanks to the presence of surrounding intersections, 500 m away from the target.

As a final test, we observe the impact of the information size on the LINGER performance. In Fig. 5, we report the CDF of the distance between the information content and the target, when the content size varies from 1 to 500 Kbytes, hence its transmission requires multiple packets. It is clear that the data size has virtually no effect on the ability to keep the information content within the RoI. Indeed, we found that the limit to the content size is imposed by the duration of contact between vehicles (that bounds the maximum amount of transferable data), rather than by the localized dissemination protocol. We show the probability density function (PDF) and the complementary of the CDF of the contact duration that we observed in our simulations in Fig. 6.

VI. Performance of Localized Broadcasting

We now focus on the performance of LINGER for localized data dissemination. We consider the information content to be disseminated within the RoI to be a beacon message, which has to be periodically broadcast by the bearer with a fair delay tolerance. The beacon can be forwarded in a multihop fashion (i.e., nodes can relay it), and its range is controlled through a time to live (TTL) field, decremented at every transmission. Nodes receiving the beacon rebroadcast it provided its TTL is not 0, else they discard it (indeed, the TTL limits the number of times the beacon can be relayed). The overhead of the rebroadcast process is reduced through the Preferred Group Broadcasting (PGB) scheme [18], that exploits physical layer information to select relays that are far from the last broadcast source and whose

rebroadcast areas do not overlap. Note however that beacon relaying is unrelated to any operation of localized broadcasting, and more sophisticated techniques for range control are out of the scope of this work.

For our performance evaluation, we define an *observation area*, that is the region over which the beacon should be persistently broadcasted. Note that such an area may be larger than the RoI, in which case multiple RoIs can be defined so as to efficiently broadcast the beacon over the observation area. Therefore, the observation area is used to compute the following performance metrics related to localized broadcasting:

- *1-beacon reception probability*: the probability that a node driving through the observation area receives at least one beacon;
- *informed nodes ratio*: the ratio of the number of nodes receiving the beacon to the number of nodes that are within the observation area at each beacon transmission time;
- *first beacon reception delay*: the time elapsed between the instant a node enters the observation area and the instant at which it receives the first beacon broadcast by an information bearer.

We first compare, in the large-scale Zurich scenario, the performance of LINGER to that of HomeZone (Sec. VI-A). We then focus on the behavior of LINGER in the small-scale scenario to study localized broadcasting in presence of pedestrian mobility (Sec. VI-B).

A. Large-scale broadcasting

We consider a source node, located in downtown Zurich, that needs to disseminate a beacon message over a circular observation area with a 6 km-radius and centered at the source node itself. Such a region is marked by the wide red circle in Figs. 7(a) and 7(b). In order to reach users in the observation area, we use LINGER to distribute copies of the same information content to a set of RoIs (denoted by blue circles in the figures), configured according to the following two policies.

In the first case, our goal is to minimize the protocol overhead: we thus place the targets of copies at intersections (to promote longer caching times), set the radius R of each RoI to 500 m (to reduce the number of information handovers), and assign one copy to each RoI (to reduce the load associated to it). Also, to further reduce the number of information copies in the network, we select a subset of intersections as targets so that they are at least

1 km apart from each other, and are spread all over the desired area of dissemination: this leads to a total of 30 copies lingering on the road topology, whose RoIs are depicted as the small blue circles in Fig. 7(a). We refer to this setting as target configuration A.

In the second case, we aim at maximizing the performance in terms of 1-beacon reception probability and first beacon delay. Thus, we increase the number of information copies so that *each* intersection in the observation region is covered by one RoI (to reduce the first beacon reception delay), and we reduce the radius R of each copy to 200 m (to increase the 1-beacon reception probability). Again, only one copy is associated to each RoI. We name this deployment, depicted in Fig. 7(b), target configuration B.

In both cases, we compare the performance of a LINGER-based large-scale broadcasting to that obtained through the HomeZone scheme. While simulating HomeZone, we set the route announcement interval to 10 s, consistently with [2], and the beaconing interval to 5 s.

Let us start our analysis from the overhead generated by the two schemes under configurations A and B. In Fig. 8, we plot the traffic generated by the LINGER and HomeZone mechanisms while performing their bearer election processes: since such transmissions do not actually disseminate the information content, they can be considered as the overhead that the two schemes induce on the network. As expected, target configuration A results in a lower overhead with respect to B, under both schemes; nonetheless, it is clear that LINGER is consistently less demanding than HomeZone, and two effects contribute to such a result. Firstly, HomeZone relies on periodic route announcements from nodes, which produce a constant traffic not required by the reactive approach of LINGER. Secondly, HomeZone RoIs are not areas, but geographical points: this means that a vehicle that just passed the RoI for the content it is carrying tends to give away its bearer role as soon as possible. Clearly, such a behavior generates a higher number of new bearer election processes with respect to what observed with LINGER, whose RoI-based approach does not disregard a bearer immediately after it has passed the actual target.

Fig. 9 reports the 1-beacon reception probability for vehicles that travel within an observation area matching the circular region centered at the information source and of 6 km of radius. The probability that a node in the observation area is informed is higher than 0.9, even for the less resource-demanding setting, namely, configuration A with beacon TTL equal to 1, and it takes similar values under LINGER and HomeZone. When increasing the

beacon TTL, the percentage of informed vehicles reaches 99%. The remaining 1% includes nodes that spend in the observation area an amount of time too small to be reached by one information broadcast.

If the difference in terms of 1-beacon reception probability is small for the two target configurations A and B, the first beacon reception delay is instead strongly affected by the targets settings. From the mean values, reported in Fig. 10, it is clear that target configuration B consistently outperforms configuration A. Increasing the beacon TTL helps however in reducing this gap, while a marginal difference is observed when using LINGER or HomeZone.

In conclusion, results prove that our infrastructure-less scheme can be successfully exploited for large-scale localized broadcasting in an urban environment, providing excellent performance at minimum resource cost. When compared with a dedicated, state-of-the-art technique as [2], LINGER yields very similar performance, without the need for knowledge and advertising of future routes.

B. Presence of pedestrian users

Streets are not only traveled by vehicles, but also populated by pedestrians, especially in urban areas. Since smart phones, PDAs, and other portable devices equipped with 802.11 interfaces are becoming more and more common, one could think of including pedestrian users in the localized broadcasting scheme.

In order to study the impact that non-vehicular nodes would have on the system, we evaluate the performance of LINGER in the scenario depicted in Fig. 11(a), where, in addition to vehicles, pedestrian users walk along two 2-km roads as well as within a square centered at the roads intersection. The vehicular lane density is set to 20 veh./km (moderate traffic) and 12 to 40 pedestrian users/km^2 are walking in the area, while R is set to 100 m. The observation area has a default radius of 200 m and is centered at the intersection.

The position of all nodes elected as bearers throughout a sample simulation is depicted in Fig. 11(b), along with their nature (either vehicular or pedestrian). We note that LINGER always selects as new information bearers nodes that are located at a distance around $\frac{R}{2}$ from the target, as they are characterized by higher lingering indices. Also, there is a tendency to prefer slow pedestrian users over fast vehicular ones, but the former are not always available.

In Fig. 12, we show the informed nodes ratio obtained when considering only vehicles and when also involving various densities of pedestrian users in the bearer selection process. It is clear that the performance of the application worsens as we rely more and more on users moving at walking speed. As a matter of fact, such slow nodes cache the information content for long periods of time, significant portions of which can be spent in areas that, although within the RoI, are far from the actual target and from the majority of users. When this occurs, a low percentage of nodes in the RoI is informed by the bearer through its beacon broadcasts. Besides, high vehicular speeds imply a higher mobility of the information content within the RoI, shortening its stay in locations unfavorable to the broadcasting application.

We can therefore conclude that mobility helps information broadcasting, and the involvement of pedestrian users in the system, although bringing some advantages in terms of overhead, dampens this positive effect. These results further motivate the use of vehicles as main actors in an infrastructure-less scheme for localized broadcasting.

VII. Multiple Copies within the RoI

Based on the above results, a further improvement in the performance of the LINGER broadcasting could be attained by making use of multiple copies of the content. This is indeed could speed up the data dissemination, cover larger geographical areas, and increase the mechanism robustness to the loss of copies. Below, we first extend LINGER so as to efficiently manage multiple copies within the RoI, then we assess the performance gain that can be obtained.

A. Local management

Let us consider that an information source injects in the network different copies of a content. Different copies can be associated either to the same or to different targets and/or RoI radiuses. In order to manage the multiple copies of an information content, we let the communication protocol behave as follows.

- A node that is already bearing an information content is allowed to participate to the bearer election process for other copies of that same content, transmitting BID messages if its index is sufficiently high (i.e., if $\Delta > 0$). In this case, however, a penalty is introduced, and the node always waits for the maximum lag time before replying to POLLs. By doing so, we guarantee that the node is selected as the new bearer if there is

no other candidate, preventing the content from delaying the arrival to its RoI, or from exiting that same region.

- When a node becomes the bearer of a copy of an information content for which it is already a bearer, it merely caches the information metadata (i.e., target, information ID and sequence number).
- A node that bears multiple copies of the same information content should retain at most one of them and try to find other bearers for the remaining ones. This promotes diversity among bearers of the different copies, a desirable condition when injecting multiple copies of the same information content in the network. To this end, each node keeps metadata about all the information copies it bears in a buffer. Upon every position update, the node scans such a buffer, computing the lingering index for each copy[2]. The copy with the highest positive[3] index is then marked as *active*, while all the others are marked as *in excess*. The *active* copy is held by the node, while an election process is periodically run for each of the copies *in excess*.

B. Effect of multiple copies

We evaluate the impact of different multi-copy strategies, in the intersection scenario of Fig. 11(a), with a vehicular lane density of 20 veh./km, beaconing interval equal to 5 s, and without considering pedestrian users. In particular, the three configurations depicted in Fig. 13 are considered: a single copy lingering in a RoI of radius R, centered at the intersection; four copies all lingering in the same RoI as the one used in the first case; four copies lingering in clover RoIs, each with radius $R/2$ and centered in a way to cover together the same region considered in the other two cases.

The 1-beacon reception probability is plotted in Fig. 14, for the different configurations of Fig. 13. We note that for smaller RoI radiuses ($R = 200$) the three configurations perform similarly, for any observation area radius. However, when R grows, the configurations yield different results: multiple copies in a concentric RoI consistently inform all nodes, whereas the single copy configuration outperforms multiple copies in clover RoIs for small observation area radiuses, while being in turn outperformed for large radiuses. As a matter of fact, the

[2]Since the copies can be associated to different targets and/or RoI radiuses, they can have a different value of lingering index, even if handled by the same information bearer.

[3]If all copies have negative indices, none of them is marked as *active*.

less constrained mobility of copies in the single and concentric case leads to a uniform probability of finding each copy over the area of radius R. Thus, the performance of the localized broadcasting is similar in all observation areas. Of course, employing four copies improves the results in terms of number of informed nodes. On the contrary, the clover configuration yields a population of copies that is denser around targets, and thus farther from the intersection (see Fig. 13(c)). When considering larger RoI radiuses ($R = 600$), this behavior results in a smaller likelihood of nodes to be informed while closer to the intersection (i.e., for smaller observation area radiuses), but increases it for far away nodes (e.g., for larger observation areas).

The advantage of employing clover RoIs can instead be observed in terms of first beacon reception delay. In Fig. 15, the clover configuration informs nodes faster than the others for smaller RoI radiuses ($R = 200$), independently of the observation area radius. As R increases, the clover configuration tends to perform worse for small observation areas, but better in large ones, increasing in this latter case its advantage over the other schemes. Indeed, the results for larger observation areas are more significant, as they provide a glimpse of how quickly a node approaching the intersection from the outside is informed by the broadcasting application: we thus rule the clover scheme as a hands-down winner in terms of rapidity in informing nodes.

Such a result comes however at some cost in terms of overhead: it is clear that maintaining a clover configuration leads to a higher number of information transfers, which results in a higher overhead. As depicted in Fig. 16, the clover configuration generates an overhead that, although small in absolute terms, is from two to three times higher than that observed with the concentric configuration. Obviously, the single copy case results in an even lower overhead, since a single information lingers in the RoI.

In summary, using multiple copies can improve the performance of localized broadcast, but at a cost in terms of overhead. Also, complex RoIs configurations pay off when the information is to be broadcast over large areas, while for smaller RoIs the simple presence of multiple copies in the same area is sufficient.

VIII. Conclusions

We considered a vehicular network where an information provider wishes to dispatch specific information to localized areas. Through a cooperative approach, we designed the

LINGER protocol which selects "information bearers" and lets them act as data broadcasters within a region of interest, without requiring knowledge of privacy-sensitive information on the vehicles. Simulation results in vehicular scenarios with realistic node mobility show that LINGER is a viable, effective alternative to infrastructure-based localized broadcasting. Indeed, it manages to successfully and efficiently carry an information content toward a region of interest, confine the content there, and inform a very high percentage of passing-by vehicles, under any road traffic conditions. We also showed that a significant further increase in performance can be achieved by efficiently managing multiple copies of the content.

REFERENCES

[1] G. Karlsson, V. Lenders, M. May, "Delay-tolerant Broadcasting," *SIGCOMM, Workshop on Challenged Networks*, 2006.

[2] I. Leontiadis, P. Costa, C. Mascolo, "Persistent Content-based Information Dissemination in Hybrid Vehicular Networks," *PerCom,* 2009.

[3] F. Granelli, G. Boato, D. Kliazovich, G. Vernazza, "Enhanced GPSR Routing in Multi-hop Vehicular Communications through Movement Awareness," *IEEE Communications Letters*, vol. 11, no. 10, 2007.

[4] FP7-Project, "GeoNet - Geographic Addressing and Routing for Vehicular Communications," STREP 216269. THEME ICT-2007.6.1, ICT for Intelligent Vehicles and Mobility Services.

[5] K.C. Lee, M. Le, J. Haerri, M. Gerla, "LOUVRE: Landmark Overlays for Urban Vehicular Routing Environments," *WiVeC*, 2008.

[6] J. Nzouonta, N. Rajgure, G. Wang, C. Borcea, "VANET Routing on City Roads using Real-Time Vehicular Traffic Information," *IEEE Trans. on Vehicular Technology*, vol. 58, no. 7, 2009.

[7] H. Wu, R. Fujimoto, R. Guensler, M. Hunter, "MDDV: A Mobility-centric Data Dissemination Algorithm for Vehicular Networks," *VANET,* 2004.

[8] J. LeBrun, C.-N. Chuah, D. Ghosal, M. Zhang, "Knowledge-Based Opportunistic Forwarding in Vehicular Wireless Ad Hoc Networks," *IEEE VTC-Spring*, 2005.

[9] C. Maihöfer, T. Leinmüller, E. Schoch, "Abiding Geocast: Time-stable Geocast for Ad Hoc Networks," *VANET*, 2005.

[10] J. Zhao, Y. Zhang, G. Cao, "Data Pouring and Buffering on The Road: A New Data Dissemination Paradigm for Vehicular Ad Hoc Networks," *IEEE Trans. on Vehicular Tech.,* vol. 56, no. 6, 2007.

[11] W.-H. Liao, Y.-C. Tseng, K.-L. Lo, J.-P. Sheu, "GeoGRID: A Geocasting Protocol for Mobile Ad Hoc Networks Based on GRID," *J. of Internet Technology*, vol. 1, no. 2, 2000.

[12] M. Jerbi, A.-L. Beylot, S.-M. Senouci, Y. Ghamri, "Geo-localized Virtual Infrastructure for VANETs: Design and Analysis," *GLOBECOM*, 2008.

[13] D. Borsetti, M. Fiore, C. Casetti, C.-F. Chiasserini, "When Mobile Services Go Local," *MSWiM*, 2009.

[14] A. Kesting, M. Treiber, D. Helbing, "Connectivity Statistics of Store-and-forward Intervehicle Communication," *IEEE Trans. on Intelligent Transportation Systems*, vol. 11, no. 1 Mar. 2010, pp. 172–181.

[15] M. Fiore, J. Härri, "The Networking Shape of Vehicular Mobility", *MobiHoc*, 2008.

[16] N. Cetin, A. Burri, K. Nagel, "A Large-scale Multi-agent Traffic Microsimulation Based on Queue Model," *Swiss Transport Research Conference (STRC)*, 2003.

[17] WINNER. D1.1.1. WINNER - Interim Channel Models, http://www.ist-winner.org/.

[18] V. Naumov, R. Baumann, T. Gross, "An Evaluation of Inter-vehicle Ad Hoc networks Based on Realistic Vehicular Traces," *MobiHoc*, 2006.

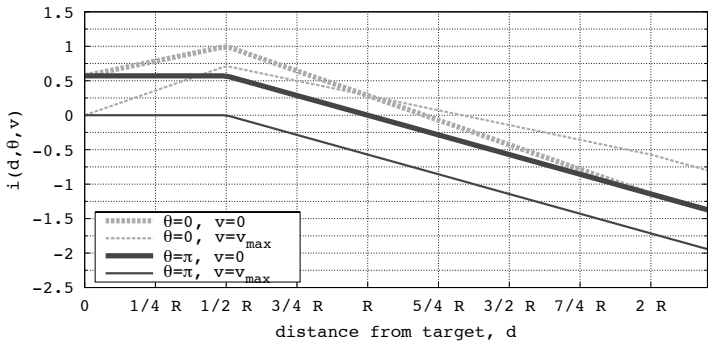

Fig. 1. Lingering index $i(d, \theta, v)$, for the extreme cases of the $\{\theta, v\}$ domain

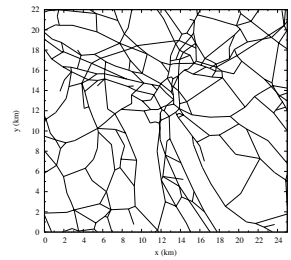

Fig. 2. The Zurich road topology

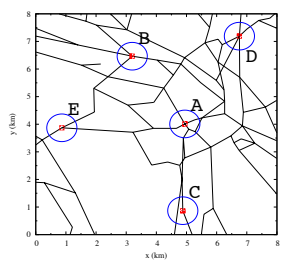

Fig. 3. Positions of information sources and targets in the Zurich scenario, indicated as red dots. RoIs are depicted as blue circles

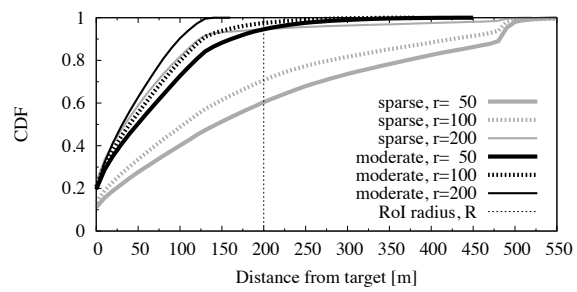

Fig. 4. CDF of the information distance from the target, for different values of node radio range and lane densities

TABLE I

TIME [MM'SS] TO REACH THE ROI (A, B, C , D, E), IN THE ZURICH SCENARIO

Source	A	B	C	D	E
A	-	1'04	0'33	0'57	3'25
B	1'19	-	2'48	2'09	1'43
C	0'30	2'53	-	1'52	2'14
D	1'54	1'04	3'57	-	5'21
E	2'23	3'01	3'22	2'56	-

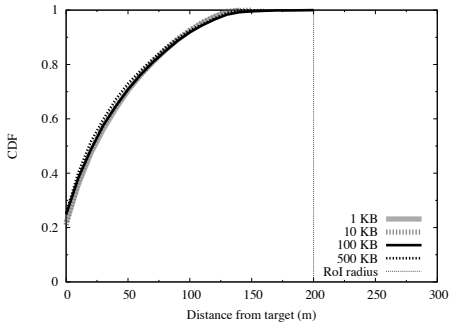

Fig. 5. CDF of the information distance from the target

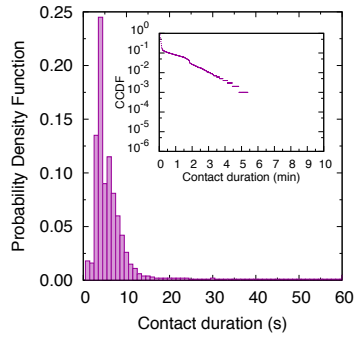

Fig. 6. PDF and complementary CDF of the contact duration between two vehicles

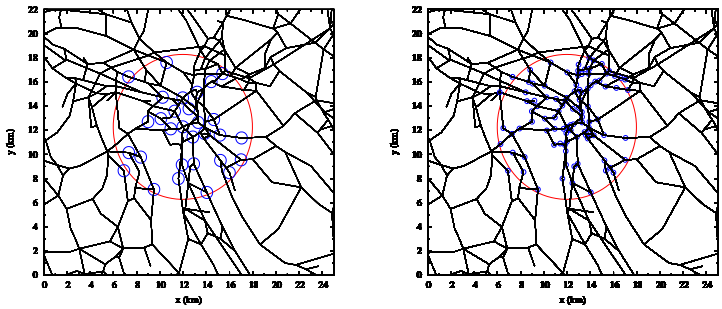

(a) Configuration A

(b) Configuration B

Fig. 7. Zurich scenario: observation area and RoIs for the two configurations A and B

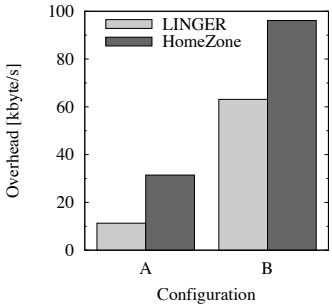

Fig. 8. Zurich scenario: overhead

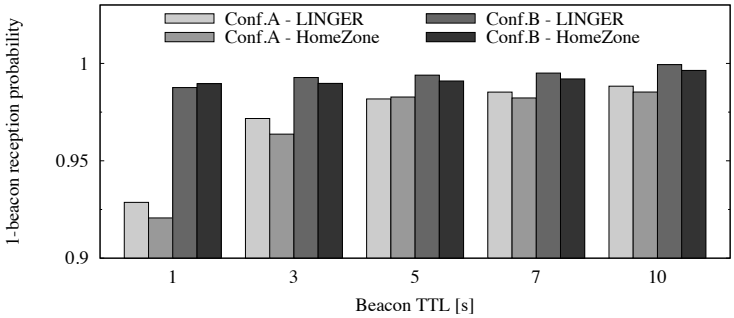

Fig. 9. Zurich scenario: 1-beacon reception probability

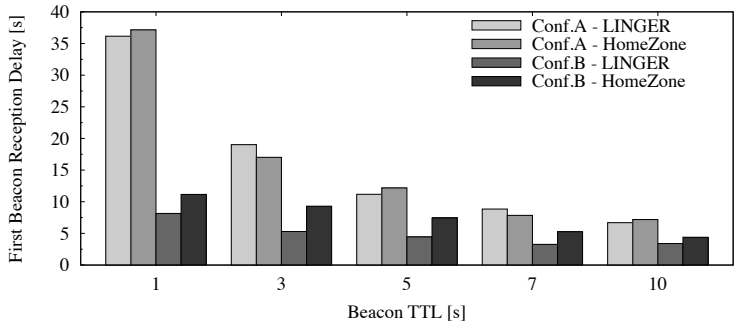

Fig. 10. Zurich scenario: mean value of the first beacon delay

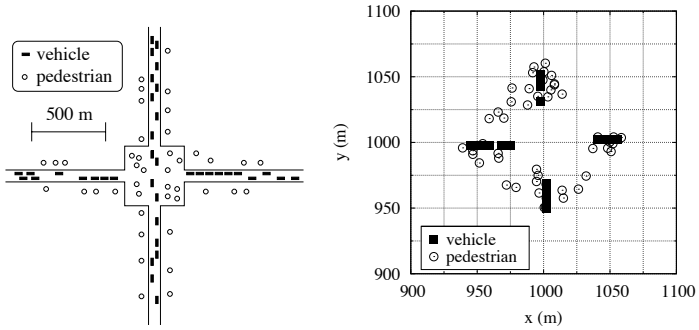

(a) Scenario (b) Bearers

Fig. 11. Scenario with vehicular and pedestrian users (left), type and position of nodes that become bearers during the simulation (right)

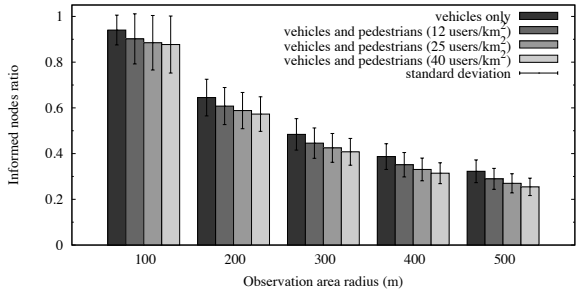

Fig. 12. Informed nodes ratio when pedestrian users are involved in the system. In all cases, $R = 100$ m, the beaconing interval and the TTL are set to 5 seconds and to 1, respectively

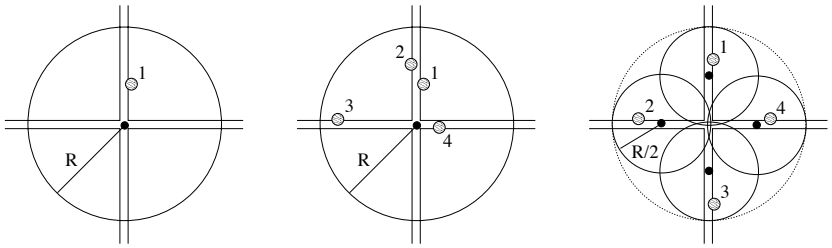

(a) Single copy (b) Concentric RoIs (c) Clover RoIs

Fig. 13. RoIs configurations with multiple copies

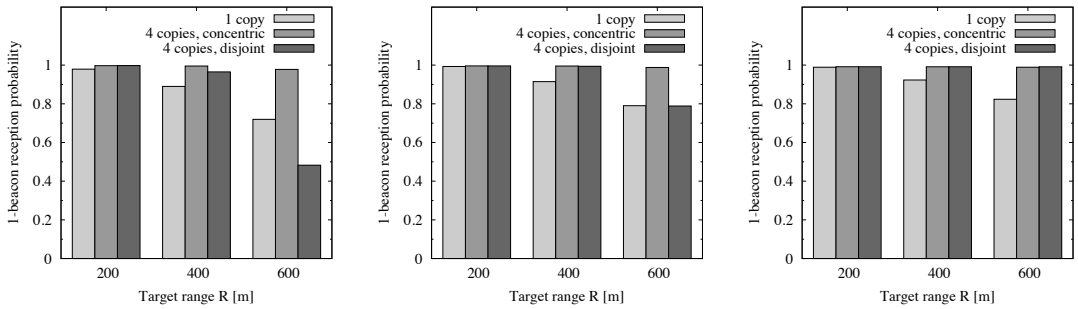

(a) Observation area radius 100 m (b) Observation area radius 200 m (c) Observation area radius 500 m

Fig. 14. 1-beacon reception probability for different RoIs configurations, with varying R and observation area radius

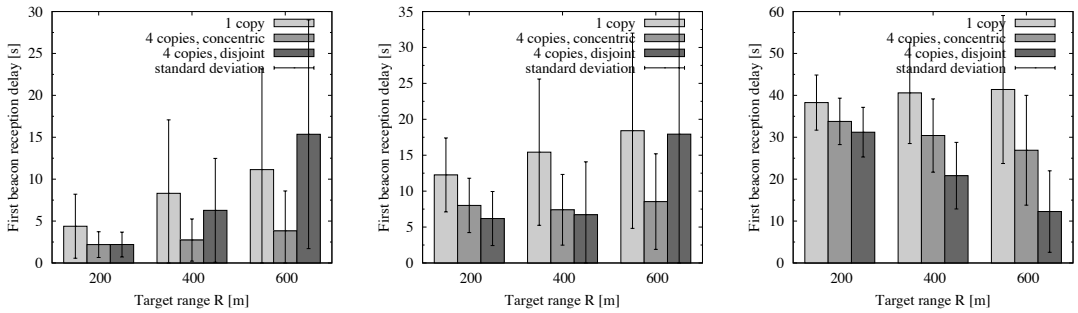

(a) Observation area radius 100 m (b) Observation area radius 200 m (c) Observation area radius 500 m

Fig. 15. First beacon reception delay for different configurations, with varying RoI radius R and observation area radius

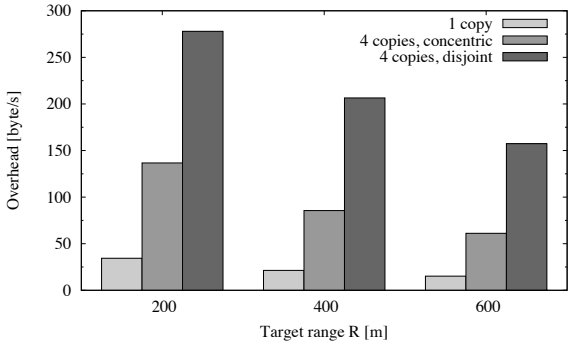

Fig. 16. Overhead (aggregating both data and control messages) for different RoIs configurations, with varying radius R