A Centralized Approach for Setting Floating Content Parameters in VANETs

Antonio Di Maio, Ridha Soua, Maria Rita Palattella, Thomas Engel
University of Luxembourg
Luxembourg
Email: {name.surname}@uni.lu

Gianluca A. Rizzo
HES-SO Valais
Switzerland
Email: gianluca.rizzo@hevs.ch

Abstract—**Floating Content (FC) has recently been proposed as an attractive application for mobile networks, such as VANETs, to operate opportunistic and distributed content sharing over a given geographic area, namely Anchor Zone (AZ). FC performances are tightly dependent on the AZ size, which in literature is classically chosen by the node that generates the floating message. In the present work, we propose a method to improve FC performances by optimizing the AZ size with the support of a Software Defined Network (SDN) controller, which collects mobility information, such as speed and position, of the vehicles in its coverage range.**

I. Introduction

One of the main technical challenges of content dissemination in VANETs is related to the high dynamism of vehicular topology and the volatility of inter-vehicular links, either Vehicle-to-Vehicle (V2V) or Vehicle-to-Infrastructure (V2I) communications, in which infrastructure typically comes in the form of Road Side Units (RSUs). Such volatility hampers the efficient spreading of the content. Floating Content (FC), a push-based communication scheme [1], was proposed to overcome the rapid vanishing of the disseminated content by a vehicle in a specific area, when no support from infrastructure is available. The goal of FC is to replicate a content in all the nodes in a limited spatial region called *Anchor Zone* (AZ), associated to that content. In this way, the content is stored probabilistically in the AZ for the content lifetime, or until its disappearance from the AZ. During such time, the content is made available to nodes traversing the AZ by means of opportunistic replication. More specifically, users that traverse that AZ and do not possess the associated content, can have a copy of it when they are within the transmission range of any vehicle owning a copy of the original content. Hence, the content ends up being available on a set of nodes within the AZ and "float" over time even after that the originating node has left the AZ. When a vehicle exits the AZ and has a copy of the associated content, it considers it as obsolete and deletes it.

From a performance viewpoint, the main challenge in designing the AZ size and shape, and in orchestrating content replications, is to strike a balance between resources utilization (e.g. storage and bandwidth) and utility of storing and replicating the content. In general, this is performed by allocating enough resources to achieve a minimum target value of an application level performance parameter, such as the average time to get content, or the percentage of nodes leaving the AZ with a copy of the message (success probability) [2].

Existing works on FC focus on modeling content availability (the percentage of nodes with content inside the AZ) and success probability in function of various mobility metrics such as node density, average speed, mean rate of node encounters and arrival rate in the AZ [1], [2], [3]. These studies assume that the generating node has perfect knowledge of those parameters, without considering how such knowledge is built. Of course, these parameters are essential for any practical application of FC, as they strongly determine shape and size of the AZ in function of the target performance. Decentralized estimation of various features of mobility patterns for VANETs has been already considered in literature [4], [5]. However, in realistic settings, typically characterized by a high spatial variability in such parameters as mean node density, mean speed, etc., the performance of decentralized estimation approaches may suffer from low accuracy.

In general, all techniques for estimation of such parameters suffer from some degree of uncertainty, which may lead to conservative and inefficient AZ dimensioning, or to poor FC performance. Conservative FC dimensioning has a cost on the system in terms of storage and bandwidth resources, but also mechanisms for mobility estimation consume system resources. Pursuing a high degree of estimation accuracy might put a higher burden on the system than a more conservative dimensioning of FC parameters based on coarser estimations. Such tradeoff depends on specific features of the mobility model, on techniques used to estimate some parameters related to it, and to node density distribution.

The ultimate goal of our work is to determine an optimal mobility estimation method that, for a target value of success probability, maximizes FC efficiency by striking a balance between centralized and distributed approach.

In this paper, we take the first steps at tackling this issue: We focus on FC in VANETs and we analyze how FC performance and resource efficiency may benefit from the availability of a centralized entity, which interacts with moving vehicles, collects relevant mobility data, and build estimates of those mobility parameters for optimal FC dimensioning. In particular, we consider an SDN-based architecture for VANETs as proposed in [6]. RSUs are SDN-enabled controllers that cooperate with vehicles and improve the management of content delivery and dissemination by setting properly the size of the AZ.

II. Network Scenario

As illustrated in Fig. 1, we consider a set of vehicles moving in a geographical area, the *Service Area* (SA), containing a certain number of RSUs. At any time, each vehicle may start disseminating a certain content in a given AZ, contained in the SA. In particular, we are interested in those setups where the density and the coverage range of RSUs do not allow a complete coverage of the SA. Vehicles communicate via V2V between them and via V2I with the RSU, using in both cases the IEEE 802.11p standard, a Medium Access Control (MAC) protocol based on the DSRC spectrum at physical layer [7]. We assume the AZ to be circular, though the analysis can be extended to more complex shapes. When a vehicle wants to create an AZ and start disseminating a content in it, it must first determine the AZ radius and the content lifetime, i.e. the time after which vehicles are allowed to discard the content. AZ radius is determined in function of target values of performance estimators such as *availability* (the mean fraction of users in the AZ with a copy of the floating content) and *success rate* (the mean fraction of users getting out of the AZ with a copy of the floating content), as defined in [2]. These estimators are strongly correlated to nodes density, to their transmission range and their speed. In addition to such performance parameters, we consider also the *FC efficiency*, which takes into account also bandwidth utilization. It is defined as the mean number of transmissions per useful content replication (an useful content replication is a transfer of content within the AZ, from a node with content to a node without it), for a given mean success rate. In an ideal setting where nodes are assumed to have perfect knowledge of mobility, and in which such knowledge comes at no cost, FC efficiency depends on node mobility patterns, on node density, and on the ratio between the AZ and vehicle range radii. In realistic settings, the message exchanges that implement the specific mechanism for mobility estimation have to be taken into account too.

In IEEE 802.11p VANETs, periodic beaconing is crucial to enable network synchronization and packets transmission orchestration. Vehicles access the channel to broadcast their beacons at least once in the Control Channel Interval. In our setting, we assume these beacons broadcast (to other vehicles and to RSUs) information about the originating vehicle's ID, position, speed and direction, as well as a global timestamp. The RSUs store these parameters in a *Mobility Information Table* (MIT). An entry in the table is updated when the RSU receives a beacon from the corresponding vehicle. Besides, an entry in the MIT is deleted if it is not updated during a specific period of time. If a vehicle wants to start floating a specific message, and it is in the RSU range, it sends a request to the RSU (possibly in the form of a Wave Service Message (WSM) [7]) for the estimated measures of average speed and density. The RSU sends back the requested data via WSM. When the vehicle is outside the RSU range, it estimates the AZ radius using a collaborative and distributed approach, as classically proposed in literature.

TABLE I
AZ Model parameters

Parameter	Definition
v_i	ID of the generic i-th vehicle in the SA
v^*	ID of the seeder vehicle in the SA
$t(v_i)$	Timestamp of the mobility information sent by v_i
$p(v_i)$	Position of the vehicle v_i at time $t(v_i)$
$s(v_i)$	Speed of the vehicle v_i at time $t(v_i)$
$R(S, D)$	Radius of the AZ started by the vehicle v^*
$RSU(v_i)$	RSU ID to which v_i is connected
$C(v_i)$	Surface of the v_i range
$C(RSU(v_i))$	Surface of the $RSU(v_i)$ range
$B(p(v^*), \delta)$	Closed disk centered in $p(v^*)$, with radius δ
$V(\alpha)$	Set of vehicles (IDs) in a neighborhood area α
$V_{MIT}(RSU(v_i))$	Set of all the vehicles (IDs) in the MIT of $RSU(v_i)$
$V_C(v^*)$	$V_{MIT}(RSU(v^*))$
$V_\delta(p(v^*), \delta)$	$V_C(v^*) \cap V(B(p(v^*), \delta))$
$V_D(v^*)$	$V(C(v^*))$
D	Density of vehicles in v^* neighborhood
$D_c(V(\alpha))$	Cardinality-based density
$D_d(V(\alpha), p(v^*))$	Distance-based density
S	Average speed of vehicles in v^* neighborhood

III. Mobility-Based AZ shaping model

In this section, we detail the proposed method to compute the AZ radius (R), used by the vehicle that wants to start disseminating the associated FC (*seeder vehicle*). The AZ radius will be function of node density and average speed in a neighborhood of the seeder. When the average speed of the vehicles close to the seeder is high, it is expected that the FC success rate drops due to the reduced time the vehicles have to exchange information. In order to increase the number of vehicles that get the content, it is possible to enlarge the AZ radius, which will imply a higher success rate. Conversely, when the density of vehicles is high, it is likely that the message will reside in several vehicles which tend to form small clusters and therefore keep floating for more time. In this case, it is also important to optimize other practical parameters aside the success rate, such as minimizing message redundancy and memory occupation, by choosing a smaller AZ radius. Table I summarizes the notation used hereafter for describing the mobility-based AZ shaping model. The relationship between AZ radius, average speed and vehicle density [2], [3] is generally in the form of:

$$R(S, D) = w\frac{S}{D} \tag{1}$$

where S is average speed in a *neighborhood* of v^*, D is vehicle density in the same neighborhood, and w is a weight which depends on the specific mobility model, to be either empirically fixed or defined using some dynamic policies. The radius is always computed locally by each seeder vehicle v^*, with S and D obtained either from an RSU, if the seeder is in the RSU range, or through estimations performed collaboratively by vehicles. In both cases, the WSMs messages

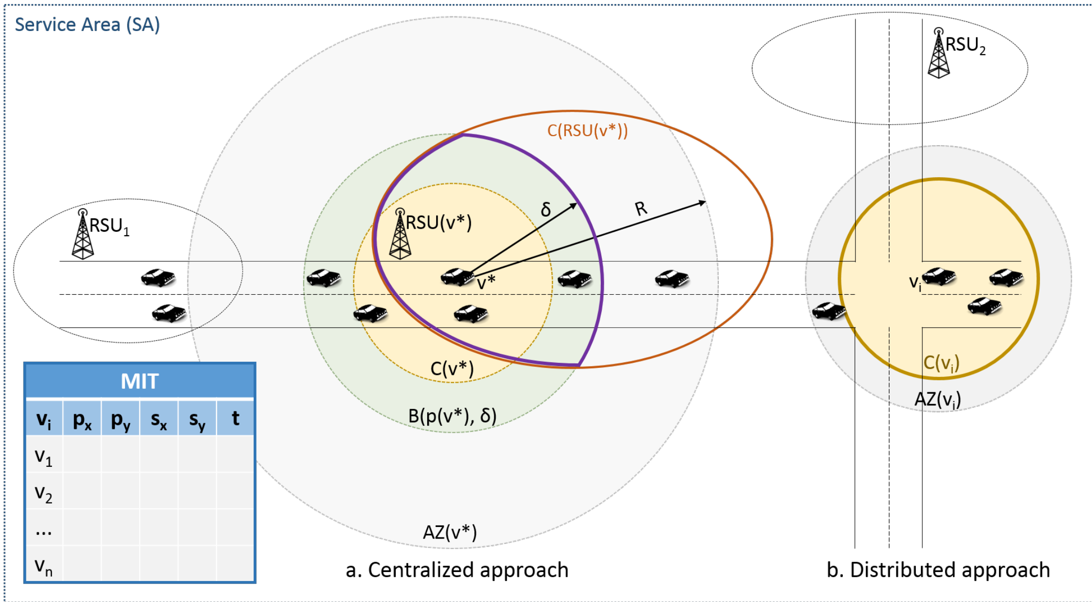

Fig. 1: An example scenario, showing ranges, neighborhoods and AZs in a generic SA. The thick lines represent the neighborhood boundaries for each approach: C-based (orange), δ-based (purple) and distributed (yellow).

sent by each vehicle contain information about its current position, speed, and timestamp:

$$p(v_i) = (p_x(v_i), p_y(v_i)) \in \mathbb{R}^2 \quad (2)$$

$$s(v_i) = (s_x(v_i), s_y(v_i)) \in \mathbb{R}^2 \quad (3)$$

$$t(v_i) \in \mathbb{R}^+ \quad (4)$$

A. Centralized approach

When a seeder vehicle is in the range of an RSU, it can rely on the mobility parameters estimated by the SDN controller. Each RSU can build its own MIT, whose generic structure is shown in Fig.1. The MIT is organized like a database table in which the primary key is the v_i field and therefore no duplicates are allowed.

Considering that the RSU memory is finite, an update policy must be implemented to keep in the MIT only the most up-to-date information about the mobility of each vehicle. For instance, an expiration time can be set so that entries older than the specified time will be discarded. In the centralized approach, a request-response mechanism is devised: the seeder vehicle v^* must choose which kind of neighborhood to use (*C-based* or *δ-based*, described in the following sections) and send to $RSU(v^*)$ the suitable information to determine the values of D, S and the set of vehicles that belong to the selected neighborhood (*neighbor vehicle set*). The $RSU(v^*)$ will compute S (Eq.12) and D (Eq.11 or Eq.10) and send their values back to the vehicle for determining the AZ radius.

1) C-based Neighborhood: In the simpler case, the whole surface of the RSU range, $C(RSU(v^*))$, can be considered as v^* neighborhood. Given that all and only the vehicles that send their mobility information to the RSU will have a valid entry in the MIT, they will also be the only ones included in the neighbor vehicle set $V_C(v^*)$. Therefore:

$$V_C(v^*) = V_{MIT}(RSU(v^*)) \subseteq V(C(RSU(v^*))) \quad (5)$$

2) δ-based Neighborhood: If the range of the RSU is too large, it is possible to define a custom-shaped neighborhood around the seeder vehicle, exploiting the mathematical definition of neighborhood. A two-dimensional closed ball (or closed disk) [8] in a metric space ($\mathbb{R}^2$, d), centered in $p(v^*)$ and with radius δ, is defined as:

$$B(p(v^*), \delta) = \left\{x \in \mathbb{R}^2 \colon d(p(v^*), x) \leq \delta\right\} \quad (6)$$

To physically shape the neighborhood as a circle, we assume as distance function d the euclidean (2-norm) distance in $\mathbb{R}^2$. Any other p-norm distance could be used as well as the Chebyshev distance (∞-norm) or a generic distance function to shape the neighborhood in a custom way.

$$d(x, y) = \|x - y\|_2 \quad (7)$$

Let $V(B(p(v^*), \delta))$ be the set of vehicles in the closed disk of radius δ and centered on the position of the seeder $p(v^*)$. We define the δ-based neighbor vehicle set $V_\delta(p(v^*), \delta)$ as:

$$V_\delta(p(v^*), \delta) = V_C(v^*) \cap V(B(p(v^*), \delta)) \quad (8)$$

The value of δ should be properly chosen. For instance, δ can be heuristically set by the RSU for all the vehicles in its range. In this case, δ could be such that the area of any $B(p(v_i), \delta)$ is a proportional fraction of the total extension of $C(RSU_i)$. It is also possible to let the seeder vehicle determine which is the most suitable value of δ, accordingly to its needs. In this case, the vehicle must communicate the desired value of δ to the RSU, along with the value of $p(v_i)$. The seeder vehicle v^* might initially use a standard value (e.g. equal to the expected or desired AZ radius) and subsequently adjust it using information provided by the RSU. A special case may consist in setting δ equal to the radius of seeder range.

B. Distributed approach

It may happen that the seeder vehicle is outside the range of any RSU in the SA, therefore it does not have any infrastructure support for determining the neighbor vehicle set $V_D(v^*)$. In this case, it can simply consider its own range as neighborhood, resulting in:

$$V_D(v^*) = V(C(v^*)) \tag{9}$$

The seeder collects the mobility information by asking all the vehicles in its range for their position and speed, in a distributed fashion. After having gathered the needed information, the density and average speed values will be computed locally by the seeder and not provided by the RSU as in the centralized approach.

C. Density Function

We propose two different methods for computing the vehicles density D in a given v^* neighborhood. The methods can be applied in both approaches (centralized and distributed), choosing the appropriate neighbor vehicle set V among V_C, V_δ, and V_D.

The density can be computed as a function of the distance between the seeder v^* and every other vehicle in the selected neighbor vehicle set, as in

$$D_d(V, p(v^*)) = \sum_{\substack{v_i \in V \\ v_i \neq v^*}} \|p(v^*) - p(v_i)\|^{-1} \tag{10}$$

Note that Eq.10 returns a *density score* that combines information about number of vehicles in the neighbor vehicle set and information on how physically close they are to the seeder.

An alternative way for estimating the density consists in considering the cardinality of the neighbor vehicle set V, normalized by the area of the v^* neighborhood.

$$D_c(V, A) = \frac{|V|}{A} \tag{11}$$

This definition of density does not necessarily require to know the position of the seeder and it is strictly dependent on the area of the chosen neighborhood. Differently from $D_d(V, p(v^*))$, it is not continuous but discrete with a minimum step of $1/A$. We might prefer using Eq.11 instead of Eq.10 when the information about the neighborhood area is available.

Note that A is equal to the area of $C(RSU(v^*))$, $C(B(p(v^*), \delta))$, or $C(v^*)$, when adopting, respectively, the centralized C-based, centralized δ-based, or distributed approach.

D. Average Speed Function

We define the average speed S as the arithmetic mean of the speed vector modules of the vehicles in a neighbor vehicle set V.

$$S(V) = \frac{1}{|V|} \sum_{v_i \in V} \|s(v_i)\|_2 \tag{12}$$

Similarly to the density, the same formula of the average speed can be used in both centralized and distributed scenarios, choosing the appropriate neighbor vehicle set of vehicles V.

IV. Future Work

The increasing availability of a large amount of information in Vehicular Ad-hoc Networks is giving rise to interest on Floating Content (FC). FC supports infrastructure-less distributed content sharing in an Anchor Zone (AZ). The performance of FC paradigm in function of specific mobility patterns and for specific applications has been at the center of recent investigations. However, the issue of how to estimate the main mobility features, on which AZ dimensioning is based, has not been considered so far. More specifically, the issue of how to maximize FC efficiency by striking a balance between estimation accuracy and conservative dimensioning of FC service is still unsolved. In this paper, we take the first steps at tackling this issue by proposing AZ dimensioning based on centralized (SDN-based) estimation, as well as a dimensioning based on information built collaboratively by vehicles through a distributed approach. The next step of the investigation will encompass the validation of the preliminary hypotheses and the derivation of an overall resource consumption model, for each of the proposed approaches. This work will lay the foundation to determine a resource-optimal strategy for FC dimensioning in VANETs.

References

[1] J. Kangasharju, J. Ott, and O. Karkulahti, "Floating content: Information availability in urban environments," in *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2010 8th IEEE International Conference on*, March 2010, pp. 804–808.

[2] S. Ali, G. Rizzo, V. Mancuso, and M. A. Marsan, "Persistence and availability of floating content in a campus environment," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, April 2015, pp. 2326–2334.

[3] E. Hyyti, J. Virtamo, P. Lassila, J. Kangasharju, and J. Ott, "When does content float? characterizing availability of anchored information in opportunistic content sharing," in *INFOCOM, 2011 Proceedings IEEE*, April 2011, pp. 3137–3145.

[4] G. V. Rossi, K. K. Leung, and A. Gkelias, "Density-based optimal transmission for throughput enhancement in vehicular ad-hoc networks," in *2015 IEEE International Conference on Communications (ICC)*, June 2015, pp. 6571–6576.

[5] S. Zemouri, S. Djahel, and J. Murphy, "A short-term vehicular density prediction scheme for enhanced beaconing control," in *2015 IEEE Global Communications Conference (GLOBECOM)*, Dec 2015, pp. 1–7.

[6] I. Ku, Y. Lu, M. Gerla, R. Gomes, F. Ongaro, and E. Cerqueira, "Towards software-defined VANET: Architecture and service," *Conference: Annual Mediterranean Ad Hoc Networking Workshop MEDHOCNET*, 2014.

[7] C. Campolo, A. Molinaro, and R. Scopigno, *Vehicular ad hoc Networks. Standards, Solutions, and Research*. Springer, 06/2015 2015, editors. [Online]. Available: http://bit.ly/1Fm4TEt

[8] M. O Searcoid, *Metric Spaces*. London: Springer London, 2007, pp. 71–82.