ODCRep: Origin–Destination-Based Content Replication for Vehicular Networks

Article *in* IEEE Transactions on Vehicular Technology · December 2015

Impact Factor: 1.98 · DOI: 10.1109/TVT.2015.2487679

6 authors, including:

Fabrício Aguiar Silva

Universidade Federal de Viçosa (UFV)

28 PUBLICATIONS **115** CITATIONS

SEE PROFILE

Fabrício Benevenuto

Federal University of Minas Gerais

76 PUBLICATIONS **2,111** CITATIONS

SEE PROFILE

Linnyer Beatrys Ruiz

Universidade Estadual de Maringá

53 PUBLICATIONS **504** CITATIONS

SEE PROFILE

Antonio Loureiro

Federal University of Minas Gerais

352 PUBLICATIONS **4,079** CITATIONS

SEE PROFILE

ODCRep: Origin-Destination-based Content Replication for Vehicular Networks

Fabrício A. Silva, Azzedine Boukerche, Thais R. M. Braga Silva,
Fabrício Benevenuto, Linnyer B. Ruiz, and Antonio A. F. Loureiro

***Abstract*—The evolving of vehicular network applications, from simple alert message exchanging to more elaborated and sophisticated systems, boosts the need for content delivery solutions. A useful technique, in such case, is the content replication, in which strategically selected vehicles replicate content and help on the delivery process. However, content replication is particularly challenging in vehicular networks, due to their special characteristics, such as highly dynamic topology, diverse density, and large-scale scenarios. While there has been progress in routing and dissemination solutions for vehicular networks, few studies have concentrated on the content replication problem. To address this issue, we propose an Origin-Destination-based Content Replication (ODCRep) solution that focuses on balancing the number of replicas across the application area. Differently from existing solutions, ODCRep relies only on the origin-destination information, and uses computationally efficient algorithms. Results from exhaustive simulations show that ODCRep can achieve high coverage, yet consuming less resource than existing solutions.**

***Index Terms*—Vehicular networks, content replication, replica allocation, content availability, content delivery**

I. INTRODUCTION

Vehicular network applications are emerging to reality with the objective of making traffic safer, less congested, more informative, and more enjoyable [1], [2], [3]. The advance of those complex applications boosts the need for content delivery solutions, since its efficient delivery is a major requirement for most of them, such as information, advertisement, and entertainment systems.

Nevertheless, the particular characteristics of vehicular networks, also referred to as vehicular ad hoc networks (VANETs), make the task of delivering content easier said than done [4]. First, vehicles move constantly, causing frequent changes in the network topology. In addition, vehicles may face different density scenarios along their way, ranging from high densities in rush hours, or in specific regions, to low densities in highways, or in low traffic hours. Moreover, VANET applications are supposed to operate in large-scale scenarios consisting of many thousands of vehicles. Finally, infrastructure stations and cellular networks may not be able to meet all demands. All these issues, together with the fact that users are willing to receive content quickly, with high quality and low cost, impose challenges to the content delivery problem in VANET scenarios.

Content replication [5], [6] is a powerful concept that may help increasing content availability in VANETs. Replicating content refers to strategically keeping content replicas in vehicles that would act as local content providers. Many benefits are achieved by replicating content in the network, such as offload of the infrastructure, decreasing in end-to-end delay, decreasing in communication cost, and increasing in the quality of services and in users' satisfaction. However, content replication imposes a cost for selecting and maintaining the replicas, a task even more complex when it comes to highly dynamic and large-scale networks.

The objective of this work is to propose and evaluate a city-wide content replication solution for VANETs. Based on the vehicles' origin-destination (O-D) points, the proposed solution, called **O**rigin-**D**estination-based **C**ontent **R**eplication (ODCRep), is flexible to be applied to different applications, and focuses on balancing the number of replicas in the entire city. Before proposing ODCRep, we first characterize a large-scale realistic vehicular mobility model to obtain insight on how appropriate vehicles can be selected as replicas. As result, we find out that some vehicles have particular characteristics, and they seem more appropriate to be content replicas than others. Based on the characterization results, we then propose and evaluate ODCRep through extensive simulation. The results reveal that ODCRep is a flexible, cost-effective solution that can provide a trade-off between content availability and communication cost.

The remainder of this text is organized as follows. Section II presents the state-of-the-art in the content replication area applied to VANET applications. Next, Section III describes the content replication problem and presents the characterization results showing which features may be considered when selecting content replicas. Based on this characterization, we then propose our solution, called ODCRep, in Section IV, and the simulation setup and results in Section V. Finally, Section VI concludes our work and presents potential future efforts.

II. RELATED WORK

Content replication is a well-known topic in the Content Delivery Network (CDN) field [5], [7]. When it comes to VANETs, replicating data may be useful for different purposes, such as content delivery [8] and cooperative systems [9]. However, unlike traditional networks (e.g., Internet) where

---

Fabrício A. Silva and Thais R. M. Braga Silva are with Federal University of Viçosa. Azzedine Boukerche is with University of Ottawa. Fabrício Benevenuto and Antonio A. F. Loureiro are with Federal University of Minas Gerais. Linnyer B. Ruiz is with State University of Maringá.

nodes are static, VANETs' highly dynamic topology poses even greater challenges to the content replication problem. Some studies in the literature have addressed this issue and are described below.

In [10], the authors propose a pro-active data replication solution that, based on mobility prediction, selects replica vehicles. Differently, in [11] content is replicated in Road Side Units (RSUs) based on analysis of real mobility traces in terms of contact rates and duration. Similarly, an optimization solution to place content into RSUs based on its demand and popularity, with the objective of maximizing the content availability, is proposed in [12]. A distributed approach is followed by [13], in which vehicles decide in their vicinity the most appropriated among them to keep replica based on information about speed, contacts, and the number of current replicas nearby. The solution described in [14] takes advantage of a phenomenon called vehicle platoon, which indicates that vehicles frequently travel in groups. The authors then propose a platoon-based data managament that considers the moving groups to prefetch and replicate data. In [15], the authors focus on replicating content for requested video streaming applications in vehicular delay tolerant networks. The replication considers the link stability and communication cost to select the most appropriate vehicles.

Some solutions use replication to schedule a delivery of a requested content. In this case, content providers orchestrate the delivery of a requested content to a client vehicle. In MobTorrent [16], the provider Access Point (AP) considers the expected contact graph to replicate chunks of content to other APs, as well as to other vehicles, with the objective of maximizing the amount of data transferred to the requesting vehicle. Similarly, in [17] the APs maintain contact maps by overhearing messages exchanged between vehicles. On the other hand, RSUs keep track of content availability in TEG-PW [18]. In Figaro [19], content management is performed by entities running on infrastructure computational systems. All these solutions are specifically to applications where clients send request to providers, which then schedule and orchestrate the response. In addition, they rely on the network topology knowledge, which is difficult to obtain in such highly dynamic network as VANETs. Finally, they depend on complex graph-based algorithms. In contrast, our solution focuses on applications where content providers are responsible for delivering a content to target vehicles, i.e., it is not required a request message, and relies only on the origin-destination points and uses computationally efficient algorithms.

Other existing solutions focus on selecting replica vehicles to keep content inside a region of interest (RoI). Examples of such studies include HomeZone [20], Abiding Geocast [21], ARM [22], Linger [23], and VTube [24]. In contrast, the objective of ODCRep is to replicate content in a city-wide scenario, where every vehicle traveling during the application lifetime must be covered.

Epidemic [25], also referred to as Flooding, is a widely used approach to deliver content in mobile ad hoc networks. In Flooding, a received content is forwarded by vehicles in order to achieve a high delivery probability, as well as low delay. However, this may generate a high network overhead, leading to a Selective Flooding strategy [26], [27], [28], [29], [30], where only a subset of the vehicles forward a content. It should be noted that this approach can be used for a city-wide content delivery. Selective Flooding is a simple but yet effective approach taken as a baseline solution in our evaluation study.

As far as we are concerned, Push-and-Track [31] is the closest existing solution to ours, since it relies on infrastructure stations and can also be used for city-wide content delivery. The goal of Push-and-Track is also the selection of vehicles to keep content replicas and help their delivery. Differently from ODCRep, Push-and-Track keeps track of the already covered vehicles, allowing it to decide when to inject new replicas based on the current coverage. We also take Push-and-Track as a baseline solution in our evaluation study as we will see in Section V.

III. CONTENT REPLICATION IN VANETS

In this section, we state the city-wide content replication problem in VANETs. We also discuss some characterization findings, from our previous work, giving us insight on how to select the appropriate replicas.

A. Problem Statement

In this work, we address the problem of delivering content in a city-wide scenario, where all vehicles in the network are expected to be covered. Applications with such demand include advertisements, broadcast news, and general notifications, for example. Content here refers to general digital files, such as alert message, audio, video, text, or image. In short, content replication is the process of selecting appropriate vehicles to keep replicas of content locally, and helping the delivery process. We refer to replica vehicles as the ones keeping content in their local storage and helping delivering it to their neighbors. Different from broadcast, in which all vehicles receiving a content forward it, replica vehicles are selected specifically to the task of delivering a content. The idea is to strategically place copies of content where all interested vehicles can receive it quickly, with quality, and with low cost. In the following, we describe the problem of replicating content in VANETs.

a) Connectivity: Let V and S be the set of vehicles and infrastructure stations, respectively. Graph $G(V \cup S, E)$ represents the connectivity between vehicles, as well as between vehicles and infrastructure stations. An edge $e_{i,j} \in E$ indicates a contact between vehicle $v_i \in V$ and either another vehicle or an infrastructure station, $vs_j \in V \cup S$, where $v_i \neq vs_j$. Two arbitrary vehicles, or a vehicle and an infrastructure station, may have been through many different contacts along their lifetime. Therefore, the total duration $d_{i,j}$ of a contact $e_{i,j}$ is defined as the sum of duration of all those contacts.

b) Content: A content C of size C_s data units[1] may be divided into C_c chunks of size $C_{cs} = \lceil \frac{C_s}{C_c} \rceil$ data units each. A content is considered valid for a period starting at t_c^i and finishing at time t_c^f. We refer to this period as the content's lifetime, within it must be delivered to the target vehicles.

---

[1] We are using the general term *data units* to represent content in terms of its storage unit (e.g., bits, Bytes, and KBytes)

c) Replication: The content replication process aims at selecting a subset $R \subseteq V \cup S$ of vehicles or infrastructure stations to act as content replicas. Many strategies may be adopted to select the most appropriate vehicles, which depend on the application requirements and information about the vehicles. In this work, for example, we propose a replica allocation solution based on the origin-destination points of the vehicles.

d) Coverage: A vehicle v_i may be covered basically by two different approaches: vehicle-to-vehicle (V2V) communication only, or with the help of vehicle-to-infrastructure (V2I) communication. Notice that V2V communication is preferred over V2I, since it is expected to cost less, as well as it should reduce the infrastructure workload. V2V and V2I present different communication capabilities, especially transmission rates, which affect the content delivery process. We define T_{V2V} and T_{V2I} the transmission rates, in data units per time unit, of V2V and V2I communications, respectively. Thus, the time required to transmit a content of size C_s using only V2V communication is $\frac{C_s}{T_{V2V}}$. Then, a vehicle v_i is covered by V2V communication only if its total contact time with replica nodes is sufficient to receive the entire content:

$$\sum_{j \in R \cap V} d_{i,j} \geq \frac{C_s}{T_{V2V}} \tag{1}$$

Nevertheless, a vehicle may also be covered with the help of infrastructure stations. In this case, a vehicle v_i is considered to be covered if its total infrastructure contact duration is sufficient to transmit the remaining parts of the content that could not be transmitted by V2V contacts. This is defined as:

$$\sum_{j \in R \cap S} d_{i,j} \geq \frac{C_s - \frac{T_{V2V}}{\sum_{j \in R \cap V} d_{i,j}}}{T_{V2I}}. \tag{2}$$

e) Time to be Covered: The time a vehicle has to wait before being covered is also an important metric. Thus, we define t_i^t as the time a vehicle v_i departs from its original point, and t_i^c as the time when it is covered. For many applications, the lower the waiting time, $t_i^c - t_i^t$, the better.

f) Delivery Cost: To represent the cost of delivering a content C to a vehicle v_i, let C_{V2I} and C_{V2V} be the cost to transmit a data unit using infrastructure (V2I) and ad hoc (V2V) communication, respectively. In general, $C_{V2I} \gg C_{V2V}$. Then, the cost to transmit a content C to a vehicle v_i is represented by

$$C_c^i = C_s \times C_{V2V}, \tag{3}$$

when only V2V communication is sufficient, and by

$$C_c^i = \frac{T_{V2V}}{\sum_{j \in R \cap V} d_{i,j}} \times C_{V2V} + C_s - \frac{T_{V2V}}{\sum_{j \in R \cap V} d_{i,j}} \times C_{V2I} \tag{4}$$

when the infrastructure stations are part of the communication solution. Thus, the lower the number of transmitted messages through V2I communication, the better.

The formalization of the content replication problem is an important part of the overall process of designing, developing, and deploying a content delivery application in VANETs. However, one question first arises: are there vehicles that present particular characteristics that turn them into better content replicas than others? To answer this question, we discuss the results of a realistic mobility trace characterization in the following.

B. Characterizing Mobility Dynamics

Some studies in the literature have studied the question of how vehicles differ in their mobility properties. Resta and Santi [32] demonstrated analytically that the fully capacity of data dissemination in mobile networks can be achieved only when the best disseminating nodes are selected. Zyba et al. [33] studied different mobility scenarios (taxis and mobile users in university campus) and showed that some mobile nodes can be considered more relevant to data dissemination than others; furthermore, they demonstrated that the dissemination capacity is expected to improve when those nodes are selected appropriately.

Also, some authors have already studied the behavior of mobility traces to discover helpful insight. Xia et al. [34] characterized a mobility trace composed of 12,096 taxis from Beijing during a week, focusing on contact and clustering characteristics. Ahmed et al. [35] analyzed 1,200 buses in Seattle and evaluated their contact duration, inter-contact intervals, and clustering characteristics. Monteiro et al. [36] characterized simplified urban ($4\,\text{km}^2$) and highway ($25\,\text{km}$) scenarios. Uppoor et al. [37] performed characterization of a realistic urban scenario, by analyzing a trace from Cologne, Germany, and evaluating the flow of vehicles and network density.

Nevertheless, other aspects of the network should also be studied to help answering the posed question. Thus, in our previous work [38], we went further and characterized a realistic, large-scale vehicular mobility trace to show that some vehicles present special characteristics related to complex networks and mobility aspects. Among several mobility traces found in the literature, the one described in [39] was chosen, since it represents a realistic, large-scale urban mobility scenario from Cologne, Germany. This trace encompasses the mobility of over 120,000 vehicles between 6:00 am and 8:00 am of a week day with 1 second granularity.

To assess the characterization results, the network was modeled as a graph $G = (V, E)$ in which V represents the set of vehicles and $E \subseteq V \times V$ is the set of edges (e.g., contacts among vehicles). An edge $e = (v_i, v_j)$ between vehicles v_i and v_j exists if the distance $d(v_i, v_j)$ between v_i and v_j was lower than 100 meters for a period of time. The value of 100 meters was chosen based on experimental results from [40]. Although other authors have shown that it was possible to achieve a higher distance in vehicular networks, this coverage range is considered appropriate for real scenarios with obstacles.

In the following, we present the most important findings obtained in our previous work [38]:

- some vehicles present a significantly higher number of contacts than most of others;

- centrality analyzes show that some vehicles present higher values of *betweenness* and *closeness*, indicating they are more important to the network structure than others;
- clustering coefficient analysis indicates that some vehicles are more connected to their neighbors than others;
- a small fraction of vehicles travel for longer periods than most of others;
- a small fraction of connections between vehicles last for longer periods of time than most of others.

All these findings reinforce the assumption that some vehicles are more likely to be better content replicas than others. However, another question arises: what features should be considered to select the best vehicles? To answer this question, we investigated what happens to the network structure when some vehicles with specific characteristics are removed from it. The results reveal that clustering coefficient, degree, and travel time features affect more the network structure.

As result, we conclude that some vehicles have special characteristics that turn them more likely to be good replicas than others, and that some characteristics, such as degree, clustering coefficient, and travel time, are relevant to the identification of such vehicles. However, topology metrics, such as vehicle's degree and clustering coefficient, are difficult to predict before having the entire contact graph. On the other hand, the vehicle's travel time is easily estimated by online navigation systems given the origin-destination (O-D) points, especially with the help of online traffic information applications. Thus, in the next section we propose a solution to select appropriate replica vehicles that relies only on the O-D information.

IV. ODCRep: Origin-Destination-based Content Replication

As stated earlier, some vehicles present particular characteristics that seem to turn them into better replica candidates than others. It turns out that some of these characteristics are difficult to obtain prior to the end of the network lifetime, such as degree, betweenness, closeness, and clustering coefficient. On the other hand, the vehicle's travel time, another important metric, is easily estimated by considering the O-D points and, naturally, an online navigation system. Therefore, we propose a solution called **O**rigin-**D**estination-based **C**ontent **R**eplication (ODCRep) that relies on this information only.

It should be clear that the selection of replicas may impose a communication cost, since replica vehicles will communicate through V2V link with their peers. This may lead to network congestion, bandwidth waste, packet loss, and, consequently, a poor quality of service. Furthermore, covering all vehicles in a large-scale, city-wide scenario is not a trivial task, given the mobility dynamics, and the diverse conditions faced by vehicles. In other words, there is a trade-off between coverage and cost. Therefore, one requirement of our proposed solution is to be flexible to be used by a variety of applications with different demands in terms of content availability and communication cost.

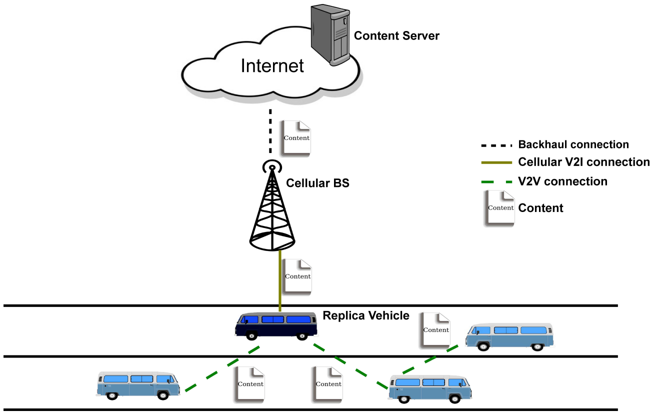

Fig. 1. System model architecture. Vehicles are capable of communicating through cellular (V2I) and V2V networks. Replica vehicles deliver content using V2V only.

A. System Model

Figure 1 illustrates the assumed network architecture. We adopt a centralized approach consisting of a content server with high computational capabilities in terms of memory, disk, and processing units. All vehicles are capable of communicating with the server through cellular network (V2I), as well as with other vehicles through V2V short-distance network (e.g., 802.11p). This is a reasonable assumption, since vehicles are expected to have cellular modules embedded in the near future. Because of commercial and technology constraints, the cellular communication (V2I) costs significantly more, in terms of money, than the V2V solution. This leads to the objective of giving priority to V2V over V2I communication, whenever possible. Only vehicles selected as replicas require a V2I connectivity, since target vehicles receive content through V2V communication. Therefore, the impact when a selected replica losses its connectivity with the infrastructure is that it will not start acting as replica until it recovers its connectivity to receive the content.

Content is assumed to be static, and should be delivered to all vehicles in the network during its lifetime. At departing time, vehicles use their navigation system to estimate the travel time based on their O-D points. The estimated time is then sent to the server through the cellular network, encapsulated in an *ENTER* message. Based on the received information for each round, the server selects the vehicles expected to be good replicas. The server then sends, to the selected vehicles, the content encapsulated into a *CONTENT* message, also through cellular communication. Upon receiving the content, the replica vehicles periodically disseminate it to their peers, through V2V communication.

B. Replica Selection

The ODCRep replica selection process runs in the server, which is responsible for deciding whether an entering vehicle should become a replica. Given that it is not possible to know, in advance, which vehicles will travel for longer periods, ODCRep estimates the probability of a vehicle being a good replica based on current information only. Three metrics contribute to this probability: the vehicle's estimated travel time, the content lifetime, and the vehicle's departure and arrival locations. Notice that these metrics are easily obtained from the content specification, and from the O-D points. Furthermore, each contribution can be parametrized according to the application's demand in terms of cost and content availability. In the following, we describe how each metric contributes to the selection of replica vehicles.

g) Travel time contribution: The vehicles' travel time is an important metric, as presented in [38] and summarized in Section III. ODCRep aims at selecting the vehicles expected to travel for longer periods. As result, we expect high content availability, since those vehicles can contribute significantly in the delivery process.

To achieve this goal, the ODCRep server receives and stores all vehicles' estimated travel time. A vehicle is then considered a replica candidate only if its travel time is higher than a threshold, which separates all estimations into two groups: the first one where $\phi\%$ of the values are smaller, and the second one where $(100-\phi)\%$ are higher than the threshold, considering $\phi \in [0, 100)$ is an application-defined value. Vehicles expected to travel for periods shorter than the threshold are not even considered as replica candidates.

An important issue here is how to efficiently compute the threshold considering that the application is expected to operate in large-scale scenarios. To this end, the estimated travel time of all vehicles are stored in two heap data structures called *maxHeap* and *minHeap*. The *maxHeap* stores the $\phi\%$ of the smallest values, while the *minHeap* stores the $(100-\phi)\%$ highest values. Both heaps satisfy the heap property meaning that the first element of *minHeap* is the smallest value among the $(100-\phi)\%$ highest ones. Similarly, the first element of *maxHeap* is the highest value among the $\phi\%$ smallest ones. Then, the threshold is easily picked as one of the first element of any heap.

In terms of computational complexity, picking the threshold takes $\mathrm{O}\big(1\big)$ operations, given that the heaps satisfy the heap property. To satisfy the heap property, an insertion takes $\mathrm{O}\big(\log\ n\big)$ operations. In addition, after a number of new inserted values, the heaps must be balanced to keep $\phi\%$ of the values in the *maxHeap*, and $(100-\phi)\%$ of them in the *minHeap*, which also takes $\mathrm{O}\big(log\, n\big)$. Thus, the overall process requires $\mathrm{O}\big(\log\ n\big)$ operations.

The parameter $\phi\%$ is used to balance content availability and cost. The highest this value is, the less likely for an arbitrary vehicle to be considered as replica candidate. Consequently, a small number of replicas may be selected, which may reduce the content availability and save communication resources. In contrast, more vehicles may be selected as replicas when $\phi\%$ is low. Thus, the application designer should set $\phi\%$ based on the application demands.

Notice that the threshold approximates the real value (i.e., the threshold for all vehicles) as new vehicles enter the network. Therefore, this value may be not accurate in the beginning of the application, which may lead to the selection of not so good replicas. To help ease this problem, other metrics also contribute to the replica selection, as described next.

h) Content lifetime contribution: The purpose of ODCRep is to select replica vehicles to help deliver content to all vehicles traveling during the content's lifetime. Therefore, the content expiration time plays an important role on the replica selection process.

City-wide content delivery applications may demand different delivery coverage, relative to the content lifetime. While some applications may require that all vehicles, no matter their departure time, must be covered, others may be more flexible, and are not affected by uncovered vehicles departing close to the content's lifetime expiration. Therefore, our solution should be compliant with this diversity of application demands.

To this end, the content lifetime contribution has three categories. The first one represents the period encompassing less than 25% of the content lifetime. The second category represents the period between 25% and 75% of elapsed lifetime. Finally, the last one represents the period from 75% of elapsed time to the expiration time. Thus, the application's designer can configure the content lifetime contribution according to its expected delivery over time, as explained below.

Let $t_e = \frac{t - t_c^i}{t_c^f - t_c^i}$ be the fraction of time elapsed after the content lifetime began, where t is the current time, and t_c^i and t_c^f define the content lifetime interval. Thus, we define a linear function $w_t(t_e)$ representing how this metric contributes to the selection process in terms of the percentage of the elapsed time, t_e, after the content lifetime begins:

$$w_t(t_e) = \begin{cases} \alpha_t \times t_e + \beta_{t1} & \text{if } 0 < t_e < 0.25 \\ \alpha_t \times t_e + \beta_{t2} & \text{if } 0.25 \leq t_e < 0.75 \\ \alpha_t \times t_e + \beta_{t3} & \text{if } t_e \geq 0.75 \end{cases}$$

where α_t indicates the changing rate in the function according to t_e, and $\beta_{t1} < \beta_{t2} < \beta_{t3}$ are the coefficients.

The parameters of this function should also be used to balance content availability and communication cost, based on the application's demand. The function $w_t(t_e)$ is flexible enough to achieve different application demands in terms of content availability over time. The parameter α_t defines the weight of this metric to the replica selection. In addition, the parameters β_{t1}, β_{t2}, and β_{t3} define the importance of each period of the content lifetime in the replica selection. For example, when content should be delivered to all vehicles, no matter their departure time, all of the three parameters should have an identical, high value. On the other hand, when content is not useful at the ending of its lifetime, the β_{t3} parameter should be assigned to a negative high value, so less vehicles will be selected as replicas in this period.

i) Departure and arrival areas contribution: The vehicle's departure (origin) and arrival (destination) areas also play a determining role in the replica selection. Since ODCRep focuses on city-wide scenarios, it attempts to balance the replica placements across the entire map to increase the content availability. Otherwise, unbalanced allocated replicas may lead to uncovered and over-covered areas.

Thus, we take advantage of the results obtained in a previous work [41], where we determine how the number of departures and arrivals are distributed across downtown and suburb regions, in terms of their demographic density. Based on those results, we estimate the number of vehicles departing from and arriving at each quadratic area of 1000-meter side, and then we calculate the number of vehicles to be selected as replica for each area as a percentage θ of the total estimated.

To this end, let d_k and a_k be the estimated number of vehicles departing from and arriving at area k, respectively. For each of the vehicles departing from or arriving at area k, we want to select up to $\theta\%$ of them as content replicas, according to the application demands. In addition, let d_k^r and a_k^r be the number of already selected replicas departing from and arriving at area k, respectively. Thus, $d_k^r \leq d_k \times \theta$ and $a_k^r \leq a_k \times \theta$.

We define two functions, $w_d(k)$ and $w_a(j)$, to calculate the weight to be assigned to a vehicle in terms of its departure area k, and its arrival area j, as:

$$w_d(k) = -\alpha_d \times (\frac{d_k^r}{d_k \times \theta} \times 100) + \beta_d \tag{5}$$

and

$$w_a(j) = -\alpha_a \times (\frac{a_j^r}{a_j \times \theta} \times 100) + \beta_a \tag{6}$$

where constants α_d and α_a must be in $(0, 1]$ range, and β_d and β_a must be assigned a value in the $(0, 100]$ range.

In summary, since both functions $w_d(k)$ and $w_a(j)$ are negatives, the higher the number of already selected vehicles departing from or arriving at an area, the lower the chance of a vehicle to be selected. The parameter θ is also used to help balance the number of replicas over space and time. A higher value will lead to more replicas, and, consequently, higher communication costs. A lower value, on the other hand, will reduce communication costs, as well as content availability. Therefore, θ should be assigned to a value to meet the application demands.

j) Putting all together: After defining the contributions of the vehicles' travel time, content lifetime, and vehicles' origin-destination areas to the replica allocation process, we now put them together. Algorithm 1 shows the procedure running in the server after receiving an *ENTER* message from vehicle v.

C. *Content Delivery*

The vehicles selected as replicas must act as local providers, helping delivering content to their peers using V2V communication only. To this end, upon receiving the content from the server, a replica vehicle delivers it periodically to its peers, every δ seconds. Given the dynamic topology of VANETs, a contact between a replica vehicle and a client interested in content may last for a short period. Therefore, the value for parameter δ should be chosen to increase the chance of a client to be covered, even for a short duration contact with the replica. On the other hand, notice that the lower the δ, the higher the number of messages.

V. PERFORMANCE EVALUATION

We conduct extensive simulations to assess the performance of ODCRep when compared with two existing solutions. Given the large-scale nature of VANETs, it is important to evaluate our solution under a large-scale mobility scenario. Due to the computational complexity of existing network simulators [42], we evaluate ODCRep through two complementary approaches: large-scale and network-enabled. In the large-scale approach, we adopt a realistic large-scale mobility scenario to measure how ODCRep performs when a large number of vehicles is involved, in an ideal network scenario, where packets are not lost and the communication conditions are always satisfied. To this end, we implemented a large-scale simulator in the *R* environment. On the other hand, in the network-enabled approach, we adopt the OMNET++ network simulator, in which vehicles implement the WAVE (wireless access for vehicular environment) suite [43] that includes the IEEE 802.11p standard [44] as MAC and physical layers and the IEEE 1609 protocol suite to define the upper-layer operations. In the network-enabled study, we adopt a realistic mobility scenario of the city of Ottawa, Canada.

A. *Baseline Solutions*

We compare ODCRep with two existing solutions that are also applied to city-wide scenarios: Push-and-Track [31] and Selective Flooding [25]. Push-and-Track presents similar objective and system model architecture as ODCRep. Selective Flooding, on the other hand, is a totally distributed approach that does not rely on infrastructure stations. Thus, by comparing ODCRep against both of them, we can measure how well it performs compared to a similar infrastructure-based architecture, as well as to a different distributed architecture.

In Push-and-Track, the infrastructure server, accessible through the cellular network, keeps track of all target vehicles already covered (i.e., received the content). A vehicle sends to the server an *ENTER* message when it becomes a target for a content, and a *LEAVE* message when it is no longer a target. By the time a vehicle receives the content, it sends an *ACK* message to the server. These messages are transmitted only through the cellular network. Based on the received *ACK* messages, the server is aware of how many vehicles are still uncovered. Furthermore, the server expects a linear coverage behavior, where at least $p\%$ of the target vehicles are expected to be covered after $p\%$ of the content lifetime has elapsed. If this coverage expectation is not satisfied, the server randomly selects new replicas and sends the content to them using the cellular network. The replicas deliver content periodically every δ seconds. Push-and-Track also defines a panic zone starting some time before the content expiration time, when all uncovered vehicles receive the content through the cellular network.

Selective flooding is well adopted in the literature and refers to the idea of vehicles selectively forwarding a content to their

Algorithm 1 - Server Replica Selection

Input: ENTER message msg from vehicle v (containing its travel time $travelTime$, departure area $dArea$ and arrival area $aArea$), $maxHeap$ and elapsed content lifetime t_e

Output: CONTENT message only if v is selected as replica.

```
procedure IsReplica(v)
    $threshold \leftarrow maxHeap[1]$
    if $msg.travelTime \geq threshold$ then                          ▷ Is $v$ a replica candidate?
        $wLifetime \leftarrow w_t(t_e)$
        $wDeparture \leftarrow w_d(msg.dArea)$
        $wArrival \leftarrow w_a(msg.aArea)$
        $probReplica \leftarrow (wLifetime + wDeparture + wArrival)/3$
        if $probReplica \geq Random(0, 100)$ then       ▷ $v$ is replica with probability $probReplica$
            send CONTENT message to $v$
        end if
    end if
end procedure
```

neighbors. By selectively, we mean that only vehicles expected to provide good dissemination performance and coverage will forward a content. Then, upon receiving a content, a vehicle decides whether to forward it to its neighbors. Different strategies may be used to decide forwarding a content, such as the ones based on the vehicle position, network density and randomness. In this work, we adopt a random approach in which vehicles only forward a content with f_p probability. Furthermore, the parameter θ is also used here to indicate the number of vehicles to start the application acting as local providers. A server in this case is responsible for sending, through cellular network, the content to these vehicles.

B. Simulation Setup

1) Configuration Parameters: As described in Section IV, ODCRep was proposed to operate in a hybrid architecture model where vehicles are able to communicate through cellular network (V2I), as well as among themselves using ad-hoc communication (V2V). For the latter, the IEEE 802.11p standard is assumed to be implemented in all vehicles. The entire map is assumed to be covered by the cellular network that is able to transmit data at a rate up to $T_{V2I} = 1\,\mathrm{Mbps}$. Based on the results obtained in [40], the V2V communication range is assumed to be of 100 meters and the transmission rate up to $T_{V2V} = 3\,\mathrm{Mbps}$. When it comes to transmission cost, it is assumed that $C_{V2I} \gg C_{V2V}$. Therefore, the objective is, in general, to reduce the V2I communication, which is the most expensive one.

The application aims at delivering a content of size $C_s = 100\,\mathrm{KBytes}$ represented by a single fragment of the same size. Thus, a contact lasting for one second is sufficient to the entire delivery. To evaluate the performance in long running applications, the content lifetime is assumed to last for the entire running period of the application. Given the city-wide coverage demand, all vehicles are expected to be covered.

Table I presents the ODCRep specific parameters, as well as the ones for the baseline solutions, configured for the simulation studies. These values were chosen with the objective of balancing the trade-off between content availability and communication cost. To demonstrate the flexibility of our solution, we suppose a low effect on the application's performance for uncovered vehicles departing close to the content expiration lifetime. This is represented by the negative, high value for parameter β_{t3}. The other values were chosen after plotting and observing the behavior of the functions $w_t(t_e)$, $w_d(k)$, and $w_a(j)$ under different inputs.

TABLE I
SIMULATION CONFIGURATION PARAMETERS.

Parameter	**Value**
ϕ	98.0%
α_t	4.0
β_{t1}	0.0
β_{t2}	100.0
β_{t3}	-400.0
α_a	0.8
β_a	80.0
α_d	0.8
β_d	80.0
δ	1 second
f_p (Selective Flooding only)	0.1
Panic Zone (Push-and-Track only)	10 s prior to the end
θ	[10%, 50%] for large-scale study 30% for network-enabled study
Simulation time	7200 s for large-scale study 3600 s for network-enabled study
Simulation area	400 km^2 for large-scale study 9 km^2 for network-enabled study

2) Evaluation Metrics: To measure the content availability and the delivery cost, the following metrics are adopted:

- coverage: the percentage of covered vehicles;
- time to be covered: waiting time before being covered;
- number of messages: V2I and V2V exchanged messages;
- redundancy: total of unnecessary content delivered.

In summary, the higher the coverage, and the lower the time to be covered, the redundancy, and the messages exchanged, the better.

C. Large-scale study

The objective here is to measure how ODCRep performs under a large-scale scenario. Given the computational constraints of existing network simulators [42], we have implemented our own simulation environment in the R platform, considering a large-scale mobility trace, and ideal network conditions (i.e., guarantee of packet delivery, and desired transmission rates). The adopted trace [39] encompasses over 120,000 vehicles traveling from 6:00 am to 8:00 am (i.e., 7200 s) of a week day in the city of Cologne, Germany. We vary θ from $[0.1, 0.5]$ (i.e., $[10\%, 50\%]$) to evaluate the impact of the number of replicas from each departing and arriving areas. All results represent the average and the 95% confidence interval from 33 simulations.

Figure 2 presents the content availability results. Notice in Figure 2(a) that, the higher the θ, the higher the coverage for ODCRep, since more vehicles are selected as replicas, according to functions $w_d(k)$ and $w_a(j)$. In addition, the coverage achieved by ODCRep reaches 96% when $\theta = 50\%$. Not surprisingly, Push-and-Track achieved 100% coverage, since all uncovered vehicles are covered by the server during the panic zone. Selective Flooding also achieved high coverage, because of its flooding spread mechanism. However, there is a cost to be paid by those solutions, as discussed later.

Figure 2(b) demonstrates the flexibility of our solution. As stated earlier, we configured ODCRep to be flexible in the delivery for late departing vehicles, by setting the parameter β_{t3} to a high negative value. In fact, this demand was accomplished, as shown in Figure 2(b), since most uncovered vehicles depart after 75% of the content lifetime has elapsed.

Another important metric is the time to be covered, which measures how long a vehicle waits before being covered. Figure 2(c) shows the Complimentary Cumulative Distribution Function (CCDF) of the waiting time of all covered vehicles, assuming $\theta = 0.5$. Notice that vehicles in ODCRep and Selective Flooding have higher probabilities of being quickly covered, when compared to Push-and-Track. In fact, many vehicles in Push-and-Track are only covered in the panic zone period, due to its unbalanced replica allocation. Again, the balanced replica allocation achieved by ODCRep leads to a higher chance of uncovered vehicles to be in contact with replicas. This is also true in Selective Flooding.

The communication cost is also an important metric, particularly when it comes to V2I (i.e., infrastructure) transmissions, which is more expensive than the V2V (i.e., ad-hoc) one. Since *ENTER* (Push-and-Track and ODCRep), *LEAVE* and *ACK* (Push-and-Track) messages are relatively short, when compared to content, they are not considered here for the sake of simplicity. However, it is known that ODCRep and Push-and-Track require as many *ENTER* messages as the number of vehicles, and Push-and-Track requires as many *ACK* messages as the number of covered vehicles. As illustrated in Figures 3(a) and 3(b), ODCRep requires less V2I and V2V messages than the baseline solutions. In addition, θ does not affect significantly the results of ODCRep. Notice that the higher the θ, the higher the number of infrastructure messages exchanged in Selective Flooding, since the number of initial replica vehicles increases with θ.

Redundant messages refer to duplicated messages received by vehicles, which lead to a waste of network resources. The number of redundant messages is significantly higher for Selective Flooding and Push-and-Track, when compared to ODCRep, as illustrated in Figure 3(c). ODCRep reduces the number of redundant data by balancing the content replication according to vehicles' departing and arriving areas. On the other hand, Selective Flooding and Push-and-Track do not use such information, which may lead to unbalanced replica allocation.

In general, we can state that ODCRep could deliver content quickly to a high number of target vehicles (around 96%), by consuming significantly less network resources than Push-and-Track and Selective Flooding.

D. Network-enabled study

The objective here is to evaluate ODCRep when vehicular specific network protocols are used. To this end, we have implemented the three solutions in the OMNET++[2] network simulator, which provides the WAVE (wireless access for vehicular environment) suite [43] that includes the IEEE 802.11p standard [44] as MAC and physical layers and the IEEE 1609 protocol suite to define the upper-layer operations. Due to computational constraints, a 9 km^2 area of Ottawa, in Canada, was adopted.

To improve the realism of the physical layer even more, we also adopt the shadowing model described in [45]. This is a realistic model for urban environments based on IEEE 802.11p measurements, and simulates signal attenuation caused by buildings. To simulate realistic vehicle movements, we take advantage of the mobility model defined by SUMO[3]. We run simulations under low, medium, and high network density to evaluate different scenarios. All results represent the average and the 95% confidence interval from 33 simulations.

The coverage, shown in Figure 4(a), was nearly 100% for all solutions. ODCRep achieved over 99% of coverage for all density scenarios (from 20 to 240 $vehicles/km^2$). Similar to the large-scale results, both Push-and-Track and Selective Flooding were able to achieve practically 100% of coverage, no matter the network density. In terms of time to be covered (Figure 4(b)), our solution was able to deliver content to target vehicles as soon as they enter the network for all scenarios, except for the lower-density one (20 $vehicles/km^2$). This is due to the fact that ODCRep balances the replicas over time and space, leading to uncovered areas in the ending of the application. In contrast, Push-and-Track leads to very high time to be covered for the high-density scenario (240 $vehicles/km^2$) because of its unbalanced replica allocation. In fact, many vehicles were only covered during the panic zone period. Selective Flooding, on the other hand, was very effective in this metric no matter the network density.

ODCRep also required less V2I messages than the baseline solutions, as shown in Figure 5(a). This is an important result, given that V2I transmissions are more expensive than V2V.

[2] http://www.omnetpp.org

[3] http://sumo-sim.org

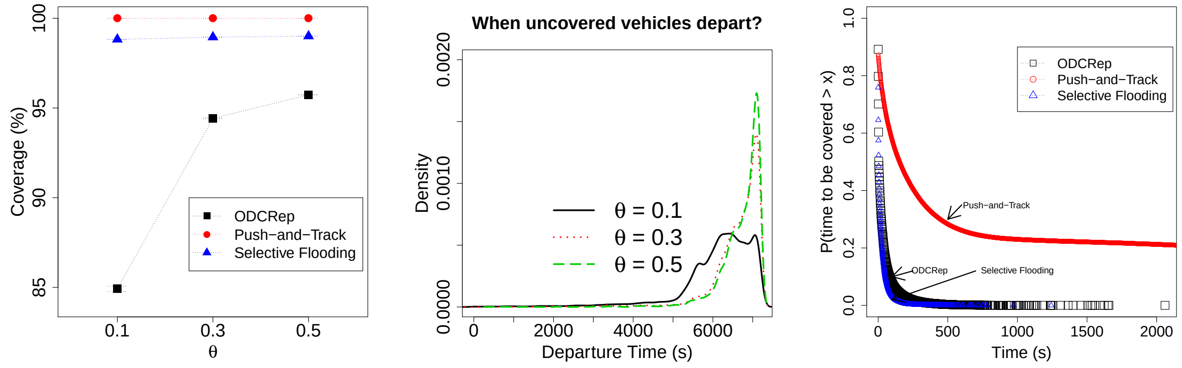

(a) Coverage

(b) ODCRep Uncovered Vehicles Departure Time

(c) Time to be Covered

Fig. 2. Content availability results for the network-enabled study. ODCRep achieved good coverage results together with low time to be covered. In addition, ODCRep's flexibility is demonstrated since most uncovered vehicles depart close to the end of the content lifetime.

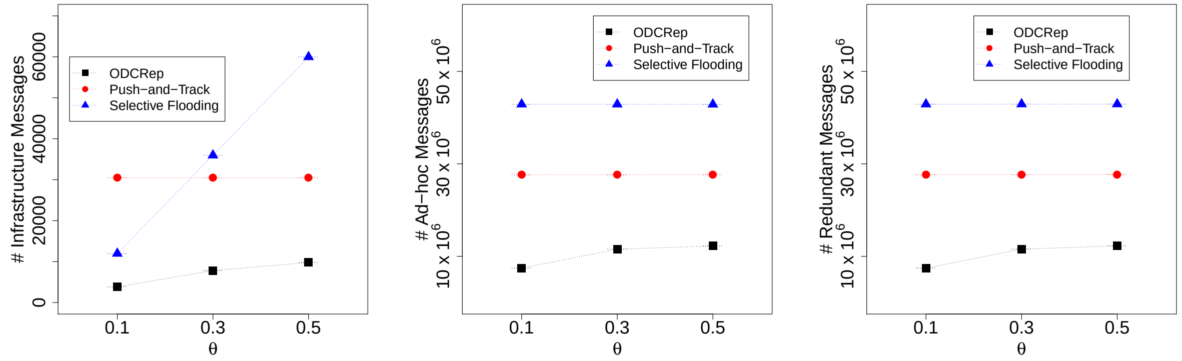

(a) Infrastructure (V2I) Messages

(b) Ad-hoc (V2) Messages

(c) Redundant Messages

Fig. 3. Delivery cost in terms of messages exchanged and redundant messages for the large-scale study. OCDRep outperformed the baseline solutions since it balances the number of replicas over time and space.

In addition, the higher the number of vehicles, the higher the number of V2I messages, since more replicas are selected. This is also true for the Selective Flooding approach. Push-and-Track performed poorly in this metric, as well as the number of V2V exchanged messages, as shown in Figure 5(b). Again, this is due to the fact that replicas are selected in a way that they are balanced across the areas. On the other hand, Push-an-Track requires more messages because of its replica allocation based on the current coverage, and also due to unbalanced selection.

The amount of redundant messages increases with the number of vehicles for all solutions, as illustrated in Figure 5(c). Push-and-Track presents the worst results when it comes to redundant messages for lower-density scenarios, due to its unbalanced replica allocation. However, Push-and-Track reduces the number of redundant messages for the higher-density scenario because many vehicles are only covered during the panic zone period by V2I communication.

VI. CONCLUSION

Content replication is a very attractive idea to increase content availability in VANET applications. However, specific characteristics of VANETs, such as highly dynamic topology and large-scale scenarios, impose challenging issues to this task. To tackle this problem, we propose and evaluate a content replication solution, called ODCRep, that relies on the vehicles' origin-destination points and on efficient algorithms. Through extensive simulation, we have shown that ODCRep could lead to high coverage results by balancing the replica placements, yet consuming less network resources, when compared to two existing solutions.

In this work, we advance the state of the art in the VANET content replication area. Nevertheless, effort should be done before having efficient, cost-effective, and large-scale solutions deployed in real scenarios. To this end, it is important to make ODCRep dynamic to adapt in diverse situations, such as during a long period of traffic congestion caused by unexpected events. In addition, we intend to evaluate ODCRep for different application demands, like delay-sensitive and geo-localized content. Finally, other existing proposals should be used as baseline solutions to help improving ODCRep.

VII. ACKNOWLEDGEMENT

This work was partially supported by NSERC Grant, Canada Research Chair Program, CNPq, CAPES and FAPEMIG. Both Fabrício Silva and Antonio Loureiro have been visiting University of Ottawa.

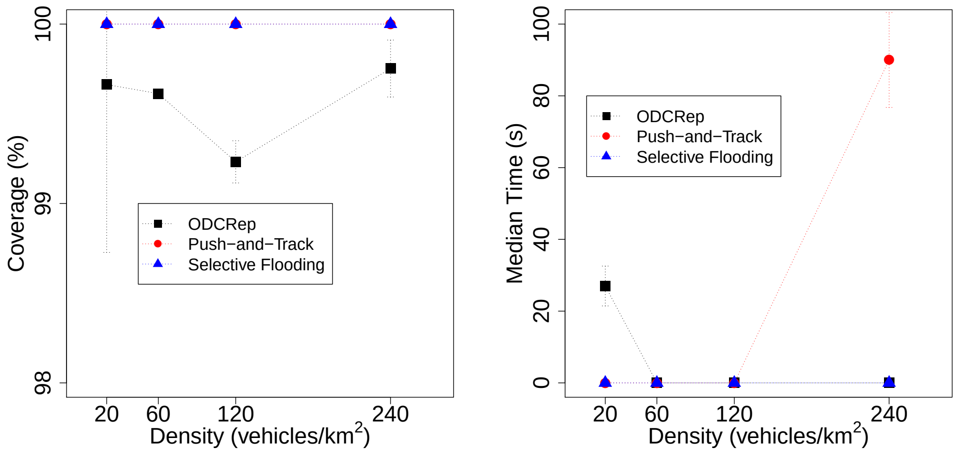

(a) Coverage for network-enabled study.

(b) Time to be covered for network-enabled study.

Fig. 4. Coverage and time to be covered results for the network-enabled scenario.

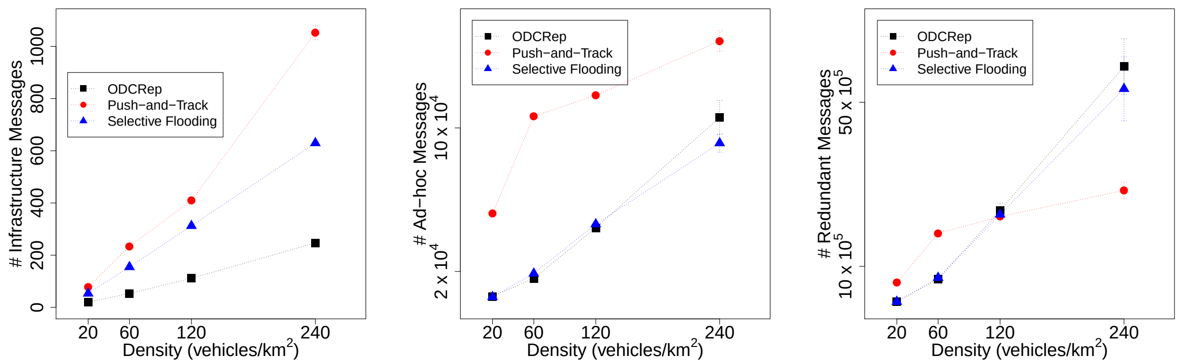

(a) Infrastructure (V2I) messages

(b) Ad-hoc (V2V) messages

(c) Redundant messages

Fig. 5. Delivering cost in terms of messages exchange for the network-enabled study.

REFERENCES

[1] E. Lee, E.-k. Lee, and M. Gerla, “Vehicular Cloud Networking: Architecture and Design Principles,” *IEEE Communications Magazine*, no. 2, pp. 148–155, 2014.

[2] S. Zeadally, R. Hunt, Y.-S. Chen, A. Irwin, and A. Hassan, “Vehicular ad hoc networks (VANETS): status, results, and challenges,” *Telecommunication Systems*, vol. 50, pp. 217–241, Aug. 2012.

[3] T. Willke, P. Tientrakool, and N. Maxemchuk, “A survey of inter-vehicle communication protocols and their applications,” *IEEE Communications Surveys & Tutorials*, vol. 11, no. 2, pp. 3–20, 2009.

[4] M. Gerla, C. Wu, G. Pau, and X. Zhu, “Content distribution in VANETs,” *Vehicular Communications*, vol. 1, no. 1, pp. 3–12, 2014.

[5] G. Peng, “CDN: Content Distribution Network,” *arXiv preprint cs/0411069*, 2004.

[6] C.-a. La, P. Michiardi, C. Casetti, C.-f. Chiasserini, and M. Fiore, “Content Replication in Mobile Networks,” *IEEE Journal on Selected Areas in Communications*, vol. 30, pp. 1762–1770, Oct. 2012.

[7] A. Pathan and R. Buyya, “A Taxonomy and Survey of Content Delivery Networks,” tech. rep., Grid Computing and Distributed Systems Laboratory, Melbourne, 2007.

[8] F. Silva, A. Boukerche, T. R. M. Silva, L. B. Ruiz, E. Cerqueira, and A. A. Loureiro, “Content replication and delivery in vehicular networks,” in *ACM International Symposium on Development and Analysis of Intelligent Vehicular Networks and Applications*, DIVANet ’14, (New York, NY, USA), pp. 127–132, ACM, 2014.

[9] J. Dias, J. Rodrigues, C. Mavromoustakis, and F. Xia, “A cooperative watchdog system to detect misbehavior nodes in vehicular delay-tolerant networks,” *Industrial Electronics, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2015.

[10] J. Gossa, A. G. Janecek, K. Hummel, W. N. Gansterer, J.-M. Pierson, *et al.*, “Proactive replica placement using mobility prediction,” in *Mobile Data Management Workshops, 2008. MDMW 2008. Ninth International Conference on*, pp. 182–189, IEEE, 2008.

[11] Y. Li, D. Jin, P. Hui, and S. Chen, “Contact-aware data replication in roadside unit aided vehicular delay tolerant networks,” *Mobile Computing, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2015.

[12] F. Bruno, M. Cesana, M. Gerla, G. Mauri, and G. Verticale, “Optimal content placement in icn vehicular networks,” in *5th International Conference on Network of the Future*, IEEE, 2015.

[13] N. Kumar, J. J. Rodrigues, J. Lloret, and A. Dua, “Replication-aware data dissemination for vehicular ad hoc networks using location determination,” *Mobile Networks and Applications*, pp. 1–17, 2015.

[14] Y. Zhang and G. Cao, “V-pada: Vehicle-platoon-aware data access in vanets,” *Vehicular Technology, IEEE Transactions on*, vol. 60, no. 5, pp. 2326–2339, 2011.

[15] N. Kumar and J. Kim, “Probabilistic trust aware data replica placement strategy for online video streaming applications in vehicular delay tolerant networks,” *Mathematical and Computer Modelling*, vol. 58, no. 1, pp. 3–14, 2013.

[16] B. B. Chen and M. C. Chan, “MobTorrent: A Framework for Mobile Internet Access from Vehicles,” in *Conference on Computer Communications*, pp. 1404–1412, IEEE, Apr. 2009.

[17] O. Trullols-Cruces, M. Fiore, and J. Barcelo-Ordinas, “Cooperative Download in Vehicular Environments,” *IEEE Transactions on Mobile Computing*, vol. 11, pp. 663–678, Apr. 2012.

[18] F. Malandrino, C. Casetti, C.-F. Chiasserini, and M. Fiore, “Content download in vehicular networks in presence of noisy mobility prediction,” *Mobile Computing, IEEE Transactions on*, vol. 13, no. 5, pp. 1007–1021, 2014.

[19] F. Malandrino, C. Casetti, and C.-f. Chiasserini, “Content Discovery and Caching in Mobile Networks with Infrastructure,” *IEEE Transactions on Computers*, vol. 61, pp. 1507–1520, Oct. 2012.

[20] I. Leontiadis, “Publish/subscribe notification middleware for vehicular networks,” in *Middleware Doctoral Symposium*, pp. 1–6, ACM, 2007.

[21] C. Maihöfer, T. Leinmüller, and E. Schoch, “Abiding Geocast: Time–stable Geocast for Ad Hoc Networks,” in *International Workshop on Vehicular Ad Hoc Networks*, pp. 20–29, ACM, 2005.

[22] D. Borsetti, M. Fiore, C. Casetti, and C.-F. Chiasserini, “An application-level framework for information dissemination and collection in vehicular networks,” *Performance Evaluation*, vol. 68, pp. 876–896, Sept. 2011.

[23] M. Fiore, C. Casetti, C.-f. Chiasserini, and D. Borsetti, “Persistent Localized Broadcasting in VANETs,” *IEEE Journal on Selected Areas in Communications*, vol. 31, pp. 480–490, Sept. 2013.

[24] T. Luan, L. Cai, J. Chen, X. Shen, and F. Bai, “Engineering a distributed infrastructure for large-scale cost-effective content dissemination over urban vehicular networks,” *Vehicular Technology, IEEE Transactions on*, vol. 63, pp. 1419–1435, March 2014.

[25] A. Vahdat, D. Becker, *et al.*, “Epidemic routing for partially connected ad hoc networks,” tech. rep., Technical Report CS-200006, Duke University, 2000.

[26] C. Rezende, A. Mammeri, A. Boukerche, and A. A. Loureiro, “A receiver-based video dissemination solution for vehicular networks with content transmissions decoupled from relay node selection,” *Ad Hoc Networks*, vol. 17, no. June, pp. 1–17, 2014.

[27] G. Maia, C. Rezende, L. A. Villas, A. Boukerche, A. C. Viana, A. L. Aquino, and A. A. Loureiro, “Traffic aware video dissemination over vehicular ad hoc networks,” in *International conference on Modeling, analysis & simulation of wireless and mobile systems*, pp. 419–426, ACM, 2013.

[28] L. A. Villas, A. Boukerche, R. B. Araujo, A. A. Loureiro, and J. Ueyama, “Network partition-aware geographical data dissemination,” in *International Conference on Communications*, pp. 1439–1443, IEEE, 2013.

[29] Y. Sung and M. Lee, “Light-Weight Reliable Broadcast Message Delivery for Vehicular Ad-Hoc Networks,” in *Vehicular Technology Conference (VTC Spring)*, (Yokohama), pp. 1–6, IEEE, May 2012.

[30] P. Ruiz, B. Dorronsoro, P. Bouvry, and L. Tardón, “Information dissemination in VANETs based upon a tree topology,” *Ad Hoc Networks*, vol. 10, pp. 111–127, Jan. 2012.

[31] J. Whitbeck, Y. Lopez, J. Leguay, V. Conan, and M. D. de Amorim, “Push-and-track: Saving infrastructure bandwidth through opportunistic forwarding,” *Pervasive and Mobile Computing*, vol. 8, pp. 682–697, Oct. 2012.

[32] G. Resta and P. Santi, “On the Fundamental Limits of Broadcasting in Wireless Mobile Networks,” in *INFOCOM*, (San Diego, CA), pp. 1–5, IEEE, Mar. 2010.

[33] G. Zyba, G. M. Voelker, S. Ioannidis, and C. Diot, “Dissemination in opportunistic mobile ad-hoc networks: The power of the crowd,” in *INFOCOM*, (Shanghai), pp. 1179–1187, IEEE, Apr. 2011.

[34] C. Xia, D. Liang, H. Wang, M. Luo, and W. Lv, “Characterization and modeling in large-scale urban DTNs,” in *IEEE 37th Conference on Local Computer Networks (LCN)*, pp. 352–359, 2012.

[35] S. Ahmed and S. K. Salil, “Characterization of a large-scale delay tolerant network,” in *IEEE 35th Conference on Local Computer Networks (LCN)*, (Denver, CO), pp. 56–63, Oct. 2010.

[36] R. Monteiro, S. Sargento, W. Viriyasitavat, and O. K. Tonguz, “Improving VANET protocols via network science,” in *2012 IEEE Vehicular Networking Conference (VNC)*, pp. 17–24, IEEE, Nov. 2012.

[37] S. Uppoor and M. Fiore, “Insights on metropolitan-scale vehicular mobility from a networking perspective,” in *International Workshop on Hot Topics in Planet-scale Measurement*, (Low Wood Bay, Lake District, UK), pp. 39–44, ACM, 2012.

[38] F. Silva, T. R. M. B. Silva, F. Benevenuto, L. B. Ruiz, and A. A. F. Loureiro, “Improving information dissemination in vehicular networks by selecting appropriate disseminators,” in *International Conference on Advanced Information Networking and Application (AINA)*, (Victoria, Canada), pp. 681–688, IEEE, May 2014.

[39] S. Uppoor and M. Fiore, “Large-scale urban vehicular mobility for networking research,” in *2011 IEEE Vehicular Networking Conference (VNC)*, pp. 62–69, IEEE, Nov. 2011.

[40] L. Cheng, B. Henty, D. Stancil, F. Bai, and P. Mudalige, “Mobile vehicle-to-vehicle narrow-band channel measurement and characterization of the 5.9 ghz dedicated short range communication (dsrc) frequency band,” *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 8, pp. 1501–1516, 2007.

[41] F. Silva, T. Silva, R. Vicente, L. Ruiz, and A. Loureiro, “On the improvement of vehicular macroscopic mobility models,” in *Vehicular Technology Conference (VTC Fall), 2013 IEEE 78th*, pp. 1–5, Sept 2013.

[42] S. Joerer, F. Dressler, and C. Sommer, “Comparing apples and oranges?,” in *ACM International Workshop on Vehicular Inter-networking, Systems and Applications*, (New York, USA), pp. 27–32, 2012.

[43] Y. L. Morgan, “Notes on dsrc & wave standards suite: Its architecture, design, and characteristics,” *Communications Surveys & Tutorials, IEEE*, vol. 12, no. 4, pp. 504–518, 2010.

[44] D. Jiang and L. Delgrossi, “Ieee 802.11p: Towards an international standard for wireless access in vehicular environments,” in *Vehicular Technology Conference, 2008. VTC Spring 2008. IEEE*, pp. 2036–2040, May 2008.

[45] C. Sommer, D. Eckhoff, R. German, and F. Dressler, “A Computationally Inexpensive Empirical Model of IEEE 802.11p Radio Shadowing in Urban Environments,” Technical Report CS-2010-06, University of Erlangen, Dept. of Computer Science, September 2010.

Fabrício A. Silva received his B.S. degree in Computer Science from Federal University of Minas Gerais, Brazil, in 2004. He received his M.Sc. degree in 2006 and is currently a Ph.D. Candidate in the Computer Science Department at the Federal University of Minas Gerais. He worked as a visiting researcher at PARADISE Research Laboratory, University of Ottawa. His research interests are in the areas of mobile and ad hoc networks, particularly vehicular ad hoc networks.

Azzedine Boukerche (FIEEE, FEiC, FCAE,FAAAS) is a full professor and holds a Canada Research Chair position at the University of Ottawa (Ottawa). He is the founding director of the PARADISE Research Laboratory, School of Information Technology and Engineering (SITE), Ottawa. Prior to this, he held a faculty position at the University of North Texas, and he was a senior scientist at the Simulation Sciences Division, Metron Corp., San Diego. He was also employed as a faculty member in the School of Computer Science, McGill University, and taught at the Polytechnic of Montreal. He spent an year at the JPL/NASA-California Institute of Technology, where he contributed to a project centered about the specification and verification of the software used to control interplanetary spacecraft operated by JPL/NASA Laboratory. His current research interests include Wireless Ad Hoc and sensor networks, wireless networks, mobile and pervasive computing, wireless multimedia, QoS service provisioning, performance evaluation and modeling of large-scale distributed systems, distributed computing, large-scale distributed interactive simulation, and parallel discrete-event simulation. He has published several research papers in these areas. He served as a guest editor for the Journal of Parallel and Distributed Computing (special issue for routing for mobile ad hoc, special issue for wireless communication and mobile computing, and special issue for mobile ad hoc networking and computing), ACM/Kluwer Wireless Networks, ACM/Kluwer Mobile Networks Applications, and Journal of Wireless Communication and Mobile Computing. He has been serving as an Associate Editor of ACM Computing Surveys, IEEE Transactions on Parallel and Distributed systems, IEEE Transactions on Vehicular Technology, Elsevier Ad Hoc Networks, Wiley International Journal of Wireless Communication and Mobile Computing, Wiley Security and Communication Network Journal, Elsevier Pervasive and Mobile Computing Journal, IEEE Wireless Communication Magazine, Elsevier Journal of Parallel and Distributed Computing, and SCS Transactions on Simulation. He was the recipient of the Best Research Paper Award at IEEE/ACM PADS 1997, ACM MobiWac 2006, ICC 2008, ICC 2009 and IWCMC 2009, and the recipient of the Third National Award for Telecommunication Software in 1999 for his work on a distributed security systems on mobile phone operations. He has been nominated for the Best Paper Award at the IEEE/ACM PADS 1999 and ACM MSWiM 2001. He is a recipient of an Ontario Early Research Excellence Award (previously known as Premier of Ontario Research Excellence Award), Ontario Distinguished Researcher Award, Glinski Research Excellence Award, IEEE CS Golden Core Award, IEEE Canada Gotlieb Medal Award, IEEE ComSoc Expectional Leadership Award, IEEE TCPP Excpetional Leadership Award. He is a cofounder of the QShine International Conference on Quality of Service for Wireless/Wired Heterogeneous Networks (QShine 2004). He served as the general chair for the Eighth ACM/IEEE Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems, and the Ninth ACM/IEEE Symposium on Distributed Simulation and Real-Time Application (DSRT), the program chair for the ACM Workshop on QoS and Security for Wireless and Mobile Networks, ACM/IFIPS Europar 2002 Conference, IEEE/SCS Annual Simulation Symposium (ANNS 2002), ACM WWW 2002, IEEE MWCN 2002, IEEE/ACM MASCOTS 2002, IEEE Wireless Local Networks WLN 0304; IEEE WMAN 0405, and ACM MSWiM 9899, and a TPC member of numerous IEEE and ACM sponsored conferences. He served as the vice general chair for the Third IEEE Distributed Computing for Sensor Networks (DCOSS) Conference in 2007, as the program cochair for GLOBECOM 20072008 Symposium on Wireless Ad Hoc and Sensor Networks, and for the 14th IEEE ISCC 2009 Symposium on Computer and Communication Symposium, and as the finance chair for ACM Multimedia 2008. He also serves as a Steering Committee chair for the ACM Modeling, Analysis and Simulation for Wireless and Mobile Systems Conference, the ACM Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks, and IEEE/ACM DSRT.

Thais R. M. B. Silva holds a B.Sc., M.Sc., and Ph.D. in Computer Science from Federal University of Minas Gerais (UFMG), Brazil. Currently she is a Professor of Computer Science at Federal University of Viçosa, Brazil. Her research interests include ubiquitous and mobile networks, with emphasis on context-aware computing.

Fabrício Benevenuto is associate professor in the Computer Science Department of Federal University at Minas Gerais (UFMG). He received a Ph.D. in Computer Science in 2010, and his thesis received the most important Brazilian thesis award (from CAPES). During his Ph.D., he held research intern positions at HP Labs (Palo Alto/CA) and MPI-SWS (Max Planck Institute for Software Systems), Germany. His research plans involve data mining and machine learning with the focus on the development of innovative systems. His work on this field has led to a number of important publications and widely cited papers. He has been serving on the TPC of many top conferences, such as ACM KDD, ACM WSDM, WWW, ACM SIGIR, ACM COSN, and ICWSM. In 2013 Fabrício was elected a member of the Brazilian Academy of Science. The innovative aspect of his research has attracted the attention of several specialized blogs and mass media channels such as The New York Times, BBC, The Hunffington Post, MIT Tech Review, and New Scientist.

Linnyer B. Ruiz is a Computer Science Professor at State University of Maringá (UEM), Brazil. She holds a Ph.D. in Computer Science received from Federal University of Minas Gerais, Brazil (2003) and a M.Sc. degree in Electrical Engineering and Industrial Informatics received from Technological University of Paraná (1996). Her main research efforts are devoted to various aspects of embedded systems, wireless sensor networks, ubiquitous computing, ubiquitous healthcare, and vehicular networks. She has published several papers in the area of wireless ad hoc and sensor networks.

Antonio A. F. Loureiro received his B.Sc. and M.Sc. degrees in Computer Science from the Federal University of Minas Gerais (UFMG), Brazil, and the Ph.D. degree in Computer Science from the University of British Columbia, Canada. Currently, he is a full professor of Computer Science at UFMG, where he leads the research group in mobile ad hoc networks. His main research areas are mobile computing, vehicular networks, wireless sensor networks, and distributed algorithms. In the last 15 years he has published regularly in international conferences and journals related to those areas, and also presented keynotes and tutorials at international conferences.