

1
2
3
4
5
6
7
8
9 Effect of content popularity, number of contents and a
10 cellular backup network on the performance of content
11 distribution protocols in urban VANET scenarios
12
13
14

15 Manuel Uruña^{a,*}, Ignacio Soto^a, Isaias Martinez-Yelmo^b, Maria Calderon^a

16
17 ^a*Departamento de Ingeniería Telemática, Universidad Carlos III de Madrid, Avda.*
18 *Universidad, 30, 28911 Leganes (Madrid), Spain*

19 ^b*Departamento de Automática, Universidad de Alcalá, Campus Universitario N-II, km.*
20 *32, 28871 Alcalá de Henares (Madrid), Spain*
21
22
23
24

25
26 **Abstract**

27
28 Different types of infotainment and traffic efficiency applications make use
29 of content distribution protocols in vehicular networks. Current research on
30 content distribution has been focused on designing efficient protocols, but it
31 has put little effort in modeling the nature of the applications that consume
32 and generate those contents, or in studying the internetworking with cellular
33 networks. This paper studies the effects of application characteristics on
34 the performance of push- and pull-based content distribution protocols in
35 VANETs. In particular, it considers the total number of contents being
36 distributed, the popularity of those contents, and the utilization of a cellular
37 backup network to guarantee a level of service for delay-bounded applications.
38 We also propose the Multi-Hop To Infostation (MHTI) protocol: a pull-
39 based, multi-hop protocol that sends content requests towards the closest
40 infostation. Requests can be satisfied before reaching the infostation by any
41 vehicle in the path that has cached the requested content. Our analysis
42 indicates that the performance of push- and pull-based protocols is only
43 satisfactory in scenarios with a low number of contents or highly popular
44 ones, while MHTI also exhibits a good performance with a large number of
45 contents, and it takes advantage from different content popularities to obtain
46
47
48
49
50
51

52 *Corresponding author

53 *Email addresses:* muruenya@it.uc3m.es (Manuel Uruña), isoto@it.uc3m.es
54 (Ignacio Soto), isaias.martinezy@uah.es (Isaias Martinez-Yelmo), maria@it.uc3m.es
55 (Maria Calderon)
56
57

1
2
3
4
5
6
7
8
9 a high number of contents through the VANET.

10
11 *Keywords:* Vehicular ad hoc network (VANET), Cellular offload, Content
12 distribution protocol, Content popularity
13

14 15 16 **1. Introduction**

17
18 Research on Vehicular Ad hoc Networks (VANETs) was initially driven by
19 safety applications, but there is an increasing interest in also using VANETs
20 for traffic efficiency and infotainment applications by taking advantage of
21 connectivity with infrastructure networks (i.e., Internet) [1]. In many of
22 these applications, specially in urban environments, the VANET is used for
23 the dissemination of contents that can be of interest for some or all vehicles
24 in the VANET. In these applications, contents will typically be relevant in a
25 certain geographical area and during a certain period of time (i.e., contents
26 have an associated lifetime) [2, 3]. Therefore, although users typically may
27 accept some delay in getting a content, they will also have a deadline from
28 which the content is not useful for them anymore.
29

30
31 There are many different types of traffic efficiency [4, 5] and infotainment
32 applications [6, 7, 8] that can be based on content distribution. For
33 instance, they may involve distributing, in a certain geographical area, cur-
34 rent traffic information or points of interest with dynamic characteristics,
35 such as prices in nearby gas stations, parking information, cultural and sport
36 events, or advertisements with offers from shops or restaurants in the area.
37 Each application type imposes different requirements to content distribution
38 mechanisms. Advertisement applications [6, 8] are focused on reaching as
39 many users as possible, although delay is not a critical factor. On the other
40 hand, a user looking for the cheapest gas station is expecting to get such in-
41 formation in time to be useful. Moreover, probably all vehicles are interested
42 in traffic information updates, but only few of them are looking for modern
43 art exhibitions.
44

45
46 However, we argue that in realistic deployments it makes little sense that
47 each application implements its own dissemination mechanism. Instead, a
48 better approach would be employing a common content distribution frame-
49 work that handles multiple contents with different popularity. Also, as the
50 VANET performance depends on vehicle density, and consequently it can-
51 not be guaranteed (e.g., vehicle traffic at 4:00 AM will be sparse), in order
52 to offer predictability, this content delivery system will need a wireless wide
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 area network, such as a cellular network, as a backup to the VANET. This
10 allows that, if the content cannot be obtained through the VANET in the
11 expected time, it can be requested using the cellular network [9]. We could
12 think in offering the service to vehicles just using the cellular network, but
13 experience has shown that having alternatives to avoid overloading the cel-
14 lular network is critical to achieve a cost-effective solution to offer wireless
15 access [10]. This will be even more so with the growth in number of users
16 expected in 5G networks (e.g., with the integration of Internet of Things ap-
17 plications). Therefore, the combination of vehicular networks based on short
18 range communications and cellular networks to offer content delivery services
19 to vehicles is a good approach from the point of view of efficiency in resource
20 usage. With high vehicle densities (without network congestion) the VANET
21 performs better and most contents can be obtained without relying on the
22 cellular network, whereas a low vehicle density also means a low load for the
23 cellular network. Thus, we can save the expensive and scarce resources of the
24 cellular network when there is more demand. The use of the cellular network
25 also has a seeding effect in the VANET, by introducing new contents in it,
26 which has a significant impact on its performance that cannot be ignored
27 when studying protocols for VANETs.

28
29
30
31
32
33 Several protocols have been proposed in the literature for non-safety data
34 dissemination in VANETs [5, 6, 8, 11, 12]. The merits of those protocols are
35 studied by means of performance parameters such as overhead, percentage of
36 cars that receive a content, or average content distribution delay. However,
37 less effort has been put in studying and modeling the nature of the applica-
38 tions that consume and generate those contents. We can identify two main
39 families of protocols for content distribution in VANETs: push and pull pro-
40 tocols. In this paper we show that the relative performance of each family of
41 protocols is heavily influenced by the number of contents to be distributed
42 and their popularity, which depends on the nature of the content distribution
43 application. In fact, we show that a family of protocols that can achieve very
44 good performance for some applications, can perform poorly in other cir-
45 cumstances. In order to deal with these challenges, we propose a pull-based,
46 multi-hop protocol, called Multi-Hop To Infostation (MHTI), that forwards
47 content requests towards the closest infostation (a node connected to the
48 fixed infrastructure), but that can be also replied by any node in the path
49 that has the requested content in its cache. This kind of multi-hop protocol
50 is not common in the literature of content distribution in VANETs, because
51 the assumption is that other kinds of simpler, one-hop protocols can pro-
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

vide enough performance. We show in this paper that, although this is true for certain content distribution applications with few contents, this is false in many other situations, where we require a multi-hop protocol to achieve better performance. This improvement in performance translates to significant savings in cellular network usage. Our work has been validated using a simulation environment based on OMNET++ [13] that includes a real urban map, buildings that affect signal propagation, and a faithful micro vehicle traffic management representation thanks to the use of SUMO [14, 15]. All the tools and datasets that we have employed, such as the implementation of the protocols, the map with obstacles, and the vehicle traces are publicly available¹ to foster future developments in this research field.

The rest of the article is organized as follows: in Section 2 we review the key protocols in the literature for content distribution in VANETs, also identifying the conditions in which the evaluation of the protocols was done; in Section 3 we describe three different reference protocols for content distribution, classified in different families, including the proposed MHTI protocol; in Section 4 we introduce our simulation environment; in Section 5 we present the quantitative results of our experiments with the three families of content distribution protocols in VANETs; and finally, in Section 6, we summarize the conclusions of our work.

2. State of the art in VANET content distribution protocols

In the literature there are many protocols proposed to support multi-hop dissemination of contents, but most of them are tailored for safety-related applications, which require a strict time-constrained data dissemination. These protocols mainly concentrate on how to choose the next forwarder to efficiently deliver a safety-related content (e.g., a warning of a collision ahead) to all the vehicles in an area, while optimizing the dissemination delay (i.e., intelligent flooding schemes). Note that this safety-related scenario is quite different from our target scenario. Thus, in the related work that follows, we leave out the dissemination schemes tailored for safety-related applications.

Our work is focused on scenarios where traffic efficiency and infotainment applications use the VANET to disseminate non-safety contents that can be of interest for some or all vehicles. These contents have an associated lifetime,

¹<http://www.it.uc3m.es/muruenya/research/MHTI.tar.gz>

1
2
3
4
5
6
7
8
9 but users typically may accept some delay in getting a content. In general,
10 contents can be obtained from infostations located at fixed points of the
11 infrastructure, or from other vehicles that already have the content. In the
12 literature there are several protocols proposed to cope with this non-safety
13 data dissemination scenario, which we review below. The usual approach is to
14 make use of caching to improve the efficiency of dissemination. Thus, content
15 replicas are stored in vehicles for its future distribution to other vehicles.
16 This store-carry-and-forward data dissemination model is considered valid
17 for both infotainment and traffic efficiency applications.
18

19
20 Two main families of dissemination protocols can be identified in the
21 literature: push- and pull-based. In the push model, data are periodically
22 broadcast to neighbors. In this way, a content is broadcast to potentially
23 interested users by an infostation, or by any another vehicle that has previ-
24 ously obtained the content and has it in its cache. In the pull model, data
25 are delivered on demand, i.e., a content is explicitly requested by a vehi-
26 cle. Therefore, the interested vehicle first sends a content request, and it
27 eventually receives the content if such request gets to some node that has
28 the content, either an infostation or another vehicle that has the required
29 content in its cache.
30

31
32 As it was previously mentioned, we argue that some key characteristics
33 of the applications should not be ignored. Namely, the total number of con-
34 tents, their popularity, and the utilization of a cellular network as a backup
35 mechanism to guarantee a bounded delay. Thus, we review the main dis-
36 semination protocols for vehicular networks, paying special attention to the
37 conditions in which the performance of these protocols has been evaluated.
38

39
40 A significant number of dissemination protocols belong to the push-based
41 family [2, 3, 4, 5, 6, 11, 16]. Among them, an interesting subset of these
42 works [2, 6] focus on the concept of *abiding geocast*, analyzing different mech-
43 anisms to guarantee that data keep being distributed in an area for a certain
44 period of time (i.e., a content is interesting for some time and so it is stic-
45 ked to the area during this time). However, these protocols [2, 6] have been
46 evaluated considering that only 1 content has to be distributed in the area,
47 which is a simplistic assumption in realistic scenarios. On the contrary, an-
48 other subset of these works [4, 5] evaluate the proposed protocols considering
49 as many contents as vehicles exist in the scenario. In particular, in [4] a rel-
50 evance function is used to determine what contents should be cached, while
51 [5] focuses on analyzing the impact of having an adaptive timer, instead
52 of a periodic one, to broadcast contents to neighbors. In the case of [16],
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

apart from the neighbors that periodically broadcast contents, the existence of some standalone infostations, located at intersections, which buffer and rebroadcast contents, is considered as well. Interestingly, the number of contents considered in the evaluation of this protocol ranges from 10 to 600, in a simulation setup with 210 vehicles, and the experiments show that the number of contents has a strong impact on the performance of push protocols, result that is consistent with our work. However, our study covers the different families of content distribution protocols (push, pull and pull multi-hop), extends the number of contents up to 100,000 finding an asymptotic behavior, and considers not only the number of contents but also their popularity.

A second group of papers has explored the approach of pull-based protocols, where contents are explicitly requested by vehicles. A first possibility is that this interest is spread one-hop away, as it is proposed in [12], where vehicles periodically disseminate their interests to their one-hop neighbors, and eventually the content is received by the interested node if this interest reaches a neighbor (vehicle or infostation) that has it in its cache. However, this one-hop pull protocol was evaluated assuming the existence of a single content to be disseminated in the area. A second possibility, considered in [17, 18], is that the requests are forwarded several hops away from the requesting vehicle. However, these interesting works are only partially relevant to ours, since they mainly deals with applying the paradigm of content-centric networking to vehicular networks. Although [17] is a pioneering work in taking into account the number of contents (4 and 32 contents are considered in the evaluation) and their popularity (two different popularity values are considered), it does not follow a systematic approach to analyze the impact of these aspects on the different dissemination protocols, as our work does, nor considers a cellular backup network.

A recent paper [8] proposes Starfish, a protocol that combines a push-based approach with the infostation pushing contents (100-KB files) using a round robin strategy, and a one-hop pull approach with vehicles asking for the missing contents. Contrary to our work, Starfish does not consider delay-bounded applications (the same contents are delivered for several days), and popularity is not taken into account. Starfish has been evaluated through real experiments using a roadside infostation and several vehicles moving around, which is very valuable but limits the evaluation scenarios.

Moreover, none of the aforementioned works considers the internetworking of cellular networks and short-range communications (i.e., the VANET).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Recently, there exist some works that advocate about their combination in vehicular communications [9, 19, 20, 21, 22, 23, 24], although they do not deal specifically with the content distribution problem in store-carry-and-forward networks. In particular, in [22] the authors propose the use of a 3G cellular network to gather and distribute relevant control information (e.g., connectivity in the different roads) with the goal of improving the routing process over the VANET. Similarly, in [24] LTE is used to organize the vehicular network into clusters and to maintain this structure over the time. On the contrary, in [23] the cellular network is used to exchange data when there is not a multi-hop VANET path between the source and destination vehicles. In [9, 21] the VANET is used to offload 3G cellular network in vehicle-to-Internet scenarios by delivering data originally targeted to cellular networks via WiFi, with the goal of alleviating the congestion of cellular networks. Some recent studies [19, 20] focus on content distribution with peer-to-peer (P2P) cooperation. In common with these works, we also propose the combination of cellular networks and short-range communications. However, while these works use the store-and-forward communication paradigm, in our research the vehicular network applies a store-carry-and-forward approach for content distribution. Therefore, these works can be considered orthogonal to ours.

Closer to our work, there is a recent paper [25] that considers the hybrid use of short-range and cellular radios to optimize the dissemination of delay-tolerant contents using a store-carry-and-forward approach, although the cellular network is only used for seeding, not for guaranteeing a bounded delay. This work analyzes mathematically the push-based content dissemination process, formulating it as an optimization problem from the perspective of the content provider. It is worth mentioning that its analytical model considers the effect of interest popularity, concluding that this parameter has an important impact on the performance of push content distribution protocols, result that is consistent with our work. However, we provide experimental results by means of simulations, and we cover the different families of content distribution protocols (push, pull and pull multi-hop) finding that popularity affects differently each family.

Table 1 summarizes the conditions in which the evaluation of the aforementioned content distribution protocols has been performed (for works based on simulation), including MHTI, the protocol that we propose in this paper. Namely, the total number of contents, their popularity, and the utilization of cellular networks as a backup mechanism to guarantee a bounded

Table 1: Analysis of VANET content distribution protocols

Protocol	Type	# Contents	Popularity	Cellular network
[2]	Push/Pull	1	No	No
[3]	Push	40	No	No
[4]	Push	N^*	No	No
[5]	Push	$\sim N^*$	No	No
[6]	Push	1	No	No
[11]	Push	N^* per sec	No	No
[12]	Pull	1	% interested cars	No
[16]	Push	[10-600]	No	No
[18]	Pull	1	No	No
[17]	Pull	{4, 32}	Zipf $\alpha=\{0.0, 0.8\}$	No
MHTI	Pull	[1-100,000]	Zipf $\alpha=[0.0-2.0]$	Yes

* N is the number of vehicles

delay. As we have seen, many works only consider one or few contents (e.g. < 50), or assume that all vehicles are equally interested in all contents. Moreover, none of the reviewed works include a cellular backup network in their experimental simulations. From this analysis we came to the conclusion that more research was needed to understand what is the impact of application characteristics and a cellular backup on the performance of VANET content distribution protocols.

3. Protocols for content distribution in urban VANET scenarios

As described in the previous section, we can identify two main families of VANET content distribution protocols: push- and pull-based. Furthermore, the pull protocols can follow a one-hop approach asking the contents only to neighbors, or a multi-hop approach in which contents can be obtained from nodes farther away through several hops.

Since we are not interested in studying a particular protocol proposal but we are trying to study the effect of applications characteristics and the use of a cellular backup network in general, we have implemented three generic protocols to represent these three different families of content distribution protocols, including the proposed Multi-Hop To Infostation (MHTI) mechanism that follows the philosophy of sending multi-hop requests towards the closest infostation. Each protocol also has a number of optimizations that

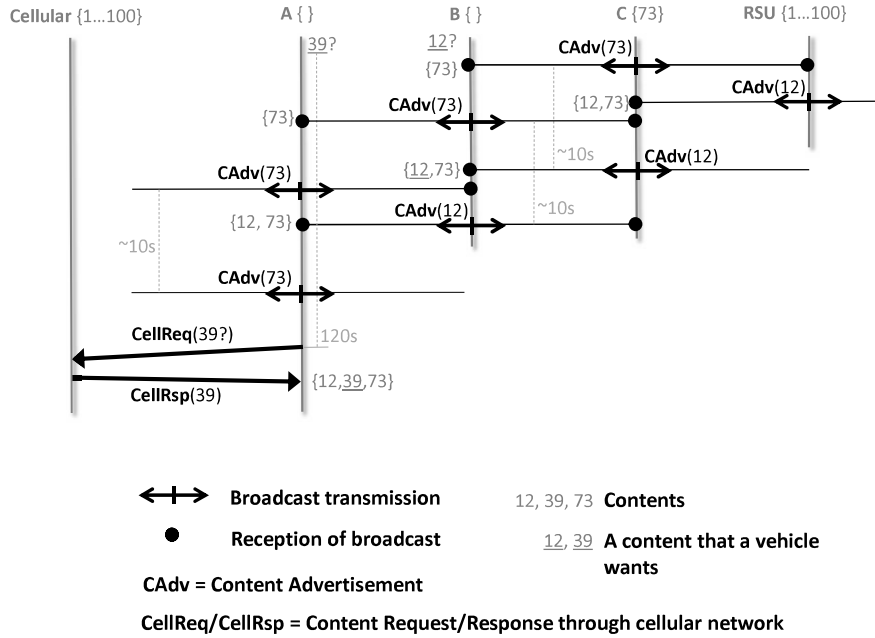


Figure 1: The push protocol

are enabled or disabled in order to compare their effects.

In the so-called push protocol (“Push”), every ten seconds² each node (car or infostation) chooses a random content from its cache and broadcast it to all its one-hop neighbors (see Figure 1). Whenever an advertised content is received, the vehicle stores it in the local cache (for simplicity we assume the cache is able to store all contents, thus no replacement policy is necessary), and then checks whether the local application is interested in this new content in order to return it immediately. Figure 1 also shows that when a vehicle fails to get a content through the VANET, after some timeout that depends on application requirements (120s in the figure), the vehicle will request the content through a cellular network. In the rest of the protocols that we describe in this section, we also use the cellular network as a backup when a content cannot be obtained through the VANET.

²A 10s timer allows the vehicle to move 139m at 50 km/h, and thus reach other neighbors in next transmission (assuming a transmission range of 200m).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52

In the optimized push protocol (“Push-opt”), the next advertised content is not chosen randomly, but following a least recently used (LRU) policy. For instance to avoid a group of cars advertising the same content over and over again. Each car keeps for each cached content a timestamp that is updated whenever the content is either advertised or listened from other cars. Then, whenever a content has to be advertised, the one with the oldest timestamp is chosen (or a random one if multiple contents have exactly the same timestamp). The contents just downloaded through the cellular network have a zero timestamp, so they are advertised next.

Similarly, in the one-hop pull protocol (“Pull”), every ten seconds each car that has not obtained its content yet, and has not given up (due to an application timeout), broadcasts a request for a content to all its direct neighbors (see Figure 2). If any of them has the requested content in its cache, it is replied back in unicast (with a random 100ms jitter³ to avoid collisions), otherwise no response is returned. When the application has obtained its content or has given up, the car stops sending requests, although it may keep replying to other cars’ requests.

In the optimized one-hop pull protocol (“Pull-opt”) the content responses are sent in broadcast to one-hop neighbors, so the requested content can be overheard by nearby cars, which also store it in their caches. Moreover, when a content is downloaded through the cellular network, it is also immediately advertised to one-hop neighbors, like in the Push protocol. This is done because, since it has been downloaded through the cellular network, this means that neighbor cars do not have such content (otherwise they would have responded to the request).

The third protocol is the proposed Multi-Hop To Infostation (“MHTI”) one. Similarly to previous protocols, cars that have not obtained a content they want, and have not given up (due to an application timeout), will try to get the content every 10 seconds. In MHTI, to get a content, content requests are sent towards the nearest infostation employing a simple geographical routing protocol based on greedy forwarding [26]. The operation on each hop is similar to the one-hop Pull protocol: content requests are first broadcast to one-hop neighbors (see Figure 3), but now the neighbors that do not have the

³A 100ms jitter prevents packet collisions even in high density scenarios, whereas its effect on vehicular mobility is negligible: in 100ms a car moves less than 2 meters at 50Km/h.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

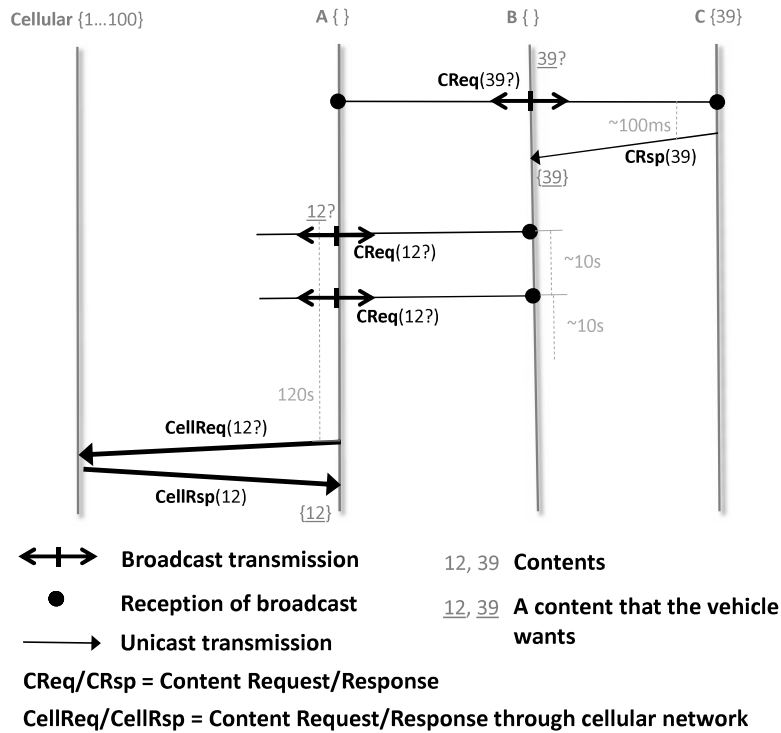


Figure 2: The pull protocol

content, but are nearer to the target infostation than the requesting car, reply with their addresses and coordinates (in a so-called Neighbor message). Thus, if the car hasn't received the content after 100ms, it forwards the content request message to the closest neighbor to the target infostation (i.e. we assume cars know the coordinates of all infostations in the area⁴ so they can target the nearest one). This process repeats until the infostation or another car replies with the requested content. The content reply follows the same multi-hop path (which has been recorded in the content request message) in reverse, and it is also stored in the content cache of all intermediate nodes. Although the movement of vehicles can potentially affect the stability of multi-hop paths, in practice the speed of vehicles (at most 50 Km/h or 13.89

⁴The location of all infostations in the area could be included in the digital maps of GPS navigators, or be downloaded through the cellular network.

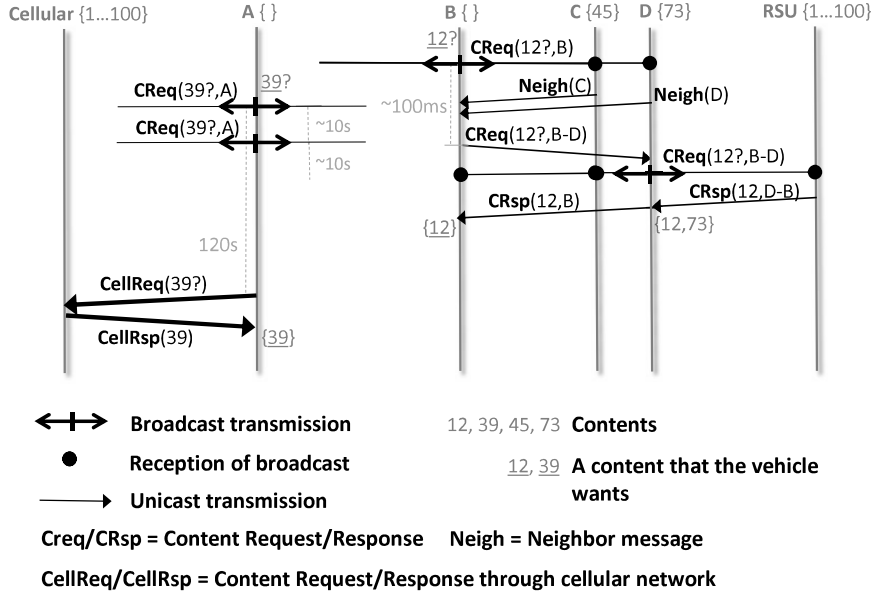


Figure 3: The MHTI protocol

m/s in an urban environment) means a vehicle moves not more than just a few meters in the time needed to send a message (a request) and receive a reply (a content). Therefore, the multi-hop path is, in general, stable during a MHTI message exchange.

The optimized version of the MHTI protocol (“MHTI-opt”), as the optimized one-hop Pull-opt one, also features broadcast responses and advertises contents downloaded through the cellular network to one-hop neighbors. Moreover, in order to reduce the number of Neighbor messages, they are only sent if there has not been a content response yet. Finally, the potential Neighbor responses are sent according to their distance to the requesting car. This way, cars far away (max. 175m) reply in broadcast with their coordinates first, and other neighbors only send theirs if they are closer to the target infostation than previous ones.

4. Simulation scenario

For our simulation study, we need a representative urban scenario where different traffic conditions can be defined. We also need a flexible tool that

Table 2: Parameter settings in the simulations

Parameter	Value
Roadmap [size]	Madrid [4.5x4.5 km^2]
Number of vehicles	[2,000-9,000]
Number of infostations	4
Simulation time	3600 s.
Max speed	50 km/h
MAC/PHY	IEEE 802.11b/g 2.4 GHz
WiFi channel rate	2 Mbps
WiFi transmission range	200 m
Shadow fading model	[28]
Cellular channel bandwidth	100 Mbps per car
Content size	1,000 bytes
No. of contents	[1-100,000]
Contents' popularity	Zipf $\alpha=[0.0-2.0]$

is able to model realistic communications and evaluate the desired protocols properly. Hence, the Veins 2.0 platform [27] was selected to conduct our simulation study. The Veins (Vehicles in Network Simulation) platform integrates the OMNeT++ [13] event-based network simulator with the SUMO (Simulation of Urban MObility) road traffic simulator [14, 15]. OMNeT++ allows the implementation and simulation of the protocols under study, using the MiXiM toolkit (that is integrated in Veins) for modeling the physical and MAC layers (IEEE 802.11b/g) of the network communications. Furthermore, Veins includes additional enhancements, such as a shadow fading model for buildings in urban and suburban scenarios [28]. SUMO provides an accurate modeling of microscopic vehicular traffic, simulating the behavior of drivers⁵. The integration of OMNeT++ and SUMO is made through the Traffic Control Interface (TraCI) provided by SUMO.

Table 2 shows the main parameters used for the simulations, which are discussed in next subsections.

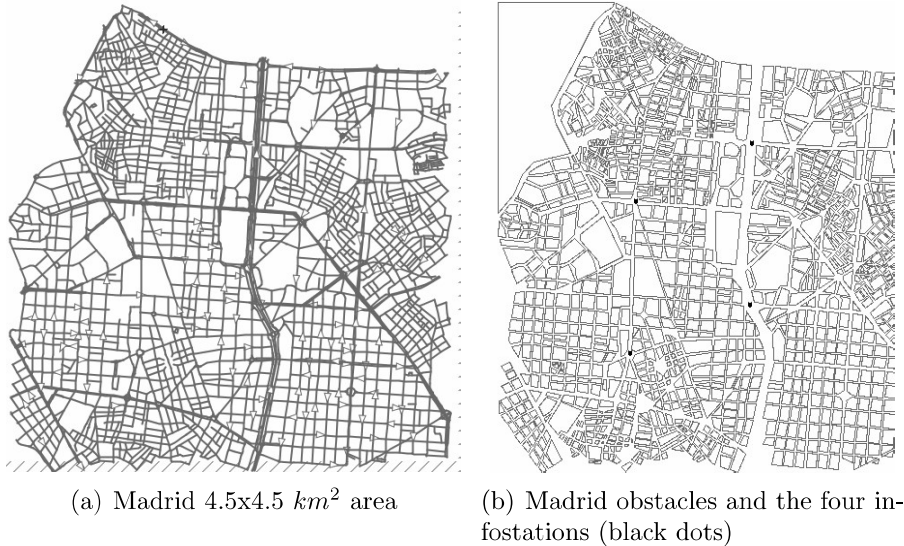


Figure 4: Madrid map and its associated obstacles and infostations locations

4.1. Urban area map

In order to conduct our experiments, we need a representative urban area. We have selected an area of $4.5 \times 4.5 \text{ km}^2$ from Madrid city center. The streets and traffic data of this urban area were obtained from OpenStreetMap (OSM) [30], using the Java OpenStreetMap editor (JOSM) [31] in order to extract the desired area and verify its junctions and connections. The maximum speed of all streets has been set to 50 km/h. Then, the final map was converted to SUMO XML format. The urban area map is shown in Figure 4(a). In addition, we also consider the effects of shadowing due to buildings in wireless data transmissions. We achieve this using the additional models provided by Veins to the OMNeT++/MiXiM framework [28]. However, we first needed to create the obstacles associated to the buildings in the map. We have generated those obstacles manually, with a custom developed tool, by considering each block of buildings like a single obstacle in order to accelerate the simulations. This procedure generates a worst case scenario, since it leads to greater propagation losses when information is transmitted.

⁵In our simulations we use the default models and parameters in SUMO. In particular, the car-following model is the SUMO extended Krauss model [15], and the lane changing model is the one described in [29].

1
2
3
4
5
6
7
8
9 The considered obstacles are shown in Figure 4(b). In our experiments we
10 use four infostations (except where specified otherwise). Their locations are
11 shown in Figure 4(b). The infostations are placed at key map locations,
12 such as the intersections of main avenues. We chose the number of infosta-
13 tions, which is a relatively small scale deployment for the map area, because
14 we want to study the benefits of using the VANET for content distribution
15 without requiring an expensive infrastructure deployment.
16
17

18 19 *4.2. Traffic conditions*

20
21 After defining the map, it is necessary to create an input dataset con-
22 taining different traffic conditions. A real-life, trace-based dataset could not
23 cover all desired conditions, and we would be limited to use the specific urban
24 area of the available traffic dataset. Therefore, we have generated a synthetic
25 traffic dataset.
26

27 In particular, we consider between 2,000 and 9,000 cars making trips in
28 our urban map. The insertion time of new cars is uniformly distributed along
29 the one-hour simulation time. The different trips of the cars are generated
30 between two random points of the map, although we require a minimum
31 distance per trip of 1 km. The cars leave the simulation as soon as they reach
32 their destinations. The assignment of routes to trips is not a trivial issue,
33 because it depends on the distance, user preferences and traffic conditions.
34 Therefore, we use the Dynamic User Assignment (DUA) algorithm [32, 33] to
35 assign a route to each trip, which uses the A-star algorithm [34] to calculate
36 the relative distances in an iterative process that we limit to 30 repetitions.
37

38
39 As result of the above process, we generate five different SUMO XML
40 route files for each scenario with a different total number of cars. Then,
41 a one-hour simulation (i.e. 3,600s) is run for each scenario with the five
42 different generated routes in order to calculate 95% confidence intervals for
43 all experimental results.
44
45

46 Table 3 shows the average trip distance, stay time (i.e., trip duration)
47 and speed of all cars in the five simulations of each scenario, where N is the
48 average number of cars in the scenario. Notice that the average speed of
49 cars decreases, and the average stay time increases, with traffic density, as
50 expected. The average trip distance decreases slightly because, as the time
51 cars need to complete their trips increases, there are more cars at the end of
52 the simulation that have not finished their trips.
53
54
55
56
57
58
59
60
61
62
63
64
65

Table 3: Traffic characteristics of the different scenarios

N	Distance (m)	Stay time (s)	Speed (km/h)
2,013	3,223.21	324.54	35.75
3,012	3,221.97	326.96	35.48
4,011	3,227.02	332.83	34.90
5,035	3,229.56	339.91	34.20
6,082	3,189.14	369.39	31.08
7,087	3,183.98	396.00	28.95
8,165	3,121.01	442.50	25.39
9,220	3,089.42	474.15	23.46

4.3. Network capabilities of vehicles

We consider that all vehicles have a dedicated short-range wireless communication system for non-emergency applications, employing the IEEE 802.11b/g standard. We think that using general purpose wireless technologies is appropriate for a content distribution system as the one we propose, since this approach enables the usage of general purpose devices for participating in the system, whereas specialized wireless technologies, such as IEEE 802.11p, will probably be reserved for safety applications. The transmission rate for the chosen IEEE 80.11b/g technology is fixed to 2 Mbps in order to maximize the communication distance among vehicles (around 200m without obstacles) and robustness (communications are mostly broadcast transmissions). Higher rates would allow having more bandwidth available in the VANET without affecting the conclusions of our work, as we are not exploring saturation scenarios in the VANET. In addition, all vehicles also have cellular communications capabilities. In particular, we simulate the use of a 5G-like cellular network with an available bidirectional rate of 100 Mbps per car, which is consistent with what we can expect in future 5G networks. Nevertheless, we prioritize the use of the VANET, offloading as much load as possible from the cellular network because, even with the increased bandwidth in 5G, it will still be an expensive and scarce resource due to the growth in users and bandwidth requirements (video traffic, Internet of Things, connected vehicles, etc.) [10].

4.4. Number and popularity of contents

To understand the impact of application characteristics in content distribution protocols, we have performed a thorough simulation study where we

1
2
3
4
5
6
7
8
9 analyze the impact of the total number of contents (M). We vary M from 1
10 to 100,000, as well as the popularity distribution of those contents, by using a
11 Zipf distribution with an α exponent ranging from 0.0 (i.e. uniform distribu-
12 tion) to 2.0 (i.e. few highly popular contents). Therefore, at the beginning of
13 the simulation, M contents are generated and a Zipf popularity distribution
14 is computed with the given α exponent. The popularity of each content is
15 the probability of being chosen by a vehicle as its content of interest (each
16 car is interested in a single content during its trip). For simplicity, we assume
17 that all contents have a fixed size of 1,000 bytes, so they fit in a single IEEE
18 802.11 frame.
19

20
21
22 The bounded delay for obtaining the content of interest is implemented
23 with a 120s timer, which starts as soon as each car starts its trip. If the
24 content is not obtained before its timeout, the content is downloaded by 5G
25 and stored in the cache. This timer represents the maximum delay applica-
26 tions can wait for the contents. A too small value gives few opportunities to
27 get contents through the VANET, while a large value might be unacceptable
28 for many applications. Note that the 120 s value is roughly a third of the
29 average stay time of cars in the studied scenario, so that this is the time
30 the vehicles have to obtain the content through the VANET. The rest of the
31 time, vehicles just keep distributing their cached contents to other vehicles.
32 It makes sense to relate stay time and maximum delay, as it is reasonable
33 that if a vehicle is in an area for some time, the contents, which are typically
34 related to the geographical area, must be obtained in time to be useful.
35
36
37
38
39

40 **5. Performance of VANET content distribution protocols**

41 *5.1. Effect of car density*

42
43 In our first scenario we have just one content to be distributed. The
44 success ratio is measured as the average number of contents that cars obtain
45 through the VANET or, since each car is interested in a single content, it can
46 also be seen as the percentage of vehicles that receive their desired content
47 through the VANET. Notice that the contents not obtained through the
48 VANET after 120s will still be downloaded through the 5G network and
49 stored in the local cache.
50
51

52 Figure 5 shows the VANET success ratio for different vehicle densities
53 for the three different content distribution protocols (Push, Pull and MHTI),
54
55
56
57
58




Figure 5: Distribution of one content with different vehicle densities ($M = 1$ content)

including its optimized versions (Push-opt, Pull-opt and MHTI-opt)⁶. Since there is no data congestion, success ratio improves with density as expected, albeit very little given the high success ratio (88%) even at the lowest density. The difference between the three protocols is thus quite small, and it does not depend on vehicle density nor on optimizations. Therefore, with a single content, even the simplest protocols are able to obtain almost a 100% VANET delivery ratio (it does not reach 100% because the simulation stops with vehicles that have just arrived to the scenario and thus they have not been able to get the content yet).

The effect of optimizations can be only seen in Table 4, which shows the relative overhead of each protocol. This relative overhead value is computed as the ratio between the number of bytes transmitted by a car, including

⁶Note that when we have just one content ($M=1$), there is not difference between the Push and the Push-opt protocols.

Table 4: Relative overhead of distributing one content with different vehicle densities ($M = 1$ content, N vehicles)

N	Push	Push-opt	Pull	Pull-opt	MHTI	MHTI-opt
2,013	31.89	31.85	1.70	1.40	1.82	1.55
3,012	31.29	31.29	1.95	1.40	2.04	1.49
4,011	31.71	31.68	2.25	1.45	2.33	1.50
5,035	32.46	32.45	2.61	1.50	2.68	1.55
6,082	35.48	35.46	3.53	1.61	3.57	1.63
7,087	38.29	38.29	4.04	1.67	4.08	1.66
8,165	43.15	43.14	5.12	1.80	5.14	1.77
9,220	46.44	46.44	5.94	1.89	5.96	1.85

both the payload and headers as well as any control messages (e.g., ACKs), divided by the number of contents obtained through the VANET, times 1,000 bytes (i.e. the size of each content). Therefore, the relative overhead value roughly defines the average number of times cars have to transmit a content in order to deliver it to an interested vehicle. A relative overhead of 1.0 would be the optimum one, because it would mean only one transmission, with only the content bytes, has been necessary to deliver the content. Clearly Push protocols that periodically transmit the cached content have a much higher relative overhead than Pull ones. The overhead of Push protocols increases with density because the more cars, the earlier the content is received, and thus it is transmitted more times. Pull and MHTI protocols have the same performance and overhead because, in this scenario, the MHTI protocol usually gets the content from the first hop (i.e. as the one-hop Pull protocol). The optimized versions of both pull-based protocols greatly reduce the overhead, because they prevent several neighbors from responding to the same request with the same content.

Therefore, if there is only one content to be distributed, which is a common scenario evaluated by VANET researchers (see Table 1), any content delivery protocol provides a good performance, albeit Push protocols have a larger overhead than Pull ones (especially with the response filtering optimization).

5.2. Effect of the total number of contents

In our second scenario we explore the impact of the total number of contents. Each vehicle is still interested in one content, but now there are




Figure 6: Distribution of multiple contents with uniform popularity ($N = 6,000$ cars)

multiple contents being distributed in the VANET, and thus different vehicles are interested in different contents. Contents have a uniform distributed popularity (i.e. all contents are equally popular). From the previous results we have chosen a middle vehicle density of $N = 6,000$ cars, so all protocols can achieve a reasonable performance from the density point of view. Figure 6 shows the VANET success ratio, i.e. the percentage of vehicles that obtain the content they are interested in through the VANET, for different number of contents ($M = 1$ to 100,000 contents).

As it can be seen in Figure 6, the VANET success ratio decreases abruptly with the number of contents. Content distribution protocols cannot cope with a large number of contents, and even having 10 or 100 contents results in a low number of contents obtained through the VANET. The Push protocol seems to achieve a better performance than the Pull one with fewer contents but, as the total number of contents increases, its performance drops almost to zero, while the Pull protocol keeps a baseline VANET success ratio of

Table 5: Relative overhead of distributing multiple contents with uniform popularity ($N = 6,000$ cars, M contents)

M	Push	Push-opt	Pull	Pull-opt	MHTI	MHTI-opt
1	35.48	35.46	3.53	1.61	3.57	1.63
10	45.01	41.56	1.52	1.55	3.53	2.24
100	167.99	133.91	1.53	2.94	8.09	5.27
1,000	1,427.24	1,100.61	1.66	5.69	11.50	9.91
10,000	18,397.56	12,637.97	1.69	6.64	12.16	11.27
100,000	104,161.62	91,131.74	1.68	6.72	12.32	11.48

16%. The MHTI protocol is the best performer of the three families across the whole range of contents, although its performance also degrades with the number of contents until it stabilizes in a 35% success ratio.

For a low number of contents, the Push protocol outperforms the Pull one due to a more aggressive dissemination of contents that are cached in the vehicles. This helps to have the small number of contents available when needed. However, when the number of contents increases, the Push protocol is unable to distribute all contents in the available time, and thus the probability of receiving the desired content from a neighbor drops to almost zero. This low performance makes the relative overhead (shown in Table 5) to be extremely high.

The Pull protocol performs worse initially because vehicles only ask for the content they are interested in, so dissemination of contents in the VANET is much slower: a vehicle or infostation only transmits a content if it is asked for it and already has it. On the other hand, when the number of contents is very large and equally probable, it is almost impossible to find a particular content in the cache of a neighbor car. Nevertheless, the Pull protocol can still achieve some success ratio because a vehicle that travels close to a infostation can ask it for the desired content. The 16% success ratio in Figure 6 for Pull protocols with more than 1,000 contents is explained because of the vehicles in the simulation that travel through the coverage area of the 4 infostations. We have validated this hypothesis by performing additional simulations with just two infostations and, as expected, the success ratio is halved. The same reasoning applies to the MHTI protocol, although the asymptotic performance is two times better because more cars are able to reach the closest infostation through intermediate hops. However, these extra hops imply multiplying the relative overhead of the MHTI protocol compared

Table 6: Relative overhead of distributing multiple contents with Zipf popularity ($M = 1,000$ contents, $N = 6,000$ cars)

$Zipf\alpha$	Push	Push-opt	Pull	Pull-opt	MHTI	MHTI-opt
0.00	1,427.24	1,100.61	1.66	5.69	11.51	9.91
0.50	1,034.32	803.16	1.64	5.11	11.07	8.98
1.00	323.05	265.13	1.56	2.76	7.37	5.31
1.50	133.35	106.38	1.89	1.83	4.21	2.87
2.00	95.02	77.70	2.34	1.67	3.51	2.18

to the one-hop Pull one. In any case, with many contents, a greater number of infostations would be beneficial for the performance of Pull and MHTI protocols, but it will not be a significant help for Push ones.

The important conclusion is that the number of contents has a significant effect on the performance of content distribution protocols. In fact, the effect is much larger than the optimizations or the type of protocol itself. Nevertheless, as we have seen in Section 2, little research in this field considers more than a small number of contents.

5.3. Effect of content popularity

In our third scenario we explore the impact of having contents with different popularity, so vehicles will be more likely to be interested in some contents than others. In this scenario we assume that each vehicle is interested in one content out of $M = 1,000$ possible ones. Vehicle density is kept at $N = 6,000$ cars. To model content popularity we use a Zipf distribution, and we analyze different values of its α skewness exponent, ranging from 0.0 (i.e. all contents are equally probable) to 2.0 (i.e. most vehicles are interested in few, highly popular contents).

The VANET success ratios and overheads for this scenario are shown in Figure 7 and Table 6. As it can be seen, the performance of the three families of protocols increases with the popularity skewness of contents, because the more popular a content is, the easier is to find it in neighboring cars. However, for Push protocols this improvement with popularity is limited due to the large number of contents, which makes cars to waste bandwidth broadcasting less popular contents. Nevertheless, Push protocols still benefit from popularity because, although the Push protocol itself treats all contents in a similar way, popular contents are downloaded more times through 5G, and thus they are seeded and advertised more times than unpopular ones.




Figure 7: Distribution of multiple contents with Zipf popularity ($M = 1,000$ contents, $N = 6,000$ cars)

The effect of popularity is much more significant in the case of the Pull and MHTI protocols, going from a low performance and high overhead for uniform popularity ($Zipf \alpha = 0.0$) to a good performance level for highly skewed popularity distributions (e.g., $Zipf \alpha = 2.0$). In these protocols, vehicles ask explicitly for the contents they want, so popular contents are broadly distributed in the VANET and available in the cache of many cars, while little time is wasted in transmitting unpopular contents. Although for Pull protocols this means less popular contents rarely can be fetched through the VANET and are downloaded via 5G, the multi-hop MHTI protocol allows to get these low popularity contents from the nearest infostation in many cases, and thus it is the best performer.

The important conclusion is that popularity also has a great impact on the performance of content distribution protocols in VANETs, being extremely beneficial for the performance of pull protocols, while having little effect in push ones. Therefore, content popularity should not be ignored when




Figure 8: Effect of the 5G network with content popularity ($M = 1,000$ contents, $N = 6,000$ cars)

analyzing VANET protocols, because it may have a significant influence on the performance as compared with other protocols or optimizations.

5.4. Effect of the cellular backup network

Given that VANET content delivery protocols can fail, and that the performance is very dependent on the traffic scenario, the cellular backup network plays a key role in our proposal because it allows offering a guaranteed service: users can always get the contents they need in a bounded time. However, the use of a cellular backup network also affects the performance of the VANET content delivery protocols themselves due to its seeding effect: contents that are not initially available in a particular geographical area can then be provided to that area through the cellular network.

The effect of the cellular backup network for different protocols and different content popularities can be seen in Figure 8. In this figure we compare the results presented in Figure 7 for the protocols with optimizations, with

1
2
3
4
5
6
7
8
9 the results in the same scenarios but without using the 5G backup network.
10 An obvious first effect is that, without the 5G network, the VANET delivery
11 ratio is also the total content delivery ratio, as contents cannot be obtained
12 using the cellular network when the VANET protocol fails. But we can also
13 observe an effect on the percentage of contents that can be obtained through
14 the VANET, which is better when using a cellular backup network than with-
15 out it. This is due to the seeding effect: contents downloaded through the
16 5G network can be later served using the VANET to other vehicles.

17
18 The seeding effect is critical for the improvement of performance with
19 popularity in Push protocols. Push protocols by themselves lack mecha-
20 nisms to benefit from popularity, since cars just advertise all cached contents
21 with the same probability. Without the cellular backup network, popular
22 contents are not cached or distributed more than any other contents. So, as
23 we see in Figure 8, the percentage of contents obtained through the VANET
24 does not improve even for skewed popularities (i.e. high *Zipf* α). However,
25 when using a cellular backup network, popular contents are downloaded via
26 the cellular network with higher probability than unpopular ones, and thus
27 there are more copies in the VANET. Thus, the seeding effect of the cellu-
28 lar backup network has a significant impact on the performance of the Push
29 protocols, specially with highly skewed popularity distributions, because now
30 the VANET performance increases with popularity.

31
32 The cellular network seeding effect with the Pull and MHTI protocols is
33 different from the Push case. In both Pull and MHTI protocols, contents are
34 requested explicitly in the VANET, so when we have popular contents, these
35 are the contents that most vehicles are looking for, and thus they are dis-
36 tributed more. Therefore, both protocols already benefit from content popu-
37 larity, i.e., when α increases the performance of the protocols improves, even
38 without using a cellular backup network. Still, the cellular network seeding
39 complements this effect helping to improve the performance by introducing
40 new contents in VANET areas far from infostations. Nevertheless, the effect
41 is not as significant as in the Push protocols because, although the benefits of
42 downloading popular contents by cellular would also be greater with highly
43 skewed popularities, in those scenarios contents are already widely replicated
44 in the VANET by the content distribution protocols.

53 54 *5.5. Performance with Internet-like contents*

55 The previous subsections have provided a systematic analysis of how the
56 number of contents and their popularity affect the performance of VANET
57




Figure 9: Distribution of Internet-like contents under different vehicle densities ($M = 100,000$ contents, $Zipf \alpha = 0.8$)

content distribution protocols, as well as the need of a cellular backup network. Although it is still unclear what will be the characteristics of VANET applications, and probably they will greatly vary, we think many infotainment applications will provide Internet contents to drivers and passengers, and thus it is worthy to evaluate the performance of VANET content distribution protocols using similar characteristics to current Internet contents [35].

For that purpose we have repeated the simulations with the maximum number of contents ($M = 100,000$ contents), a fixed Zipf popularity ($Zipf \alpha = 0.8$), and a 5G backup network, while varying the number of cars in order to study the effect of traffic density. Figure 9 shows the VANET content delivery ratio, and Table 7 shows the relative overhead of the three different families of protocols.

As already suggested by previous experiments, the combination of a high number of contents with a moderate popularity reduces the performance of

Table 7: Relative overhead of distributing Internet-like contents under different vehicle densities ($M = 100,000$ contents, *Zipf* $\alpha = 0.8$, N vehicles)

N	Push	Push-opt	Pull	Pull-opt	MHTI	MHTI-opt
2,013	9,099.42	5,554.51	1.59	5.38	5.40	8.28
3,012	7,265.96	6,068.16	1.59	5.45	6.91	8.97
4,011	4,714.99	3,454.89	1.61	5.49	8.30	9.29
5,035	4,232.32	3,366.89	1.64	5.64	9.61	9.82
6,082	4,389.02	2,947.84	1.66	5.92	11.62	10.40
7,087	3,451.90	2,650.76	1.70	5.95	12.60	10.41
8,165	3,484.27	2,957.74	1.77	6.46	15.22	11.13
9,220	3,179.22	2,738.20	1.84	6.64	17.14	11.43

all VANET content delivery protocols, although in different degrees. In particular, Push protocols are almost useless with such high number of contents, even with cellular network seeding, because it is rather difficult that someone advertises the exact content you are interested in. The one-hop Pull protocols only achieve a 20% delivery ratio due to the rather flat popularity, thus VANET contents are mostly obtained when driving close to an infostation. Interestingly, its performance decreases with car density because, in order to avoid traffic jams, the DUA module spreads the cars more over the map (as real drivers do), skipping the large avenues where the four infostations are placed.

The protocol that better copes with such Internet-like scenario is the MHTI one, since it is able to request unpopular contents to the closest infostation. However, its performance greatly depends on traffic density so multi-hop paths can occur. In low traffic densities its performance is slightly better than the one-hop Pull protocol but, with higher densities, it almost reaches a 50% VANET delivery ratio. Table 8 shows the average number of hops needed by the MHTI and the MHTI-opt protocols to get the contents from an infostation or an intermediate car. It is interesting to see that, as the number of average hops increases with higher densities, the difference between the performance of the Pull protocol and the MHTI protocol increases (see Figure 9) as the multi-hop paths are the basis of the advantage of the MHTI protocol compared with the Pull protocol: the MHTI protocol can reach farther to fetch the less popular contents. Additionally, we can see in Table 8 that the MHTI-opt protocol performs less hops than the MHTI protocol, because the optimizations allow finding contents in an intermediate

Table 8: Average number of hops used by the MHTI protocol to get contents ($M = 100,000$ contents, *Zipf* $\alpha = 0.8$, N vehicles)

N	MHTI	MHTI-opt
2,013	1.70	1.60
3,012	2.06	1.78
4,011	2.50	2.01
5,035	2.74	2.17
6,082	2.92	2.26
7,087	3.02	2.27
8,165	3.13	2.25
9,220	3.24	2.24

car closer to the requesting vehicle.

Therefore, in the MHTI protocol, performance increases with car density. This is a convenient behavior because, when there are few cars, they can safely use the cellular backup network to fetch the contents, whereas during peak hours a VANET using the MHTI protocol is able to offload almost half of the traffic from the overloaded cellular network.

6. Conclusions

Content distribution is considered an appropriate communication paradigm to build infotainment and traffic efficiency applications in VANETs. Therefore, in the literature there are several proposals of content distribution protocols for VANETs. The evaluation of these proposals has been mainly focused on performance parameters, such as the percentage of cars that receive a content or the overhead of the proposed protocols. Nevertheless, a review of the literature shows that less effort has been dedicated to model the applications that consume the contents being distributed, and to study how the characteristics of the contents influence the performance of these protocols. Another aspect that has not received enough attention is the combination of the VANET with a cellular network in order to be able to provide a guaranteed service.

In this paper, we have thoroughly studied the influence on the performance of different VANET content distribution protocols of the number and popularity of contents, and the presence of a cellular backup network. Our study has been performed with a realistic simulation model that includes a

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

real urban map, buildings that affect signal propagation, and a faithful micro mobility vehicle simulation. We also consider a cellular backup network so, if contents are not received through the VANET, they are requested through the cellular network. This is often ignored when studying VANET protocols, but in the real world, as the connectivity in the VANET cannot be guaranteed but the service must be, a backup network will be surely required. This means that contents are always obtained, so the main performance figures are the number of contents obtained through the VANET, and thus offloaded from the cellular network, as well as the overhead to do so.

In our study, we have found that VANET content distribution protocols are very sensitive to the total number of contents. The performance of push protocols is very good when distributing a very small number of contents (less than 10), but it is poor with an increase in the number of contents. One-hop pull protocols achieve only a slightly better performance. The popularity of contents contributes to reduce the penalty of an increase in the total number of contents, because popularity reduces the number of significant contents. Still, pull protocols are able to benefit from popularity to improve performance more efficiently than push ones. Additionally, we have shown that a cellular backup network creates a seeding effect that helps improving the performance of all VANET content distribution protocols, although this improvement varies for different families of protocols.

In this paper, we have also proposed a multi-hop pull protocol for content distribution in VANETs, called Multi-Hop To Infostation (MHTI) protocol, in which content requests are sent towards the closest infostation, although they can also be replied by any car in the path that already has the content in its cache. The performance of this protocol allows using the VANET to obtain a significant percentage of contents in peak hours, therefore saving cellular network resources, even with a large number of contents, and it further benefits from content popularities in order to obtain more contents through the VANET.

In our work, a popular content is a content that is requested by many vehicles. But we assume that the popularity of a content is not known by the vehicle requesting it. Therefore, the analyzed content distribution protocols do not use special strategies depending on the popularity of the contents to be requested. In general, knowing a priori the popularity can be very difficult, as it depends on the interests of other vehicles. On the other hand, if popularity were known a priori, it would be possible to implement more intelligent strategies, such as asking for a content through the VANET only

1
2
3
4
5
6
7
8
9 if its popularity is higher than a certain threshold, further optimizing the
10 content distribution protocols.

11 An important conclusion from our work is that any proposal for a VANET
12 content distribution protocol must consider the target applications and their
13 characteristics in terms of number of contents and popularity or, alterna-
14 tively, it must study the performance for a range of number of contents and
15 popularities, as they may significantly influence its performance. We also
16 think that a cellular backup network is also required for reliable service, and
17 its presence influences the performance of the VANET itself, so it should also
18 be considered when studying content distribution protocols.

19 In our future work we plan to investigate scenarios with heterogeneous
20 content size and different latency requirements to analyze their impact on
21 content distribution protocols. We are also interested in analyzing the sensi-
22 bility of the performance of content distribution protocols to the number of
23 infostations.

24 **Acknowledgment**

25 This work has been funded by the Spanish Ministerio de Economía y
26 Competitividad through the ADAS-ROAD project (TRA2013-48314), and by
27 Comunidad de Madrid through the MEDIANET-CM project (S-2009/TIC1468).

28 **References**

- 29
30
31
32
33
34
35
36
37
38
39 [1] V. Sandonis, I. Soto, M. Calderon, M. Urueña, Vehicle to internet com-
40 munications using the ETSI ITS GeoNetworking protocol, *Transactions*
41 *on Emerging Telecommunications Technologies*. 27 (3) (2016) 373 – 391.
42 doi:10.1002/ett.2895.
43
44
45 [2] C. Maihöfer, T. Leinmüller, E. Schoch, Abiding geocast: Time-
46 stable geocast for ad hoc networks, in: *ACM International Workshop*
47 *on Vehicular Ad Hoc Networks (VANET)*, ACM, 2005, pp. 20–29.
48 doi:10.1145/1080754.1080758.
49
50
51 [3] H. Wu, R. Fujimoto, R. Guensler, M. Hunter, MDDV: A mobility-centric
52 data dissemination algorithm for vehicular networks, in: *ACM Interna-*
53 *tional Workshop on Vehicular Ad Hoc Networks (VANET)*, ACM, 2004,
54 pp. 47–56. doi:10.1145/1023875.1023884.
55
56
57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
- [4] B. Xu, A. Ouksel, O. Wolfson, Opportunistic resource exchange in inter-vehicle ad-hoc networks, in: IEEE International Conference on Mobile Data Management. Proceedings, 2004, pp. 4–12. doi:10.1109/MDM.2004.1263038.
 - [5] L. Wischhof, A. Ebner, H. Rohling, Information dissemination in self-organizing intervehicle networks, IEEE Transactions on Intelligent Transportation Systems 6 (1) (2005) 90–101. doi:10.1109/TITS.2004.842407.
 - [6] Y. Lim, S. Ahn, K.-H. Cho, Abiding geocast for commercial ad dissemination in the vehicular ad hoc network, in: IEEE International Conference on Consumer Electronics (ICCE), 2011, pp. 115–116. doi:10.1109/ICCE.2011.5722491.
 - [7] N. H. Khiadani, S. M. Safavi Hemami, F. Hendessi, Introducing a handshake for data dissemination using rateless codes in vehicular ad hoc networks, Transactions on Emerging Telecommunications Technologies. Accepted for publication, available online. doi:10.1002/ett.2944.
 - [8] V. Kone, H. Zheng, A. Rowstron, G. O’Shea, B. Zhao, Measurement-based design of roadside content delivery systems, IEEE Transactions on Mobile Computing 12 (6) (2013) 1160–1173. doi:10.1109/TMC.2012.90.
 - [9] N. Cheng, N. Lu, N. Zhang, X. S. Shen, J. W. Mark, Vehicular WiFi offloading: Challenges and solutions, Vehicular Communications 1 (1) (2014) 13 – 21. doi:10.1016/j.vehcom.2013.11.002.
 - [10] M. Hatton, M2m growth necessitates a new approach to network planning and optimisation, Machina Research, white paper, <https://machinaresearch.com/report/m2m-growth-necessitates-a-new-approach-to-network-planning-and-optimisation/> (May, 2015).
 - [11] P. Costa, D. Frey, M. Migliavacca, L. Mottola, Towards lightweight information dissemination in inter-vehicular networks, in: ACM International Workshop on Vehicular Ad Hoc Networks (VANET), ACM, 2006, pp. 20–29. doi:10.1145/1161064.1161069.
 - [12] I. Leontiadis, P. Costa, C. Mascolo, Persistent content-based information dissemination in hybrid vehicular networks, in: IEEE International

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Conference on Pervasive Computing and Communications (PerCom),
2009, pp. 1–10. doi:10.1109/PERCOM.2009.4912758.

- [13] Omnet++ discrete event simulator, <http://www.omnetpp.org/> ([Accessed: June 27, 2016]).
- [14] Sumo – simulation of urban mobility, <http://sumo-sim.org/> ([Accessed: June 27, 2016]).
- [15] D. Krajzewicz, J. Erdmann, M. Behrisch, L. Bieker, Recent development and applications of SUMO - Simulation of Urban MObility, *International Journal on Advances in Systems and Measurements* 5 (3&4) (2012) 128–138.
- [16] J. Zhao, Y. Zhang, G. Cao, Data pouring and buffering on the road: A new data dissemination paradigm for vehicular ad hoc networks, *IEEE Transactions on Vehicular Technology* 56 (6) (2007) 3266–3277. doi:10.1109/TVT.2007.906412.
- [17] M. Amadeo, C. Campolo, A. Molinaro, CRoWN: Content-centric networking in vehicular ad hoc networks, *IEEE Communications Letters* 16 (9) (2012) 1380–1383. doi:10.1109/LCOMM.2012.072012.120282.
- [18] C. Bian, T. Zhao, X. Li, W. Yan, Boosting named data networking for data dissemination in urban VANET scenarios, *Vehicular Communications* 2 (4) (2015) 195 – 207. doi:10.1016/j.vehcom.2015.08.001.
- [19] R. Atat, E. Yaacoub, M.-S. Alouini, F. Filali, A. Abu-Dayya, Delay-sensitive content distribution via peer-to-peer collaboration in public safety vehicular ad-hoc networks, *Ad Hoc Networks* 16 (2014) 182–196. doi:10.1016/j.adhoc.2013.12.007.
- [20] M. Gerla, C. Wu, G. Pau, X. Zhu, Content distribution in VANETs, *Vehicular Communications* 1 (1) (2014) 3 – 12. doi:10.1016/j.vehcom.2013.11.001.
- [21] M. Gramaglia, C. J. Bernardos, M. Calderon, Seamless internet 3G and opportunistic WLAN vehicular connectivity, *EURASIP Journal Wireless Communications and Networking* 2011:183 (2011) 1–20. doi:10.1186/1687-1499-2011-183.

- 1
2
3
4
5
6
7
8
9
10 [22] I. Lequerica, P. Ruiz, V. Cabrera, Improvement of vehicular commu-
11 nications by using 3G capabilities to disseminate control information,
12 IEEE Network 24 (1) (2010) 32–38. doi:10.1109/MNET.2010.5395781.
13
14 [23] C. J. Bernardos, I. Soto, M. Calderon, F. Boavida, A. Azcorra, VARON:
15 Vehicular ad hoc route optimisation for NEMO, Computer Communi-
16 cations 30 (8) (2007) 1765–1784. doi:10.1016/j.comcom.2007.02.011.
17
18 [24] G. Remy, S.-M. Senouci, F. Jan, Y. Gourhant, LTE4V2X – impact of
19 high mobility in highway scenarios, in: Global Information Infrastruc-
20 ture Symposium (GIIS), 2011, pp. 1–7. doi:10.1109/GIIS.2011.6026706.
21
22 [25] J. Ahn, M. Sathiamoorthy, B. Krishnamachari, F. Bai, L. Zhang, Op-
23 timizing content dissemination in vehicular networks with radio hetero-
24 geneity, IEEE Transactions on Mobile Computing 13 (6) (2014) 1312–
25 1325. doi:10.1109/TMC.2013.100.
26
27 [26] B. Karp, H. T. Kung, Gpsr: Greedy perimeter stateless routing for wire-
28 less networks, in: Proceedings of the 6th Annual International Confer-
29 ence on Mobile Computing and Networking, MobiCom’00, ACM, 2000,
30 pp. 243–254. doi:10.1145/345910.345953.
31
32 [27] Veins – vehicles in network simulation, <http://veins.car2x.org/>
33 ([Accessed: June 27, 2016]).
34
35 [28] C. Sommer, D. Eckhoff, R. German, F. Dressler, A computationally
36 inexpensive empirical model of IEEE 802.11p radio shadowing in ur-
37 ban environments, in: Eighth International Conference on Wireless
38 On-Demand Network Systems and Services (WONS), 2011, pp. 84–90.
39 doi:10.1109/WONS.2011.5720204.
40
41 [29] D. Krajzewicz, Traffic simulation with SUMO - Simulation of Urban
42 MObility, Fundamentals of Traffic Simulation 145 (2010) 269–293.
43
44 [30] Openstreetmap, <http://www.openstreetmap.org/>
45 ([Accessed: June 27, 2016]).
46
47 [31] Java openstreetmap editor, <http://josm.openstreetmap.de/>
48 ([Accessed: June 27, 2016]).
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- [32] C. Gawron, Simulation-based traffic assignment – computing user equilibria in large street networks, Ph.D. thesis (1999).
URL <http://e-archive.informatik.uni-koeln.de/366/>
- [33] M. Behrisch, D. Krajzewicz, P. Wagner, Y. Wang, Comparison of methods for increasing the performance of a DUA computation, in: International Symposium on Dynamic Traffic Assignment (DTA), 2008, pp. 1–9.
- [34] P. Hart, N. Nilsson, B. Raphael, A formal basis for the heuristic determination of minimum cost paths, *IEEE Transactions on Systems Science and Cybernetics* 4 (2) (1968) 100–107. doi:10.1109/TSSC.1968.300136.
- [35] C. Fricker, P. Robert, J. Roberts, N. Sbihi, Impact of traffic mix on caching performance in a content-centric network, in: *IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2012, pp. 310–315. doi:10.1109/INFCOMW.2012.6193511.