Advertisement Delivery and Display in Vehicular Networks

C.F. Chiasserini†⋆, F. Malandrino†, M. Sereno ‡
† Politecnico di Torino, Italy
⋆CNR-IEIIT, Italy
‡Università degli Studi di Torino, Italy

Abstract—**Advertisement delivery is expected to play a crucial role in future vehicular networks. In this paper, we address such a problem in vehicular networks, where advertisements (ads) can be broadcasted by roadside units (RSU) as well as vehicles, and then displayed to interested users. We describe the advertisement dissemination process by means of an optimization model that aims at maximizing the number of ads that users display within the target area and validity period of the ad. We then solve the optimization problem, obtaining the optimal scheduling strategy that RSUs and vehicles should adopt for ad broadcasting. Our study highlights the important role that vehicle-to-vehicle communication will have in ads delivery. Also, it shows how coexisting vehicular and cellular networks can effectively complement each other, with vehicular networks being a very efficient means for pervasive ad dissemination.**

I. Introduction

Vehicular networking is one the most important enabling technologies for a massive set of applications related to vehicles, such as traffic management, transport efficiency, user entertainment, office on-wheels, commercial and tourist information, mobile shopping, and emergency management applications. Car manufacturers, as well as the research community, are investing considerable time and resources in designing new applications and/or protocols that intercept, anticipate, and often stimulate the development of new services for the vehicular networking world. It is easy to predict that the closer interaction that is expected in the near future, between vehicles as nodes of the future Internet, will further foster this trend.

From the networking viewpoint, an interesting direction in this context is the integration of infrastructure-to-vehicle (I2V) with vehicle-to-vehicle (V2V) communications. Such an integration can indeed enhance the dissemination of information from the network infrastructure to the vehicles, with the data collected by the vehicles themselves. We can indeed envision a highly scalable mechanism for gathering and disseminating data about traffic conditions, tourist information, (location-aware) advertisement dissemination (see, e.g., [1]–[3]).

In addition to that, there is a clear trend followed by all car manufacturers concerning the development of in-car platforms that can host sophisticated navigation and passenger entertainment systems, sensing devices and applications, aiming at increasingly integrating the car and the smartphone worlds. An obvious side effect of this trend leads to the possibility for users traveling on cars (both drivers and passengers) to access the same content, applications, and business opportunities as stationary users.

Given such an exciting scenario, in this paper we investigate how advertisement services can be effectively supported in vehicular systems. Indeed, mobile advertising already provides huge revenues, and it is expected to do so also in the vehicular context. In particular, relying on the successful business model that has been adopted for web sites, Internet blogs, and social networks, we address the delivery of *targeted* advertisements (ads for short) to cars. Targeted advertising is especially relevant as it is beneficial to both advertisers and users: advertisers can gain higher revenue by delivering an ad to users with a strong potential to purchase, and users in turn receive more pertinent and useful information that match their interests (see [4] for a survey on the economical effects of targeted advertising). Clearly, users' interests may be driven by various factors such as personality, likes and dislikes. Since we address a vehicular scenario, in this work we consider that they are driven also by geographical position and time of the day. Thus, given the users' profiles and current/predicted mobility, we identify the ads in which likely users are interested and exploit this information to select the set of ads that roadside units (e.g., base stations, WiFi access points), as well as fellow vehicles, should broadcast.

It is worth stressing that ad dissemination in vehicular networks is related to several other issues, such as alternative business models underlying the service and the kind of entities and/or markets where demand and offer for such a service can meet. Importantly, advertising must also be compliant with traffic safety and not disrupt the driver's attention. For instance, ads should be accessed by passengers only, or displayed to drivers with sufficiently low frequency and only when the car's speed is low enough, or when the car is stopped. This scenario, however, will significantly change with the advent of self-driving vehicles: in this case there will be no need for such limitations, or they could be greatly relaxed.

Data dissemination and caching in VANETs are a widely studied problem [5], [6], with most researchers seeking to exploit the largely predictable user mobility in order to improve system performance. However, we emphasize that the peculiarities of targeted advertising make the problem under study very different from classical data dissemination. In particular, the study of ad dissemination introduces issues

related to content relevance (and, hence, content value) that depends on the car location, the time of the day, and, in general, the spatial and time validity of the ads. Another related problem to ours is that of publishing content over a shared medium, where several users can post media items (messages, pictures, or videos). The latter finds application in many contexts and has been studied in, e.g., [7]–[9]. These pieces of work, however, investigate issues related to optimal timing and rates at which users should post their content over the medium to maximize visibility. Finally, relevant to our study are also the works in [10], [11], which address the delivery and display of ads on social and vehicular networks, respectively. There, the focus is on the optimal display strategy that maximizes the expected number of clicks on the delivered ads. The intriguing finding of [10] concerns the importance of targeted ads, as it shows that ad dissemination strategies that account for user interests can significantly increase the expected click rate.

The rest of the paper is organized as follows. We describe the system model under study in Section II. We then formulate the optimization problem targeting the scheduling of ad broadcasting at roadside units and vehicles, subject to the system constraints, in Section III. In Section IV we solve the problem, along with a low-complexity relaxation thereof, and show numerical results. Section V concludes the paper and discusses future work.

II. System model

We consider a vehicular reference scenario including the following entities: the service manager, the advertisements, the roadside units (RSU), and the vehicles.

The service manager is a centralized entity, which can be deployed in the cloud. Area-specific servers, however, can be envisioned, and located at the edge of the network according to the well-known Mobile Edge Computing (MEC) paradigm. The service manager will communicate with the vehicles through cellular or DSRC communication, the first being preferable due to the need for pervasive coverage.

Ads are the content items that the service manager aims at showing to vehicular users. All ads have the same size. Each ad has a target area and a target time period, i.e., a set of space *and* time coordinates at which the vehicle has to be for the ad to be relevant to its driver/passengers. Also, we associate with each ad a display time, i.e., how long it takes to display the ad.

The RSUs provide partial or full coverage of the geographical service area of interest and can broadcast ads to vehicles. Vehicles can communicate with RSUs as well as with other vehicles, using, e.g., the WAVE technology. This implies that a vehicle can receive ads from either RSUs or other fellow vehicles in its proximity. Vehicles are equipped with an enhanced on-board navigation platform that can display ads if they are relevant to the user and if the ads are still valid in space and time. Vehicles display an ad at most once, and they can display at most one ad at a time.

The service manager has to make decisions on which RSUs and which vehicles should broadcast ads and which ads they should transmit. The revenue it receives from ad broadcasting depends on the number of ads displayed by the vehicles – a piece of information that can be returned by the cars on-board platform. In order to make its decisions, the service manager leverages on the knowledge of the RSUs radio coverage, the ads' target area and time period, and the current and future vehicle positions. Current vehicle position can be easily acquired through RSUs by exploiting the beacon messages that vehicles are required to periodically transmit, while future positions can be obtained through well-known mobility prediction techniques. Also, the interests of a user (hereinafter also referred to as vehicle or car for brevity) are assumed to be known thanks to today's commonplace user profiling techniques. V2V transfers additionally require the service manager to be aware of the content of the vehicles' caches. This information can be obtained from the vehicles themselves (at the cost of a small overhead), or reconstructed with high precision from the vehicles' trajectories.

For simplicity, consider that time is divided into slots, and that ad broadcast transmissions last one time slot each. We assume that each RSU and vehicle can broadcast at most one ad per time slot, due to bandwidth limitations.

Given the above system, the advertisement process can be summarized as follows:

1) ads are *disseminated* by RSUs as determined by the service manager;
2) ads are *stored* by vehicles and re-broadcasted as determined by the service manager;
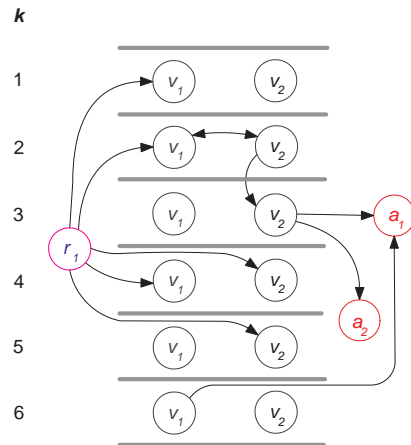3) ads are *displayed* to the vehicle driver/passengers if relevant and valid.

Fig. 1. Time-expanded graph representing the network evolution over time (k denotes the time slots). The example includes one RSU (r_1), two vehicles (v_1 and v_2), and two ads (a_1 and a_2).

III. Problem formulation

The advertisement process can be conveniently represented through an oriented time-expanded graph (TEG), in a similar

$$\max \sum_{(v_k, a)} x(v_k, a, \langle a \rangle) \tag{1}$$

$$\sum_a y(e, \langle a \rangle, k) \leq 1 \tag{2}$$

$$x(e_k, v_k, \langle a \rangle) \leq y(e, \langle a \rangle, k) \tag{3}$$

$$h(v, \langle a \rangle, k) \leq \\ \sum_{(e_k, v_k)} [x(e_k, v_k, \langle a \rangle) + x(v_{k-1}, v_k, \langle a \rangle)] \tag{4}$$

$$y(v, \langle a \rangle, k) \leq h(v, \langle a \rangle, k) \tag{5}$$

$$x(v_k, v_{k+1}, \langle a \rangle) \leq h(v, \langle a \rangle, k) \tag{6}$$

$$x(v_k, a, \langle a \rangle) \leq h(v, \langle a \rangle, k) \tag{7}$$

$$\sum_a x(v_k, v_{k+1}, \langle a \rangle) \leq C \tag{8}$$

$$\sum_k x(v_k, a, \langle a \rangle) \leq 1 \tag{9}$$

$$\sum_{h=k}^{k+d} \sum_a x(v_h, a, \langle a \rangle) \leq 1 \tag{10}$$

Fig. 2. Problem formulation: objective function (1) and constraints (2)–(10).

way to [12], [13]. As shown in Figure 1, the graph comprises three types of vertices: those representing the RSUs, the vertices corresponding to the vehicles at different time slots, and those representing the ads. Note that using different vertices to represent the same vehicle at different time slots, allows us to account for vehicle mobility and the resulting changes in communication opportunities.

Edges connecting the instances of the same vehicle at consecutive time slots represent the fact that a vehicle can carry cached ads around, over time. In particular, the capacity of these edges corresponds to the vehicle cache size (expressed in ads) and denoted by C, and the flow over these edges by $x(v_k, v_{k+1}, \langle a \rangle)$, where v_k and v_{k+1} are the vertices representing a vehicle at slot k and $k+1$, respectively, and $\langle a \rangle$ indicates the ad cached by the vehicle. Edges between instances of different vehicles at the same time slot represent the existence of a V2V link, i.e., the possibility for a vehicle to receive the ad broadcasted by the other fellow vehicles. We assume that all these edges have unitary capacity, representing the ability of each vehicle to broadcast one ad per time slot (recall that all ads have the same size). We also remark that, assuming bidirectional links, V2V edges are symmetric, hence our TEG will contain cycles. Similarly, edges from RSUs to vehicles represent infrastructure-to-vehicle communications; these edges too have unitary capacity. The flow over an edge from a vehicle (RSU, resp.) to another vehicle at time k is denoted by $x(e_k, v_k, \langle a \rangle)$, where e may be either a vehicle or an RSU. Finally, edges between vehicle and ad vertices

indicate that the ad is relevant to the vehicle and is valid at the current spatial location and time slot, i.e., if received, the ad would be displayed by the vehicle. The flow over such edges is denoted by $x(v_k, a, \langle a \rangle)$.

Recall that the service manager is in charge of selecting the ads that should be broadcasted by the RSUs and the vehicles at the different time slots. Its goal is to maximize its own revenue, i.e., the number of ads that are displayed by vehicles. Given the TEG, we can then formulate the decision problem at the service manager by associating a binary variable with each (vehicle, ad) pair and with each (RSU, ad) pair, indicating whether the ad should be transmitted or not, by the vehicle or the RSU, respectively, at that time slot. In the formulation in Figure 2, we denote such variable by $y(e, \langle a \rangle, k)$, where e is either a vehicle or an RSU. Likewise, in the formulation in Figure 2, the binary variables $x(v_k, a, \langle a \rangle)$ associated with the (vehicle, ad) pairs indicate whether a vehicle v displays ad a at time slot k.

The goal of the service manager can thus be expressed by maximizing the sum of the latter type of variables (i.e., number of displayed ads), as shown in (1) in Figure 2. The values of the decision variables have to meet the following constraints:

- each RSU and vehicle can broadcast at most one ad per time slot (Eq. (2));
- an ad reception occurs at a vehicle only if the ad is transmitted by an RSU or another vehicle (Eq. (3));
- a vehicle may have an ad in its cache only if it receives the ad from an RSU or another vehicle, or if it has stored the ad at a previous time slot (see Eq. (4) where the binary variable $h(v, \langle a \rangle, k)$ denotes whether vehicle v is storing ad a in its cache at slot k);
- a vehicle can broadcast (Eq. (5)), carry around (Eq. (6)), or display (Eq. (7)) an ad, only if the ad is stored in the vehicle's cache;
- the total number of ads cached at a vehicle cannot exceed the cache size (Eq. (8) in Figure 2);
- a vehicle can display each ad at most once (Eq. (9)), and no ad can be displayed before d time slots have elapsed from the previous one (Eq. (10)), where d is the ad display time.

Similarly to [13], we also consider an additional variable and some additional equations (not shown here for simplicity), in order to account for collisions. Specifically, we ensure that no vehicle transmits if another vehicle or RSU within its radio range is transmitting [13].

We remark that as output of the optimization problem we obtain the ads (if any) that each RSU and vehicle have to broadcast at each time slot, and, as a side product, the ads that each vehicle should cache. The above problem, however, is an Integer Linear Problem (ILP), which is NP-hard, thus impossible to use in real time for large-scale scenario. In our performance evaluation, we therefore consider an *LP relaxation* of the problem where binary variables are allowed to take real values in $[0, 1]$. The LP relaxation has a computational complexity that is cubic in the problem size in the worst case, and much lower in most practical instances [14]. However, the

solution of the relaxed problem cannot be used directly, as the decision values related to ad broadcasting/storage may now be real. Thus we interpret the "relaxed" values as priorities, i.e., intuitively, how important is that a certain RSU/vehicle broadcasts a given ad, and that a vehicle stores an ad, at a certain time slot. Such priority values will be used by vehicles and RSUs to make their own transmission/storage decisions.
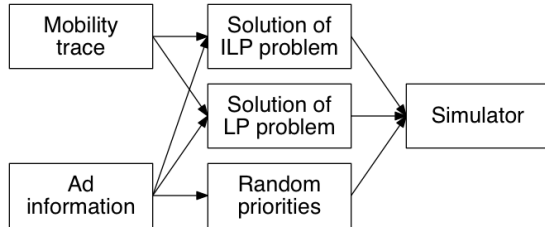

Fig. 3. Simulation process.

Once the decision/priority values have been obtained, we input them to a custom simulator, as depicted in Figure 3. The tool simulates which entities are in radio visibility of each other, who transmits, what, and when, under the scenario detailed next.

IV. NUMERICAL RESULTS

Simulations are run on the same service area as in [15], representing an 1×1.5 km^2 section of the urban area of Ingolstadt, Germany. The scenario models a total number of cars ranging between 474 and 2510, over a period of about 1 hour with a mean traffic density between 10 and 124 vehicles per km^2, respectively. The vehicular mobility is generated using the well-known Simulator of Urban MObility (SUMO), capable of reproducing real-world microscopic and macroscopic road traffic [16]. As shown in Figure 4, 20 RSUs are deployed (red dots in the figure) at the centers of as many intersections, which were randomly selected. Consistently with [13], the transmission range is assumed to be 120 m for RSUs, and 80 m for vehicles. The total fraction of the service area covered by RSUs is equal to 0.4. We also identify 10 points-of-interest (POIs), around which we draw an interest area of radius equal to 200-meter. Interest areas are highlighted in blue in the figure and can overlap in space. Each POI is associated with one ad, and we assume that all vehicles within the interest area of a POI are interested in viewing the corresponding ad. Notice that vehicles will display the ad *only* when inside the area – which of course does not prevent them from getting the ad beforehand or from storing the ad afterwards. Each time slot is one minute long and the display time of each ad is set to one time slot. Unless otherwise specified, the cache size is set to 10 ads.

As far as ad broadcasting by RSUs and vehicles is concerned, the ads to be transmitted (if any) at a given time slot are determined by adopting the following strategies:

Optimal: The simulator takes as input the optimal decisions obtained through the ILP optimization.

Relaxed: The simulator takes as input the solution of the relaxed optimization problem and use it to establish priorities in the transmission of ads by vehicles and RSUs. This case allows us to evaluate the performance loss that occurs when we reduce the problem solution complexity.

Random: Ads are chosen at random by RSUs and vehicles – of course, vehicles can only transmit ads they have in their cache. This case is used as a benchmark; its choice is motivated by the implementation simplicity that it exhibits.
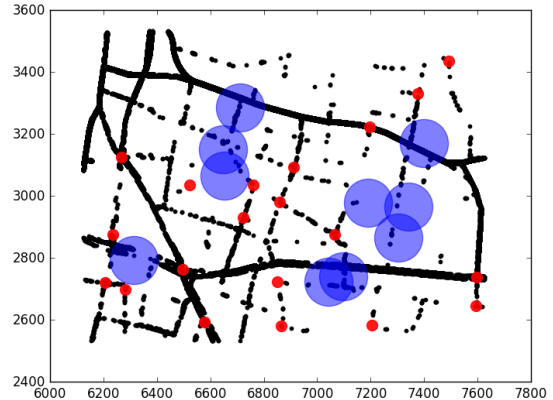

Fig. 4. Topology: Red dots correspond to RSUs while blue areas denote the interest areas of POIs. The radius of the interest area is set to 200 m.

The following plots depict our main performance metric, i.e., the cumulative number of ads displayed by the vehicles over time. While deriving the results, we considered that the advertisement service (i.e., ads transmission) starts at the first time slot. The number of displayed ads therefore increases with time since the number of vehicles, traveling by the various POIs and receiving the ads, grows as time passes.

In Figure 5 (left), the black line represents the performance that could be achieved in an ideal case, e.g., if we used a cellular network providing full coverage of the whole area and whose base stations (all of them, not only those covering the ads interest area) transmit all ads at every time slot. The red and gray lines instead represent what we obtain when only RSUs and V2V communications are exploited to support the advertisement service. In particular, the red line corresponds to the case where optimum ad scheduling decisions are fed to the simulator, i.e., we solve the full ILP problem in Figure 3, while the gray line corresponds to the case in which LP priorities are used. Finally, the blue curve refers to the Random strategy.

The plot brings very good news: (i) the vehicular network is able to reach around 50% of the users with respect to the ideal case, (ii) the performance loss we incur when moving from optimal to relaxed decisions is quite limited (it is always below 10%), (iii) both the optimal and the relaxed strategies provide way better performance than a simple random ad selection. As far as the first point is concerned, we stress that delivering an ad to 50% of the interested users with respect to the ideal case, is quite a good achievement, although the performance could improve in more favorable scenarios.

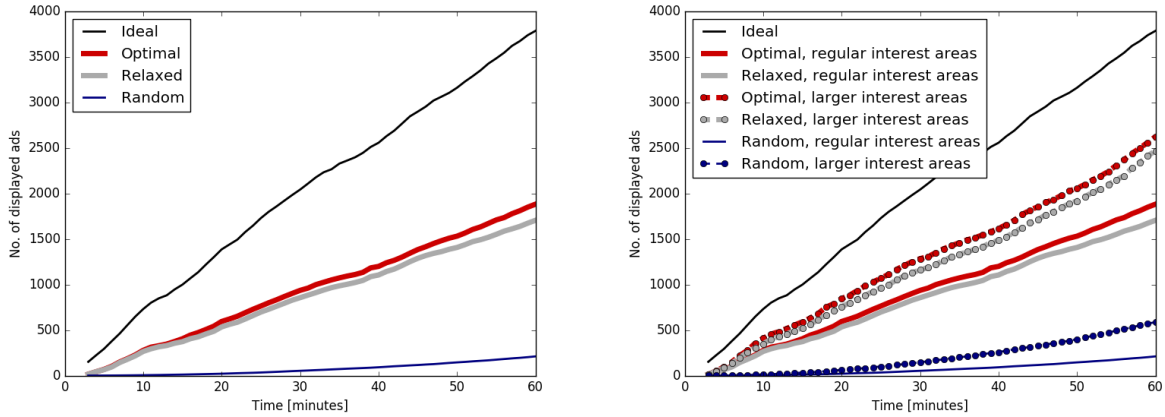An effective way to increase the number of displayed ads

Fig. 5.   Basic scenario: performance of the Optimal, Relaxed and Random strategies under the default scenario (left); effect of 400-meter interest areas (right).
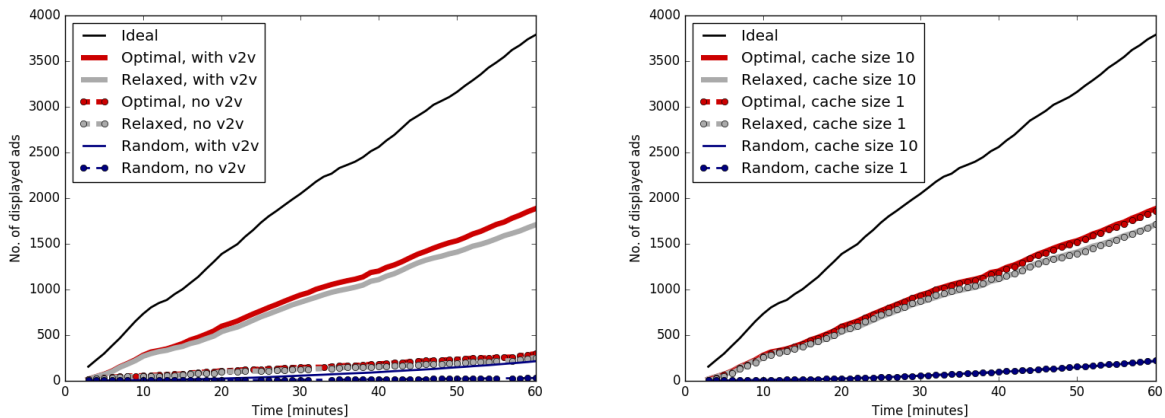

Fig. 6.   Effect on performance of forbidding V2V transfers (left) and of reducing the cache size from 10 ads to 1 ad (right).

is to further exploit the communication opportunities afforded by V2V links. This is possible by enlarging the radius of the area of interest around each POI. We do so by increasing such a parameter from 200 to 400 meters. Figure 5 (right) confirms this intuition: looking at the red curve (corresponding to the solution of the ILP formulation) and at the gray curve (corresponding to the solution of the LP formulation), we note that a larger radius has indeed the effect of substantially increasing the number of vehicles that the advertisement service can reach. In other words, vehicular networks can effectively exploit V2V communications for ad dissemination when the RSU deployment is sparse. In particular, the larger the spatial validity of the ads, the higher the chance that vehicles can disseminate ads across large geographical areas, hence the benefit of V2V ad broadcasting. Thus, if both vehicular and cellular networks are available for ad transmission, they should be used in a complementary way: the vehicular network should support ads with larger coverage areas – possibly global –, while the cellular network should broadcast local ads.

Next, in Figure 6 (left) we investigate how important the contribution of V2V transfers is to ad dissemination. Solid lines represent the case where V2V communications are allowed, while dashed lines depict the case where they are

forbidden. It is evident that in the latter case the number of displayed ads drops severely, leading to performance as poor as that of the Random scheme. This is a further proof that it is indeed V2V transfers that make vehicular networks a viable solution, even in such scenarios as ours where a significant number of RSUs are deployed.

Finally, in Figure 6 (right) we reduce the size of the cache C used by the vehicles to store the received ads: we decrease it from ten ads to one. The effect on the overall performance is surprisingly limited, with the curves for cache size equal to one that almost overlap with those corresponding to cache size equal to ten. This is again good news: it is a clear sign that, in general, vehicles do not need to carry multiple ads in their cache.

V. CONCLUSION AND DISCUSSION

The recent trend, according to which cars are becoming connected nodes of the Internet of Things, opens a huge number of new opportunities and business models. In this paper we investigated such a scenario by considering a targeted advertisement service for vehicular users, where the advertisement relevance to vehicles also depend on the car location and on the time of the day. We study the feasibility

of such a new service and the challenges that have to be taken in order to implement it effectively. In particular, we represented the process of advertisement delivery and display through a graph model that captures the time dynamics of the vehicular network, the limitations of wireless broadcasting – the foremost being collisions –, the spatial and time validity of the advertisements, as well as the users' interests. Through this model, we formulated an optimization problem that maximizes the number of advertisements that are displayed by vehicles. Our results, obtained in a realistic urban topology and mobility scenario, show that: (i) vehicle-to-vehicle communications play a primary role in advertisement delivery, and (ii) vehicular networks can be effectively combined with cellular networks for an efficient service support, with vehicular networks being particularly suitable for the delivery of advertisements with wide spatial validity.

Future research can span according to several directions. While our model already accounts for packet collisions, signal propagation and channel access issues could be added thereto in order to evaluate the control overhead implied by the proposed scheme. Indeed, while information on the vehicles' position and interests are easy to obtain in next-generation vehicular networks, some extra control data is required to be transmitted. Specifically, when V2V communication is enabled, the service manager needs to notify the vehicles about the ads they should store/broadcast (if any) at every time slot. Also, business models should be developed for assessing the popularity and the revenue that can be expected from targeted advertising in vehicular networks. Other important issues concern the definition of different management policies such as caching and selective targeting strategies that account for different mobility patterns and models of time and space-dependent user interests and advertisement validity.

References

[1] T. Economist, "The connected car: Smartphones on wheels," *The Economist: Technology Quarterly:*, vol. Q3 2014, Sep. 2014.

[2] P. Gray, "Is the car the next smartphone?" *Mobile Enterprise*, 2014, *Available at http://www.techrepublic.com/blog/mobile-enterprise/is-the-car-the-next-smartphone/*.

[3] N. Team, "From the smartphone to the smartcar," *Nurun blog*, 2014, *Available at http://www.nurun.com/en/our-thinking/internet-of-things/from-the-smartphone-to-the-smartcar/*.

[4] H. Beales, "The value of behavioral targeting," *Network Advertising Initiative*, 2010.

[5] H. Wu, R. Fujimoto, R. Guensler, and M. Hunter, "MDDV: A mobility-centric data dissemination algorithm for vehicular networks," in *Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks*, ser. ACM VANET, 2004, pp. 47–56.

[6] C. Maihofer and R. Eberhardt, "Geocast in vehicular environments: caching and transmission range control for improved efficiency," in *IEEE Intelligent Vehicles Symposium, 2004*, June 2004, pp. 951–956.

[7] Z. Lotker, B. Patt-Shamir, and M. R. Tuttle, "A game of timing and visibility," *Games and Economic Behavior*, vol. 62, no. 2, pp. 643–660, 2008.

[8] E. Altman and N. Shimkin, "Strategic Posting Times over a Shared Publication Medium," in *International Conference on Network Games Control and Optimization (NETGCOOP)*, 2014.

[9] A. Reiffers-Masson, Y. Hayel, and E. Altman, "Game theory approach for modeling competition over visibility on social networks," in *6th International Conference on Communication Systems and Networks (COMSNETS)*, 2014.

[10] Z. Abbassi, A. Bhaskara, and V. Misra, "Optimizing display advertising in online social networks," in *24th International World Wide Web Conference (WWW)*, Florence, Italy, 2015.

[11] G. Einziger, C.-F. Chiasserini, and F. Malandrino, "Effective selection of targeted advertisements for vehicular users," in *ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, 2016, pp. 11–18.

[12] F. Malandrino, C. Casetti, C. F. Chiasserini, and M. Fiore, "Content downloading in vehicular networks: What really matters," in *IEEE INFOCOM*, 2011, pp. 426–430.

[13] F. Malandrino, C. Casetti, C. F. Chiasserini, and M. Fiore, "Optimal content downloading in vehicular networks," *IEEE Transactions on Mobile Computing*, vol. 12, no. 7, pp. 1377–1391, 2012.

[14] J. Mattingley and S. Boyd, "Real-time convex optimization in signal processing," *IEEE Signal Processing Magazine*, 2010.

[15] D. Eckhoff, C. Sommer, R. German, and F. Dressler, "Cooperative awareness at low vehicle densities: How parked cars can help see through buildings," in *IEEE Global Telecommunications Conference (GLOBECOM)*, Dec. 2011, pp. 1–6.

[16] C. Sommer, R. German, and F. Dressler, "Bidirectionally coupled network and road traffic simulation for improved IVC analysis," *IEEE Transactions on Mobile Computing*, vol. 10, no. 1, pp. 3–15, Jan. 2011.