

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Advertisement Delivery and Display in Vehicular Networks

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1525515> since 2020-02-19T16:08:17Z

Publisher:

IEEE

Published version:

DOI:10.1109/VTCTFall.2015.7390833

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

This is the author's final version of the contribution published as:

C. Borgiattino, C. F. Chiasserini; Malandrino, F.; Sereno, M.. Advertisement Delivery and Display in Vehicular Networks, in: Proceedings of the 82nd IEEE Vehicular Technology Conference (VTC 2015-Fall), IEEE, 2015, 978-1-4799-8091-8, pp: 1-5.

The publisher's version is available at:

<http://xplorestaging.ieee.org/ielx7/7389774/7390768/7390833.pdf?arnumber=7390833>

When citing, please refer to the published version.

Link to this full text:

<http://hdl.handle.net/2318/1525515>

Advertisement Delivery and Display in Vehicular Networks

C. Borgiattino[†], C.F. Chiasserini[†], F. Malandrino^{*}, M. Sereno[‡]

[†] Politecnico di Torino, Italy

^{*} The Hebrew University of Jerusalem, Israel

[‡]Università degli Studi di Torino, Italy

June 27, 2016

Abstract

The role of vehicles has been rapidly expanding to become a different kind of utility, no longer just vehicles but nodes of the future Internet. The car producers and the research community are investing considerable time and resources in the design of new protocols and applications that meet customer demand, or that foster new forms of interaction between the moving customers and the rest of the world. Among the variety of new applications and business models, the spreading of advertisements is expected to play a crucial role. Indeed, advertising is already a significant source of revenue and it is currently used over many communication channels, such as the Internet and television. In this paper, we address the targeting of advertisements in vehicular networks, where advertisements are broadcasted by Access Points and then displayed to interested users. In particular, we describe the advertisement dissemination process by means of an optimization model aiming at maximizing the number of advertisements that are displayed to users within the advertisement target area and target time period. We then solve the optimization problem on an urban area, using realistic vehicular traffic traces. Our results highlight the importance of predicting vehicles mobility and the impact of the user interest distribution on the revenue that can be obtained from the advertisement service.

1 Introduction

Recently, vehicular networks have received a great deal of attention due to a steady increase of interest in safety and entertainment applications (see, e.g., [5, 7, 12]). It is easy to foretell that the integration of cars as nodes of the future Internet will foster new research challenges spanning in many directions. Examples of this trend are, for instance, the new car navigation systems that provide real time traffic information and driving assisting facilities. An interesting direction in this field is the opportunistic collection of information to build a highly scalable “zero-infrastructure” system for gathering and disseminating information about traffic conditions [6]. Other examples of applications propose the combined use of vehicular ad-hoc networking with social networks to develop innovative dissemination methods for tourist information [10].

In this paper we focus on the interesting application of advertising, which is already a significant source of revenue for most online social networks. In particular, we address the spreading of (targeted) advertisements (ads for short) to a selected set of cars. The selection can be based on

many different criteria such as car location, interests, time. Indeed, there are many studies showing that targeted ads is beneficial to both advertisers and users: advertisers can gain higher revenue by advertising to users with a strong potential to purchase, and users in turn receive more pertinent and useful ads that match their interests (see [4] for a survey on the economical effects of targeted advertising).

The problem that we address resembles to a variation of the problem of the publication of content items over a shared medium where several users can post media items such as messages, pictures, or videos. This kind of problem has been studied in many contexts, see, e.g., [2, 9, 11]. Such works investigate issues of the problem related to the optimal timing and rates at which users may post content items over the shared medium. Relevant to our work is also the study presented in [1], which addresses the delivery and display of ads on social networks. There, however, the focus is on the optimal display strategy that maximizes the expected number of clicks on the ads. The intriguing finding of [1] concerns the importance of targeting ads. In particular, it emerges that ad dissemination strategies that take care of user interests can significantly increase the expected click rate.

In our paper, we study a different ad dissemination environment, i.e., a vehicular network. It follows that, unlike previous works, we address the targeting of ads mainly in terms of location and time. In particular, an ad may have different relevance in the time and spatial domains, e.g., a shop advertises a special offer that is of interest to potential customers travelling nearby in the time period during which the offer is valid. Thus, it is important to deliver the ad to users that are expected to pass by the ad *target area* during the *target time period*. While this is a challenge that holds for all mobile users, in the case of vehicular users the very fast temporal and spatial mobility dynamics make the problem more complex. A further challenge with respect to (wired) social network systems is represented by the transmission and display of ads to mobile users. First, ads have to be transmitted via wireless resources, which are scarce and costly. Thus, *which ads* to deliver, *when* and *where* (i.e., through which cellular base stations or WiFi access points (APs)) should be carefully determined. Second, only few ads at the time (possibly, even only one) can be shown to a vehicular user and the ad(s) should be displayed for a sufficiently long time in order to ensure that the vehicular user has the chance to look at them.

In our work, we consider a geographical area covered by APs and aim to solve the problem of which ads each AP should broadcast to passing-by vehicular users and when, so that we maximize the number of ads that are displayed to interested users within the ad target area and target time period. We describe the system model under study in Section 2 and we formulate the above optimization problem subject to some system constraints in Section 3. We then solve the problem and show numerical results in Section 4. Section 5 concludes the paper and discusses future work.

2 System model

Below we describe our reference scenario and how we model it, and introduce the adopted notations.

2.1 Model entities

Our scenario and model includes the following entities: the entity in charge of the advertisement service, the advertisements, the Wi-Fi APs that act as road-side units (RSUs), and the vehicles. Ads $a \in \mathcal{A}$ are the content items that the entity providing the service aim at showing to vehicular

users. 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. The RSUs provide, partial or full, coverage of the geographical service area of interest and can broadcast ads to vehicles. By \mathcal{R} we denote the set of RSUs.

A vehicle travelling across the service area is denoted by v while the set of all vehicles in the area is denoted by \mathcal{V} . The vehicles are equipped with an enhanced navigator software that can display ads if they are relevant to the user. Time is divided in time slot $k \in \mathcal{K}$, and space is divided in locations $l \in \mathcal{L}$. At every time slot, each of the vehicles travelling over the service area occupies exactly one location. Clearly, multiple vehicles can share the same location.

2.2 Parameters

We assume that the service entity can predict the position $\lambda(k, v) \in \mathcal{L}$ of each vehicle $v \in \mathcal{V}$ at each time slot $k \in \mathcal{K}$. Furthermore, we summarize the knowledge of the service entity by means of a set of binary flags $\gamma(l, r) \in \{0, 1\}$ (with $l \in \mathcal{L}$, and $r \in \mathcal{R}$) such that $\gamma(l, r) = 1$ if the RSU r covers location l , or in other words r can transmit ads to vehicle v whose position is $\lambda(k, v) = l$. Similarly, the binary flags $\tau(a, k, l) \in \{0, 1\}$ (with $a \in \mathcal{A}$, k is a time slot, and $l \in \mathcal{L}$) take the value 1 if location l and time slot k belong to, respectively, the target area and target time period of ad a and, hence, if it makes sense to display a at location l and time k .

Additionally, we know the *display time* $\delta(a)$ of ads, i.e., the number of time slots during which they have to be displayed. The maximum number of ads we can broadcast from each RSU at each time slot depends on the RSU capacity and the monetary budget available to the service entity; we denote such number by $\mu \in \mathbb{N}$.

Finally, the service entity knows whether the driver (or, passengers) of vehicle v is interested in ad a , through the binary parameter $\iota(a, v) \in \{0, 1\}$.

3 Problem formulation

The advertisement process takes three steps:

1. ads are *disseminated* by RSUs;
2. ads are *stored* by vehicles;
3. ads are *displayed* to the vehicle driver.

Each step corresponds to a binary variable and several constraints, as detailed below.

3.1 Ad dissemination

We define variable $x(a, k, r) \in \{0, 1\}$, expressing whether ad a is broadcasted from RSU r at time slot k . All passing-by vehicles can store the disseminated ads, but are not forced to do so, as explained next.

We cannot broadcast more than μ ads from each RSU at each time slot, i.e.,

$$\sum_{a \in \mathcal{A}} x(a, k, r) \leq \mu, \forall k \in \mathcal{K}, r \in \mathcal{R}. \quad (1)$$

As mentioned before, the above constraint can represent both technological limitations (e.g., maximum amount of bandwidth that can be devoted to broadcast transmissions at RSUs) and budget constraints due to the cost that mobile operators charge the service entity for ad broadcasting.

3.2 Ad storage

Variable $y(a, k, v) \in \{0, 1\}$ expresses whether ad a is inserted into vehicle v 's cache at time slot k . Clearly, the same ad cannot be stored at a vehicle more than once:

$$\sum_{k \in \mathcal{K}} y(a, k, v) \leq 1, \forall a \in \mathcal{A}, v \in \mathcal{V}. \quad (2)$$

Also, vehicles will not store ads their driver or passengers are not interested in, i.e.,

$$y(a, k, v) \leq \iota(a, v), \forall a \in \mathcal{A}, k \in \mathcal{K}, v \in \mathcal{V}. \quad (3)$$

Last, we can store an ad only if it has been disseminated to v , i.e., it has been broadcasted by at least one RSU covering v 's position. Thus, for any ad a , time slot k and vehicle v , we impose:

$$y(a, k, v) \leq \sum_{r \in \mathcal{R}} x(a, k, r) \gamma(\lambda(k, v), r). \quad (4)$$

Eq. (4) can be read as follows: we cannot store ad a (the y -value) unless, at the current time k , the location of vehicle v (the λ -value) is covered (the γ -value) by an RSU that is disseminating ad a (the x -value).

3.3 Ad display

Binary variable $z(a, k, v) \in \{0, 1\}$ expresses whether at time slot k ad a *starts* to be displayed to the driver of vehicle v . To begin with, no ad that is not stored in vehicle v can be displayed:

$$z(a, k, v) \leq \sum_{h=1}^k y(a, h, v), \forall a \in \mathcal{A}, k \in \mathcal{K}, v \in \mathcal{V}. \quad (5)$$

Also, ads are not shown unless the vehicle is within the ad target area during the target time period:

$$z(a, k, v) \leq \tau(a, k, \lambda(k, v)), \forall a \in \mathcal{A}, k \in \mathcal{K}, v \in \mathcal{V}. \quad (6)$$

In the above equation, notice that, given the definition of the binary flag, $\tau(a, k, \lambda(k, v))$ takes 1 if at time slot k vehicle v is within the target area and target time period of ad a , and 0 otherwise.

Furthermore, if we begin to display ad a at time k , then no other ad can be shown until the display time of a is over. Thus, for any ad a , time slot k and vehicle v , we have:

$$1 - \sum_{a' \in \mathcal{A}} \sum_{h=k-\delta(a')}^{k-1} z(a', h, v) \geq z(a, k, v). \quad (7)$$

By induction and starting from the first-displayed ad, it is easy to verify that the above constraint guarantees that ads are never interrupted once they start to be displayed.

Finally, each ad can be displayed at most once per vehicle, i.e.,

$$\sum_{k \in \mathcal{K}} z(a, k, v) \leq 1, \forall a \in \mathcal{A}, v \in \mathcal{V}. \quad (8)$$

3.4 Ad queue

It is often the case that we want ads to be displayed in the same order in which they arrive at the vehicle. However, as expressed by constraint (6), it makes sense to display an ad only when and if the vehicle is in the target area of the ad and if such an event occurs during the target time period of the ad. It follows that a first-input-first-output (FIFO) policy cannot be applied as the stored ads will be served not only based on their order of arrival at the queue but also on when the conditions expressed by the flag $\tau(a, k, \lambda(k, v))$ are verified. That said, in the presence of limited storage space, it is convenient to store an ad at a vehicle only if there is time to display the ad before the vehicle leaves the corresponding target area.

Let us define the memory space taken by the ads stored at vehicle v at time slot k as

$$N(k, v) = \sum_{a \in \mathcal{A}} \left[\sum_{h=1}^k (y(a, h, v) - z(a, h, v)) s(a) \right] \quad (9)$$

where $s(a)$ is the size of ad a . Then we impose that, at any time slot k and for each vehicle v , such quantity cannot exceed the available storage space S :

$$N(k, v) \leq S. \quad (10)$$

Also, let the auxiliary variable $Q(k, v)$ be the total duration of the ads stored by vehicle v at time k and not yet displayed:

$$Q(k, v) = \sum_{a \in \mathcal{A}} \left[\sum_{h=1}^k (y(a, h, v) - z(a, h, v)) \delta(a) \right]. \quad (11)$$

The expression in (11) is a summation, for each ad, of the difference between two terms: the first takes a value equal to 1 if the ad is in the vehicle cache; the second is equal to 1 if the ad is displayed. Queued ads are the ones for which the first term is one and the second is zero, i.e., they have been stored but not displayed yet.

Let us now denote by $\eta(a, v) \in \mathcal{K}$ the time at which vehicle v leaves the target area of ad a . Then, considering that $K = |\mathcal{K}|$, we can write the following constraint, forbidding to store ads that will not be shown:

$$Q(k, v) + k - 2K(1 - y(a, k, v)) \leq \eta(a, v) \quad \forall a \in \mathcal{A}, k \in \mathcal{K}, v \in \mathcal{V}. \quad (12)$$

In order to explain the above expression, let us see what happens if $y(a, k, v) = 1$, i.e., if the content is stored: the first term reduces to $Q(k, v) + k$, i.e., the time at which all ads that are currently stored by the vehicle will be displayed. Note that, as explained above, the queue policy is not a FIFO thus the time $Q(k, v) + k$ represents a worst case, i.e., the latest time at which, given the current queue status, ad a could be displayed. The second term is simply the deadline by which content a has to be shown; of course, this second term must be larger, i.e., later in time, than the first one. If this is not the case, then we cannot store the content, so we must set $y(a, k, v) = 0$. When this happens, the first term becomes $Q(k, v) + k - 2K$, with K being the latest time we consider in our simulation, i.e., the term on the left hand side of (12) becomes a very small quantity, certainly lower than $\eta(a, v)$. In summary, constraint (12) forbids vehicles from storing ads that cannot be shown by the time the vehicle leaves the associated target area.

3.5 Objective

The service entity is interested in maximizing the value of the ads that are disseminated, i.e., its objective function can be written as:

$$\max \sum_{a \in \mathcal{A}} \sum_{k \in \mathcal{K}} \sum_{v \in \mathcal{V}} \nu(a, v) z(a, k, v). \quad (13)$$

The above expression includes (i) the decision variable $z(a, k, v)$, which indicates if an ad is shown to a vehicle driver, and (ii) the input parameter $\nu(a, v) \in \mathbb{R}$, which is the revenue value that is expected from showing an ad to a vehicular user. Note that $z(a, k, v)$ depends on the decision variables $x(a, k, r)$ and $y(a, k, v)$, which express, respectively, whether an ad should be disseminated or not, and whether this should be stored at a vehicle. Also, observe that $\nu(a, v)$ may be different for each ad-vehicle pair as it accounts for the user profile and the relevance that the ad may have to the user.

The optimization problem has to be solved under all constraints reported in the previous section. The proposed formulation is an integer linear programming (ILP) problem, which we solved by using the Gurobi software tool [8]. Although the solution can be quickly obtained in the scenario that we address in our performance evaluation, we remark that a relaxed version of the problem can be efficiently solved when large-scale scenarios have to be considered.

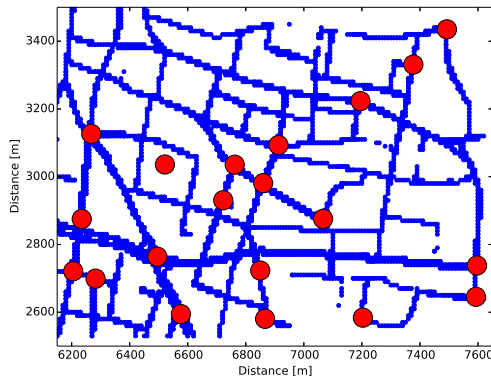


Figure 1: Road topology. Blue dots represent vehicles' paths while red dots represent RSUs

4 Numerical results

Simulations are run on a service area representing an $1 \times 1.5 \text{ km}^2$ 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. As shown in Fig. 1, 20 RSUs were deployed (red points in the figure) at the centers of random intersections. The total fraction of service area covered by RSUs is equal to 0.4. Each RSU is able to broadcast maximum 5 ads at each time slot. The total number of ads \mathcal{A} is set to 9 while the target area of

Table 1: Parameters of the ads categories.

	Class A	Class B	Class C
Target time period [minutes]	60	30	15
Target area size [km ²]	0.2	0.2	0.2

each ad is a circle with radius of 250 m, centered at a random point within the service area and such that it is all contained within the area. Each target area covers about 13% of the service area. The display time of each ad $\delta(a)$ is set to 63 s [3]. All ads have the same size $s(a) = 5$ MB. Ads are equally assigned to three categories: Class A, Class B and Class C. As reported in Table 1, all ads in these categories have the same target area size while they feature a different target period of 60, 30 and 15 minutes, respectively. In the simulation, time is divided in time slots of 30 s each, while space is divided into squared tiles, each with side of 50 m. Finally, the parameter S is set to 30 MB and ν is set to 1 for all vehicles and ads.




Figure 2: Randomly distributed user interests: (top) Total number of displayed ads per vehicle vs. vehicle density; (bottom) Average display probability for the different classes of ads and various values of vehicle density

In our simulation, we consider that only 20% of drivers are interested in each ad. Initially, we select such users randomly among all vehicles. The top plot in Figure 2 shows the average number of ads displayed to drivers normalized to the total number of vehicles present in the trace. Interestingly, even if normalized to the number of vehicles, this value increases as the vehicle density

grows. Indeed, the vehicular traffic congestion directly affects the time spent by vehicles under the coverage of an RSU and within the ads target area. In our simulations, we observed that the average time spent by a vehicle under an RSU is equal to 24.8 s for the lowest vehicle density, and it increases to 26.5 s and 30.95 s under medium and high vehicle density, respectively. This clearly has an impact on the system capability to deliver ads to the users and on the probability that they are relevant to them. In summary, the higher the vehicle density, the higher the vehicle dwell time in a given area and, thus, the chance to deliver and display ads to the users.

The bottom plot in Figure 2 shows the average display probability for different ad classes. For each class, this probability is given by the ratio of the number of displayed ads to the total number of vehicles interested in the ads belonging to that class. The probability of displaying an ad to an interested driver increases with the vehicle dwell time on the target area, hence with the vehicle density, as well as with the ad target time period. Indeed, short target time periods may not allow drivers to reach an ad target area on time, even if their vehicle was able to download the ad.

Next, we investigate the case where the vehicle interest in an ad, $\iota(a, v)$, is influenced by its expected route. Specifically, for each ad, we now select the interested users among those that are predicted to cross the target area of the ad during their trip. The obtained results, depicted in Figure 3, show a similar trend compared to the previous case. However, it is interesting to notice that the service entity is now able to significantly increase the number of displayed ads. This underlines the importance of predicting the level of user interests in location-dependent ads, hence the vehicles route.




Figure 3: User interest based on vehicle route: Total number of displayed ads per vehicle vs. vehicle density.

5 Conclusions

One of the recent trends of the automotive industry concerns the connection of cars with the communication infrastructure. Cars are becoming nodes of the Internet and automotive companies are rapidly unveiling their connected car development plans. The convergency between cars and Internet will enable a huge number of new opportunities and of new business models. In particular, in this paper we focused on an advertising service towards a (selected) set of cars, where the selection can be based on car location and user interests. We represented the dissemination process of ads as an optimization model that takes into account the additional challenges that advertising in vehicular environments poses with respect to the spreading of ads in other media. Our model indeed

accounts for the limitations of the wireless broadcasting technology and of the vehicles, as well as for the spatial and time relevance of ads to users and the fast dynamics of the scenario under study. Our results, obtained in a realistic scenario, show that vehicle mobility has a significant impact on the revenue that can be obtained from the advertising service. Furthermore, they highlight the importance of predicting vehicle mobility, hence the level of interest that users may have in location-dependent ads.

Acknowledgment

We thank Prof. Falko Dressler and Dr. Christoph Sommer for providing us with the Ingolstadt mobility trace. This work has been supported through the LIMPID project (POR FESR 2007/2013) funded by Regione Piemonte (Italy).

References

- [1] Z. Abbassi, A Bhaskara, and V. Misra. Optimizing display advertising in online social networks. In *24th International World Wide Web Conference (WWW 2015)*, Florence, Italy, 2015.
- [2] E. Altman and N. Shimkin. Strategic Posting Times over a Shared Publication Medium. In *International Conference on NETWORK Games CONTROL and OPTimization (NETGCOOP 2014)*, 2014.
- [3] M.C. Amarie, I. Pefkianakis, and H. Lundgren. Mobile Video Ad Caching on Smartphones. In *Proceedings of the 16th International Conference on Ubiquitous Computing, UbiComp '14*, pages 57–61, Seattle, WA, USA, 2014. ACM.
- [4] H. Beales. The value of behavioral targeting. *Network Advertising Initiative*, January 2010.
- [5] The Economist. The connected car: Smartphones on wheels. *The Economist: Technology Quarterly.*, Q3 2014, Sep 2014.
- [6] S. Goel, T. Imielinski, K. Ozbay, and B. Nath. Sensors on wheels – towards a zero-infrastructure solution for intelligent transportation systems. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems, SenSys '03*, 2003.
- [7] P. Gray. Is the car the next smartphone? *Mobile Enterprise*, February 2014. Available at <http://www.techrepublic.com/blog/mobile-enterprise/is-the-car-the-next-smartphone/>.
- [8] Gurobi Optimization Inc. Gurobi Optimizer Reference Manual, Version 6.0. <http://www.gurobi.com>, 2015.
- [9] Z. Lotker, B. Patt-Shamir, and M. R. Tuttle. A game of timing and visibility. *Games and Economic Behavior*, 62(2):643–660, 2008.
- [10] D. Quercia, R. Schifanella, and L. M. Aiello. The shortest path to happiness: recommending beautiful, quiet, and happy routes in the city. In *25th ACM Conference on Hypertext and Social Media, HT '14, Santiago, Chile, September 1-4, 2014*, 2014.

- [11] A. Reiffers-Masson, Y. Hayel, and E. Altman. Game theory approach for modeling competition over visibility on social networks. In *Sixth International Conference on Communication Systems and Networks, (COMSNETS)*.
- [12] Nurun 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/>.