The Role of Parked Cars in Content Downloading for Vehicular Networks

The Role of Parked Cars in Content Downloading for Vehicular Networks

F. Malandrino[†], C. Casetti[†], C. F. Chiasserini[†], C. Sommer[‡], F. Dressler[‡]

[†] Politecnico di Torino, Italy

[‡] University of Innsbruck, Austria

Abstract

When it comes to content access using Inter-Vehicle Communication (IVC), data will mostly flow through Road Side Units (RSUs), deployed in our cities. Unfortunately, the RSU coverage is expected to be rather scattered. Instead of relying on RSUs only, the paper investigate the possibility of exploiting parked vehicles to extend the RSU service coverage. Our approach leverages optimization models aiming at maximizing the freshness of content that downloaders retrieve, the efficiency in the utilization of radio resources, and the fairness in exploiting the energy resources of parked vehicles. The latter is constrained so as not to excessively drain parked vehicle batteries. Our approach provides an estimate of the system performance, even in those cases where users may only be willing to lease a limited amount of their battery capacity to extend RSU coverage. Our optimization-based results are validated by comparing them against ns-3 simulations. Performance evaluation highlights that the use of parked vehicles enhances the efficiency of the content downloading process by 25%-35% and can offload more than half the data traffic from RSUs, with respect to the case where only moving cars are used as relays. Such gains in performance come at a small cost in terms of battery utilization for the parked vehicles, and they are magnified when a backbone of parked vehicles can be formed.

I. Introduction

After more than 10 years of research in the field of Intelligent Transportation Systems (ITS), a consolidation phase can be observed. This is mainly supported by the success in standardization efforts both for the underlying communication protocols such as DSRC/WAVE (IEEE 802.11p/IEEE 1609.4 [1]) as well as for higher layer applications (e.g., ETSI cooperative awareness messages [2]). Given the multitude of emerging applications in the ITS domain, emphasis is currently given to the integration of heterogeneous Inter-Vehicle Communication (IVC) techniques. As nicely described by Willke et al. [3], applications range from time-critical safety information disseminated using short-range communication to entertainment applications primarily using cellular networks. In this spectrum, the

focus of this paper is on content downloading from centralized systems to moving vehicles [4]–[6]. In previous work, it was shown that this can also, and very efficiently, be supported using short-range communication techniques, instead of relying on 3G/4G networks only. Content downloading has become the primary application in Delay Tolerant Networks (DTNs) [7], [8]. It is obvious that vehicular networks will similarly benefit from this service primitive.

Content downloading in the scenario of this paper occurs from a central entity to moving vehicles using RSUs, frequently also called Access Points (APs). RSUs are placed along the streets following the ideas presented in [4]. Unfortunately, the RSU coverage can be expected to be rather limited due to two main reasons. First, in early stages of deployment only a few RSUs will be available. Secondly, also from a business perspective, the number of RSUs will always be rather limited if other alternatives may be used. This is mainly due to high upfront costs and considerable maintenance overhead. As a novel concept, we therefore introduce the use of *parked vehicles* to support content downloading. For a motivation and related work, please refer to Section II.

Previous work on content downloading in vehicular networks dealt with individual aspects of the process, such as the deployment of Roadside Units (RSUs) [6], [9], the performance evaluation of IVC [10], or the network connectivity [11], [12]. In [4], the actual potential of IVC-based content downloading was quantified.

This paper, which builds upon [5], aims at identifying the degree at which the use of parked vehicles helps in this process and what the associated costs are. As novel components with respect to previous work, battery drainage was incorporated in the model. This leads to observing the vehicular network in a more holistic way, also accounting for heterogeneous car batteries and their different draining process. This is in line with findings in other domains such as the network lifetime metric in sensor networks [13]. Finally, a simulation-based model validation of the model was added, along with several new results and dynamic scenarios where parked vehicles come and go at random times.

The approach taken in this work assumes the availability of mobility data on a defined area, including the location of existing RSUs. The underlying assumptions are the ability to predict mobility to a certain extent and the introduction of simplified physical and medium access control (MAC) layers. Through such assumptions, it is possible to estimate the energy consumption of vehicles engaged in content transmission and reception. In this framework, we cast and solve a data scheduling optimization problem, to carry out a qualitative study and identify an upper bound to system performance. The analysis is validated by showing that the bound is very close to the results obtained through ns-3 simulations, which use very detailed models at physical and MAC layers and imperfect knowledge of vehicle movements.

The key contributions of this paper can be therefore summarized as follows:

- To show the potential of using parked vehicles in addition to RSUs for content downloading to moving vehicles. The results confirm that the use of parked vehicles greatly benefits the content downloading process.
- To show that, even considering heterogeneous battery capacities of parked vehicles, content download is greatly enhanced by the use of parked cars.
- To carefully study the impact of the number of parked vehicles, their ability to create a dynamic backbone infrastructure, and to spend a certain amount of energy in support of content download, within the constraints of the available battery capacity.

The rest of the paper is organized as follows. After reviewing related work in Section II, we present the system scenario and the graph-based representation of the network that we use for system optimization in Section III. Our metrics of interest are then introduced in Section IV, while the optimization problems maximizing such metrics are presented in Section V. We show the benefits of involving parked vehicles in the download process in Section VI, in the different cases where the main metric of interest varies among content freshness, channel utilization, and fairness in using parked vehicles. Finally, in Section VII we draw our conclusions.

II. Related Work

Parked vehicles have been shown to perfectly complement driving vehicles and locally installed RSUs to support IVC [14]–[16]. As the authors of [15] argue, during the first years of Dedicated Short Range Communication (DSRC) deployment, the number of vehicles that are equipped and can help relay information to others will be very low – as will the number of available RSUs because of the high CapEx [17]. Taken together with findings from a study conducted in the area of Montreal [18] pointing out that roughly 70% of vehicles are parked for an average of 23 hours a day [19], we can see both the need and the potential for incorporating parked vehicles in any content downloading scheme. The same study found that only 3.7% of vehicles were found to be parking indoors, leaving others to act as what can conceptually be described as a set of dynamic Stationary Support Units (SSUs) [20].

This begs the question of how likely it is that parked vehicles can participate in a vehicular network. Technically, this can easily be solved: just like their moving counterparts, parked vehicles are already equipped with DSRC devices. Modern cars also come pre-equipped with dedicated electronics to keep certain devices powered on when the vehicle is not running – and cutting power to these devices when the battery charge drops below a certain point.

The impact of DSRC operation on battery drain is minimal. As a worst-case analysis, based on specifications of early prototypes and assuming that no energy management scheme is employed, a typical DSRC Onboard Unit (OBU) will not drain more than 1 W on average and market-ready OBUs can be expected to drain even less. The battery of a small vehicle provides about 40 Ah to 80 Wh [21], the battery of a Tesla Roadster 4400 Ah, thus allowing the OBU to run for anything between 20 days and several years. Smart energy management schemes have been presented in the literature [22] to further increase the lifetime manyfold. Thus, it can be argued that drivers will tolerate their parked vehicle to participate in the vehicular network. Further, the success of social networks and of crowd-sourcing activities demonstrates the general willingness to share information for mutual benefit. Finally, owners might not have a choice whether their vehicles participate in IVC since, not only because of their substantial safety benefits, DSRC devices might become mandatory in the near future [23].

The participation of parked vehicles to support different applications of ITS has been proposed by several other authors. None of them, though, addresses energy aspects as we do.

Liu et al. [14] presented a method to use parked vehicles as relay nodes to disseminate information in a DTN fashion. They mainly focus on connectivity and show that an ITS can greatly benefit from additional nodes.

Eckhoff et al. [15] investigated communication latency with a special focus on traffic safety, showing that parked vehicles can help improve safety when used as relay nodes and cope with radio shadowing by routing around obstacles in an urban environment.

Crepaldi et al. [16] expands on the discussion of parked cars as relay nodes to further investigate their usefulness; they propose that parked vehicles can be used to share and provide opportunistic Internet access to other vehicles.

III. System Model

Our starting point is a vehicular mobility instance, i.e., an area where we assume to know the location and movements of cars, their IVC capabilities and the location of RSUs. IVC-capable vehicles communicate with each other, as well as with RSUs, and thus form a dynamically-changing network topology. For the sake of realistic representation of such a network, we use traces that were generated using the Veins simulation framework and in particular the SUMO simulator [24]. Beside car positions, the traces also report the SINR with which each cars receives beacon transmissions from other cars.

A. The Time-Expanded Graph

We represent the network and its temporal and spatial dynamics as a Time-Expanded Graph, or TEG [4]. Specifically, time is discretized into steps, by adopting the same granularity used in the given mobility instance. Then, at any given time step, each network node (RSUs or vehicles) is represented as a vertex in the graph, and any existing wireless links as directed edges between vertices.

Each edge is associated to a finite weight, representing the network-layer data rate that can be achieved by transmitting over the link at the corresponding time step. In the following, the data rate is assumed to be a function of inter-node distance (although values based on experimental results will be used). Also, the terms link and edge are as regarded as synonyms. For brevity, the links from RSU infrastructure to moving (parked) vehicles are denoted as I2M (I2P); the links from a moving (parked) vehicle to another as M2M (P2P); the links from parked to moving vehicles as P2M.

If an edge connects vertices that model the same RSU, or parked vehicle, over time, it represents the possibility that such nodes store data for a given period. Edges modelling the same moving node over time allow us to represent the possibility that a vehicle physically carries data during its movement. Since storage capabilities of any node are assumed to be significantly larger than the content size, all edges of this "storing" type are associated to an infinite weight. Also, note that modelling the duration of contacts between network nodes, instead of considering them as atomic, facilitates the accounting for channel contention (see [4] for further details).

Finally, we model the server(s) (from which RSUs retrieve the data) as a vertex named α. Edges of infinite capacity are drawn from α to any vertex representing an RSU. The graph is completed by a virtual vertex ω to which all downloaders are connected by edges of infinite capacity. These two vertices are a practical device that will allow us to represent the source and destination of the total flow over the graph.

The TEG is therefore a weighted directed graph, representing the temporal evolution of the network topology and the flow of data through it. Figure 1 illustrates a sample instance of the TEG, where downward arrows are used to represent the temporal evolution of network nodes, while horizontal arrows indicate data flows (different line types correspond to different link types). Using the TEG model, the performance of content downloading services for vehicle users can be studied.

B. Content downloading

We assume that a set of n content items is available and that a moving vehicle becomes interested in either downloading an item for the first time, or in refreshing an already cached item, according to

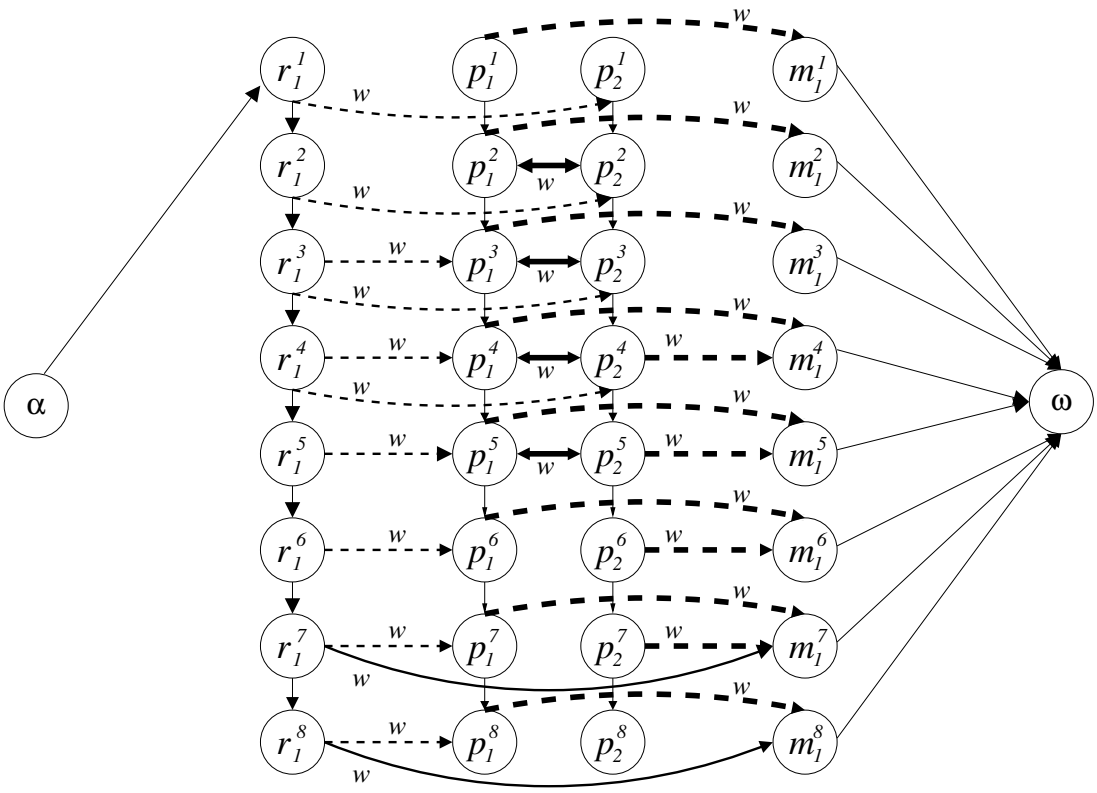

Figure 1. TEG sample: *r* represents the RSU, *p* are parked vehicles, *m* represents a moving vehicle, *w* represent edge weights; apices refer to time steps, pedices differentiate nodes of the same type.

a Poisson distribution with rate λ. We also assume that user interest lasts for a constant time period T. We refer to a vehicle that is engaged in the data retrieval process as a *downloader*.

Each content is updated at the server at time instants that are Poisson-distributed with rate ρ (the extension to time-varying update rates as well as to content-depending rates is straightforward). An updated content item is called a *version* of the content. The content version size is denoted by $S(c)$ ($c = 1, \ldots, n$), and, along with the weight of the edge representing a link between any two nodes, it defines the content transfer time over that link. If the link fails, an interrupted transfer of a content version can be resumed through other links that may become available.

The choice of the Poisson distribution (as opposed, e.g., of any power-law distribution) for content request is motivated by the general settings in which users are supposed to be active. Since content is not geolocalized, the purpose of this choice is to provide an assessment of the performance without specialized assumptions on the content request process, which are left for future work. The model, it is to be noted, can accommodate any arbitrary distribution.

IV. Performance Metrics

In the framework outlined above, three main metrics of interest are defined. These will be the subject of optimization models later on.

Content freshness. This metric is defined as the average number of new versions of content that each user manages to download throughout its trip across the road topology. More specifically, consider the generic downloader d and assume that d is currently interested in content c; then, the freshness associated to such a user is defined as:

$$\phi_{d,c} = \frac{1}{S_c} \sum_{k \in \mathcal{W}_{d,c}} \sum_{v \in \mathcal{V}_c^k} \sum_i (v \cdot x_k^v(i,d)) \tag{1}$$

where

- $\mathcal{W}_{d,c}$ is the set of time steps at which d is interested in content c,
- $\mathcal{V}_c^k$ is the set of all versions v of content c that have become available before time step k,
- i denotes the node (RSU or vehicle) such that a link from i to d exists at time step k, and
- $x_k^v(i,d)$ is the amount of data bytes, belonging to a version v, that can be transferred on such a link.

The system freshness is then obtained by summing (1) over all downloaders and content items, i.e., as $\sum_d \sum_c \phi_{d,c}$.

Radio resource utilization. The fraction of idle time observed over the channel is taken as a measure of the efficiency in radio resource utilization during data downloading. To sample such a quantity, we identify a number of "check points" across the road layout, which are located in a regular grid with adjacent points 100 m from each other. At each of these locations, a "fake" node (in our model) is purposely added. Its task is to sense the wireless channel and compute its utilization as seen at that position. I_p ($0 \leq I_p \leq 1$) denotes the fraction of channel idle time observed at check point p, and $(1 - I_p)$ indicates the channel load there.

Drainage of parked vehicle batteries. The amount of energy that parked vehicles consume while acting as relays should be limited; specifically, it should represent a small fraction of the battery capacity available at each vehicle. In order to compute the amount of energy consumed by a parked vehicle in transmitting/receiving data traffic, we proceed as follows. The amount of data transmitted and received by a generic parked vehicle is translated into transmission and reception times, respectively. Let us consider a parked vehicle i communicating at frame k with node j; then, the amount of data handled by i is given by $\sum_c x_k^c(i,j)$ and the communication data rate over the (i,j) link is $w_k(i,j)$. Note that the latter parameter is known, since it can be derived from the SINR included in the mobility trace.

We then introduce a constant factor H to account for protocol headers and the factors A^{tx} and A^{rx} ($A^{tx}, A^{rx} \geq 1$) that represent the average number of, respectively, transmission and reception attempts that are needed before a packet is correctly received. Then, at frame k, the fractions of time during which vehicle i transmits and receives can be written as:

$$T_k^{tx}(i) = A^{tx} \cdot (1+H) \cdot \sum_j \sum_c \frac{x_k^c(i,j)}{w_k(i,j)}, \tag{2}$$

$$T_k^{rx}(i) = A^{rx} \cdot (1+H) \cdot \sum_j \sum_c \frac{x_k^c(j,i)}{w_k(j,i)}. \tag{3}$$

Clearly, within step k the fraction of time that i spends in idle state (i.e., without either transmitting or receiving) is given by $(1 - T_k^{tx}(i) - T_k^{rx}(i))$.

Next, given the power consumed in transmit, receive and idle state, namely, P^{tx}, P^{rx} and P^{idle}, the total energy consumed by vehicle i is obtained as:

$$\pi_i = \sum_k \left(T_k^{tx}(i) P^{tx} + T_k^{rx}(i) P^{rx} + (1 - T_k^{tx}(i) - T_k^{rx}(i)) P^{idle} \right) \tag{4}$$

where the sum is over all frames k's. Supported by recent experimental campaigns, the values of P^{tx}, P^{rx} and P^{idle} are set to 1.6 W, 1.2 W and 0.15 W respectively. These parameters are deliberately not representative of a specific wireless card, but can be found to be in the range of those in [25]. Also, from simulation results (which will be described in Section VI-B), we found $H = 0.05$, $A^{tx} = 1.03$, $A^{rx} = 1.02$.

V. Optimizing the system performance

We now formulate three optimization models that aim at maximizing, respectively, the content freshness, the efficiency in radio resource utilization and the fairness in exploiting the battery resources of parked vehicles. The models also take into account specific system constraints that will be illustrated later.

A. Maximizing the content freshness

The first objective is the maximization of the freshness over all downloaders and all content items, i.e.,

$$\max \sum_d \sum_c \phi_{d,c} \,. \tag{5}$$

In the TEG model, the problem in (5), amounts to maximizing the flow of data bytes belonging to a content version from α to ω, for a given set of versions. Such a max-flow problem needs to be solved taking into account several constraints, briefly outlined below.

- The flow on every existing edge of the TEG must be greater than or equal to zero.
- The amount of incoming flow in every vertex of the TEG must equal the amount of outgoing flow.
- Given a specific vehicles, none of the following events can take place simultaneously: the vehicle transmits to a neighboring vehicle; a neighboring vehicle receives from any relay; the vehicle receives from a neighboring relay; a neighboring relay transmits to any vehicle; the vehicle receives from a neighboring RSU.

The problem falls in the category of mixed integer linear programming (MILP) problems. Its solution is quite standard, and follows the guidelines set forth in [4]. The result is the optimal scheduling of data traffic over all existing links. In other words, it is possible to derive the value of the variables $x_k^v(i,j)$ representing the amount of data to be transferred over the network links at any time step.

The problem in (5) can also be solved under constraints limiting the drainage of the parked vehicles batteries. In particular, we can associate to each parked vehicle i a value Π_i, expressing the maximum amount of energy that i can consume to relay traffic to others. Such an amount is typically bounded by a fraction of i's battery capacity (which is between 35 and 70 Ah at 12 V). Thus, for every parked vehicle i, the following constraint is added:

$$\pi_i \leq \Pi_i. \tag{6}$$

B. Maximizing the radio efficiency

Next, a target value of freshness, Φ, is set, and the aim becomes the minimization of the medium usage to reach such target (i.e., the maximization of radio efficiency). Specifically, the following non-linear optimization problem can be posed:

$$\min \sum_p (1 - I_p)^2 \tag{7}$$

$$s.t. \sum_d \sum_c \phi_{d,c} \geq \Phi. \tag{8}$$

The sum in the objective function is over all selected check points at which the channel idle time is observed (i.e., at the "fake nodes" included in the TEG). Conversely, the sums in (8) are over all downloaders and all content items.

Remarkably, choosing a quadratic objective function leads to a solution that reduces the channel load, i.e., $(1 - I_p)$ mainly in the most congested zones. More formally, the property below can be proven.

Property 5.1: When a choice exists, the objective (7) reduces the channel load in those areas where the load is higher.

Proof: Assume that there are two points p_i, p_j, where the current channel load is $(1 - I_{p_i,curr}) > (1 - I_{p_j,curr})$. Also, assume that the load at either p_i or p_j can be reduced by an amount ϵ. If the load is reduced at p_i, then the objective improves by $(1-I_{p_i,curr})^2 - ((1-I_{p_i,curr}) - \epsilon)^2 = 2(1-I_{p_i,curr})\epsilon - \epsilon^2$. If the load is reduced at p_j, the objective improves by $2(1-I_{p_j,curr})\epsilon - \epsilon^2$. Since $(1-I_{p_i,curr}) > (1-I_{p_j,curr})$, then the former improvement is higher than the latter, and therefore the traffic load at point p_i will be reduced, which proves the property. ■

As a final remark, the function in (7) is convex, thus the problem can be solved in polynomial time using any standard commercial solver. Similarly to what was done before, constraints such as (6) can be imposed so as to make sure that the energy resources of parked vehicles are not over-exploited.

C. Evening the drainage of parked vehicles batteries

Taking the viewpoint of parked vehicles, it is conceivable that battery drainage of such vehicles should be evenly distributed. Specifically, given a target value of content freshness, the corresponding load (i.e., the amount of data sent or received) should be shared as evenly as possible among parked vehicles. Similarly to (7), we take a quadratic function of the energy consumption of parked vehicles and formulate the following problem:

$$\min \sum_i \pi_i^2 \tag{9}$$

$$s.t. \sum_d \sum_c \phi_{d,c} \geq \Phi. \tag{10}$$

It can be remarked that Property 5.1, proven for the channel utilization, also holds for the energy consumption of parked vehicles. Hence, the quadratic formulation leads to a *fair* situation, in the sense that it reduces the energy consumption of the most exploited parked vehicles. Furthermore, the objective in (9), involving variables π_i, is independent of constraint (6). Thus, the optimization problem could be set also including constraint (6), no matter which values of Π_i are imposed.

D. Complexity analysis

Before we proceed with our performance evaluation, we give a more comprehensive look at the size of the TEG we build, and hence to the complexity of performing optimization on it.

Let us start with the number of vertices in the TEG. For each time frame k, we have one vertex for each RSU, parked vehicle and mobile vehicle; furthermore, we need to add the virtual vertices α

and ω. The number of vertices is thus at most $2+K(R+M+P)$, where K, R, M, P are the number of frames, RSUs, mobile and parked vehicles respectively. Notice that such a number is an upper-bound. Mobile vehicles are not always present on the topology, so each of them will be associated to less than K vertices.

As far as edges are concerned, we have to consider:

- KR edges from α to RSUs;
- KM edges from mobile vehicles to ω;
- at most $K(R+M+P)$ "vertical", inter-frame edges;
- at most $K(R(M+P)+(M+P)^2)$ "horizontal", intra-frame edges.

In the last item, we account for the fact that RSU don't communicate between each other while, in principle, all mobile and parked vehicles can communicate with each other. All summed, the number of edges is $K(2R+2M+P+R(M+P)+(M+P)^2)$.

Assuming that $M > P > R$, we can say that the number of vertices is $O(KM)$ (linear in the number of vehicles), and the number of edges is $O(KM^2)$ (quadratic). Notice that these are worst-case complexities, and real-world traces give substantially lower results. Even in the worst case, a number of edges (and thus of constraints) that grows quadratically with the number of vehicles is manageable with current optimization algoritms; we can therefore conclude that our model is suitable for real-world and real-time usage.

VI. Performance Evaluation

Below, the reference scenario used to derive the performance results is briefly outlined. Then, content freshness and channel idle time are computed by solving the optimization problems introduced in (5), (7) and (9), under the following conditions: (i) an exact mobility prediction is available, and (ii) simplified assumptions on the physical and MAC layers hold. These results therefore represent an upper bound to the system performance. In order to validate such analysis, some of the above results are compared to simulation results obtained via ns-3. In simulation, the same optimal scheduling of data traffic found by the model is used, but physical and MAC aspects are accurately modeled and the prediction may not match the actual mobility of cars.

A. Reference scenario

Our reference scenario is that of [15], accurately modeling road and lane topology, building layout and parking places, as well as vehicular traffic patterns in a 1.5×1 km^2 section of the urban area of Ingolstadt in Germany. As previously mentioned, the mobility traces were generated using the Veins

simulation framework and the SUMO simulator [24]. The scenario then has been carefully validated through measurement data of road traffic flows in the city of Ingolstadt, distributing parked cars according to satellite imagery [15].

The scenario models 963 vehicles travelling over the road topology, with an average trip time of 388 s and an average vehicle density of 69.28 vehicles/km^2. There are also 80 parked vehicles, which are present at each time step. We fix the number of available RSUs to 10. Figure 3 shows a snapshot of the road layout, including the positions of parked (in black) and moving (in grey) vehicles, as well as RSUs (in red).

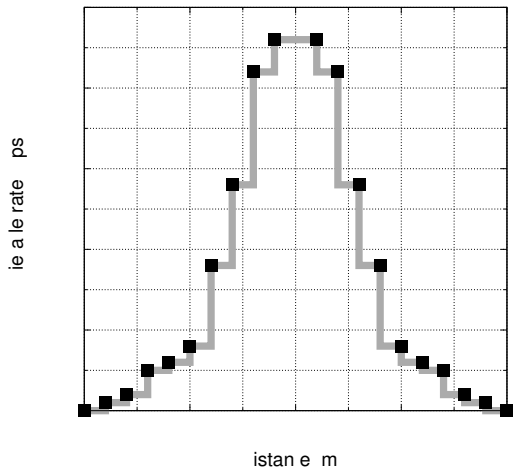

Figure 2. Characterization of the achievable network-layer rate as a function of distance, based on experimental data

All nodes use the IEEE 802.11p technology. The value of the achievable network-layer rate between any two nodes is adjusted according to the distance between them. To this end, we refer to the experimental results in [10] to derive the values shown in Figure 2, and we use them as samples of the achievable network-layer rate. Note that the maximum node transmission range is limited to 200 m, since, as stated in [10], this distance allows the establishment of a reliable communication in 80% of the cases.

Finally, for clarity, one content item only (i.e., $n = 1$) is considered, and the following settings are chosen: $\lambda = 0.005$, $\rho = 0.02$, $T = 30$ s, and $S = 10$ Mbytes. Unless otherwise specified, it is assumed that $\Pi_i = \infty \, \forall i$.

B. Optimal data scheduling

As mentioned above, we provide an upper bound to the system performance, by assuming that an exact mobility prediction is available, which is given by the aforementioned vehicular trace. Also, in

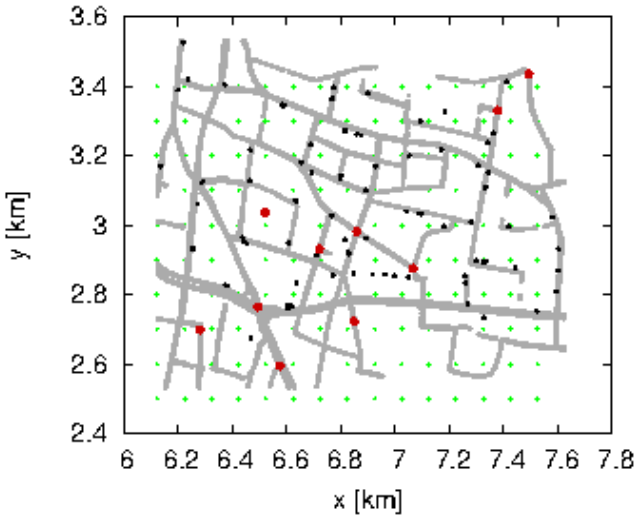

Figure 3. Road topology. Grey dots represent mobile vehicles, black dots represent parked vehicles, red ones represent RSUs. Green crosses mark the "check points".

order to easily solve the optimization problem and carry out our qualitative study, the physical and MAC layer aspects are represented in a simplified manner, through the TEG network model.

The baseline scenario is a network where all relays are mobile and no parked vehicles are enabled; thus, only I2M and M2M links can be exploited. Then, (part or all of) the parked vehicles present in the Ingolstadt trace are included and the optimization problem is solved by exploiting links also involving such vehicles. More precisely, we consider two different ways to use parked vehicles. In the former, referred to as *no backbone*, parked vehicles can receive and transfer data using, respectively, I2P and P2M only. Instead, in the latter, referred to as *with backbone*, parked vehicles can form a backbone by exploiting P2P links among themselves, i.e., data can flow from an RSU, through one or more parked vehicles, and be eventually delivered to downloaders. Describing algorithms and protocols for the backbone creation, however, is outside the scope of this paper.

Furthermore, with the aim to investigate the impact of the available M2M links, for each of the above scenarios, the optimization problems are solved over different graphs, obtained by progressively removing existing M2M links. Specifically, the problems over the full graph are solved, then *vertices* are randomly removed, so as to represent vehicles that do not request content updates but that can only relay traffic for others. Such vertices are removed till a fraction R of the original relay traffic is phased out. For a given value of R, the optimization problems are solved again, using the corresponding pruned graph.

Maximizing the content freshness. By using the TEG, we first maximize the content freshness,

i.e., the formulation in (5). The level of content freshness that is attained is shown in Figure 4 as R and the fraction of parked vehicles that can be used vary.

In the plots, the curves labeled by "no park" refer to the baseline scenario. As expected, increasing the number of involved parked vehicles and building a backbone positively affect freshness, showing the important role of parked vehicles in content downloading. In particular, when few parked vehicles are exploited, removing M2M links (i.e., $R = 0.5$) impairs the performance; however, when more than 20% of parked vehicles are involved, they easily make up for the unavailability of mobile relays.

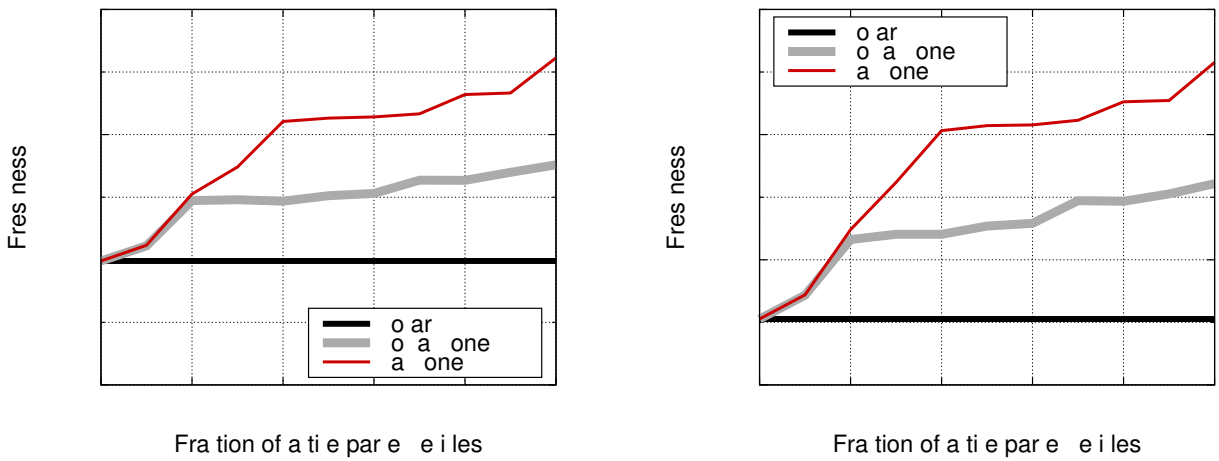

Figure 4. Freshness as a function of the maximum time period mobile vehicles can be used as relays, for $R = 0$ (left) and $R = 0.5$ (right). Parked vehicles do not change position over time.

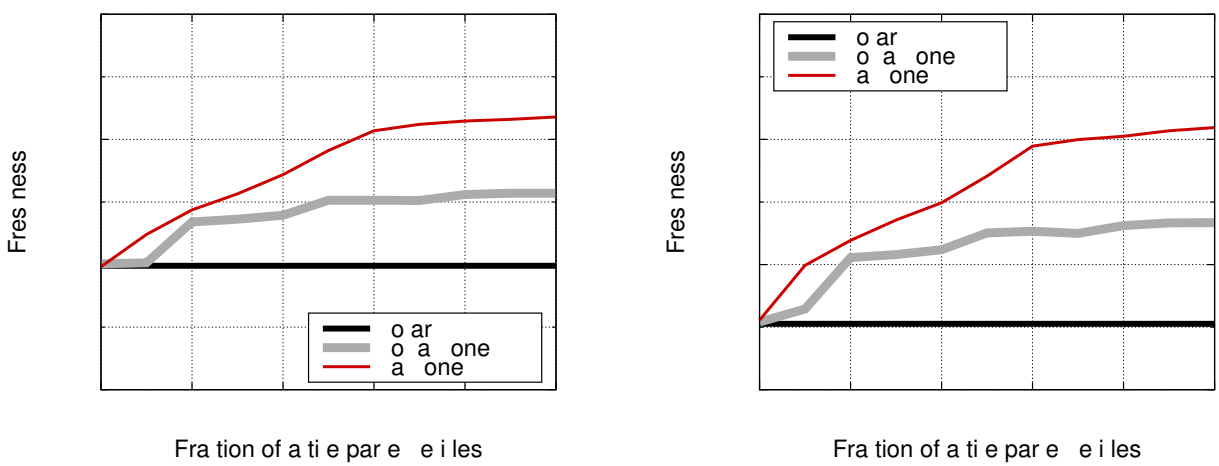

Figure 5. Freshness as a function of the maximum time period mobile vehicles can be used as relays, for $R = 0$ (left) and $R = 0.5$ (right). Parked vehicles come and go at random times.

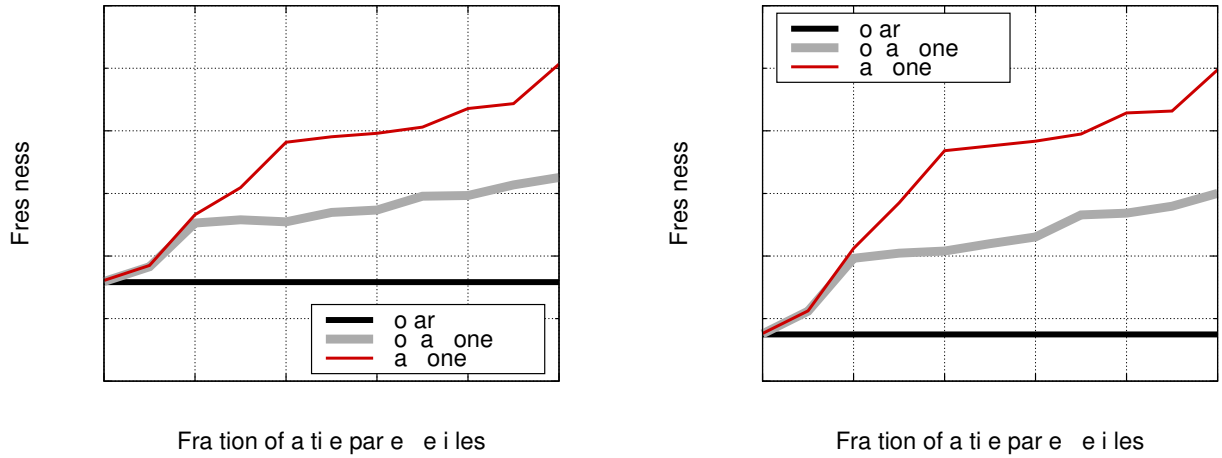

Figure 6. Freshness as a function of the maximum time period mobile vehicles can be used as relays, for $R = 0$ (left) and $R = 0.5$ (right). Parked vehicles do not change position over time; the retransmission rate is $A^{tx} = 1.05$.

In Figure 5, we study the freshness obtained when parked vehicles are not static, but come and go at random times. Two effects can be observed: first and most obviously, the freshness is lower than in the static case (Figure 4). More interestingly, allowing parked vehicles to form a backbone mitigates such an effect.

Transmission errors. In Figure 5, we have assumed that all packets are successfully delivered. However, in general this is not the case. We now set a fairly high retransmission rate $A^{tx} = 1.05$, and observe its impact on content freshness.

As we can see from Figure 6, the need for retransmissions directly translates into a lower freshness. It is also worth to notice, by comparing Figure 6 with Figure 4, that such an effect is more evident in the "no park" case; using parked vehicles ("no backbone") and allowing them to form a backbone ("with backbone") reduces the impact of packet losses. Similarly, increasing the time each parked vehicle can be used has a positive impact.

Where data come from. The plots in Figure 7 depict the amount of traffic that downloaders receive directly from RSUs, other moving vehicles, or parked vehicles. Due to the limited impact of R, here and in the following plots (unless otherwise specified) we consider that all M2M links can be exploited (i.e., $R = 0$).

Consistently with the previous figure, the plots show that when parked vehicles used, the amount of data that such vehicles manage to deliver towards downloaders increases significantly (especially when both P2M and P2P are involved). Also, as more parked vehicles can be used, less and less data traffic is originated from RSUs and the impact of mobile relays is significantly reduced.

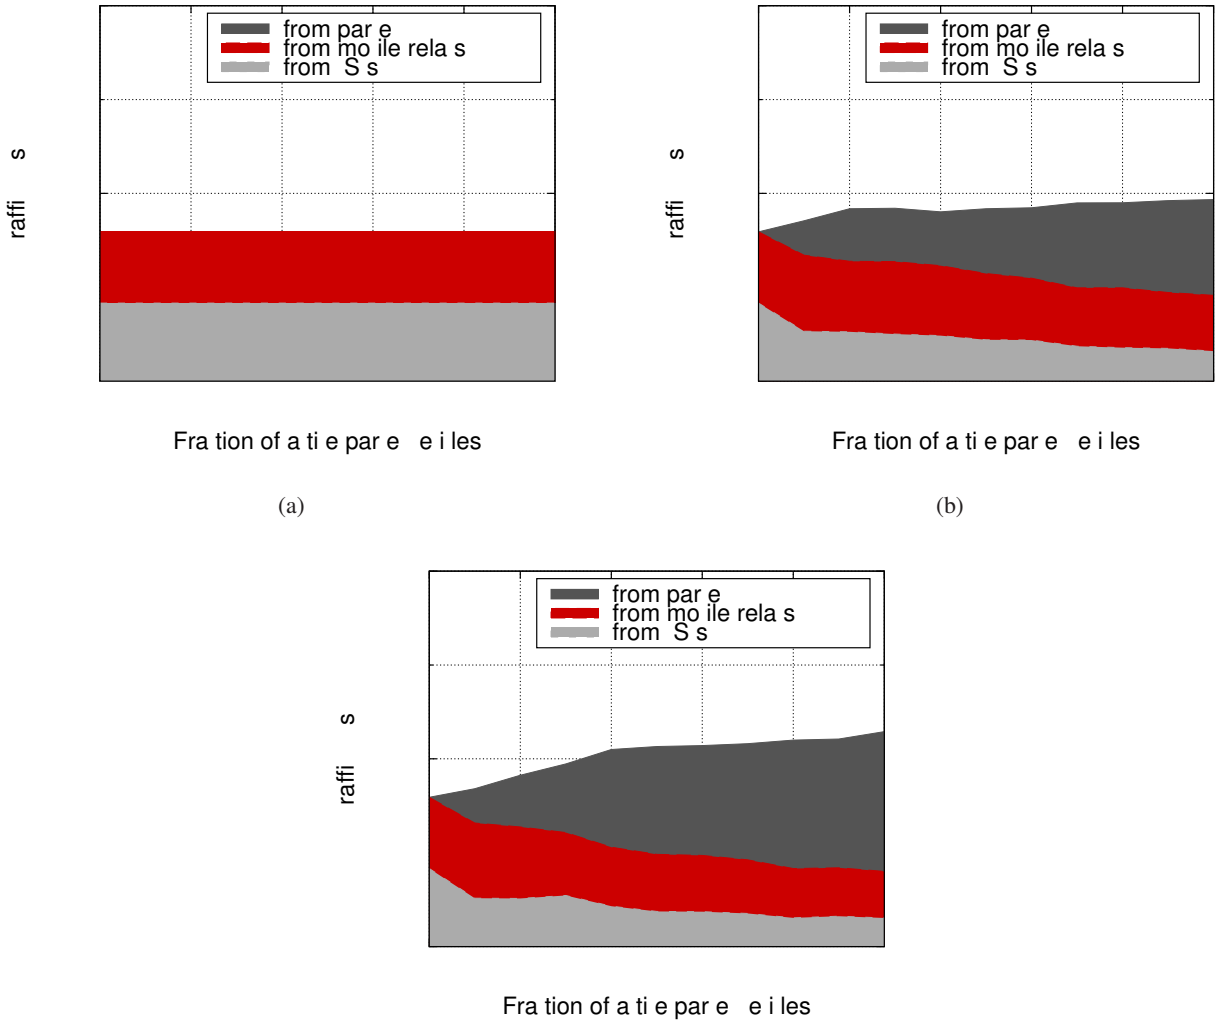

Figure 7. Traffic to downloaders for the cases "no park" (a), "no backbone" (b), and "with backbone" (c).

Figure 8 presents the amount of data transmitted by the RSUs, split into three parts: data towards parked vehicles, mobile relays and downloaders. Comparing this figure to Figures 4 and 7, it is clear that parked vehicles storing an information item (or part of it) are more effective relays than moving vehicles. Indeed, since they remain on the road topology much longer, they provide the same data to more than one downloader. It follows that the load on RSUs, as well as on mobile relays, can be greatly reduced while maintaining very good performance in terms of content freshness.

Figure 10 shows the impact on the content freshness of imposing a finite value of $\Pi_i = \Pi \, \forall i$ in constraint (6). Clearly, if $\Pi = 0$, we fall back to the "no park" scenario. By looking at Figure 10, it can be seen that the smaller the amount of data traffic that each parked vehicle can handle, the larger the amount of data that is routed through RSUs and mobile relays. However, a constraint as high as

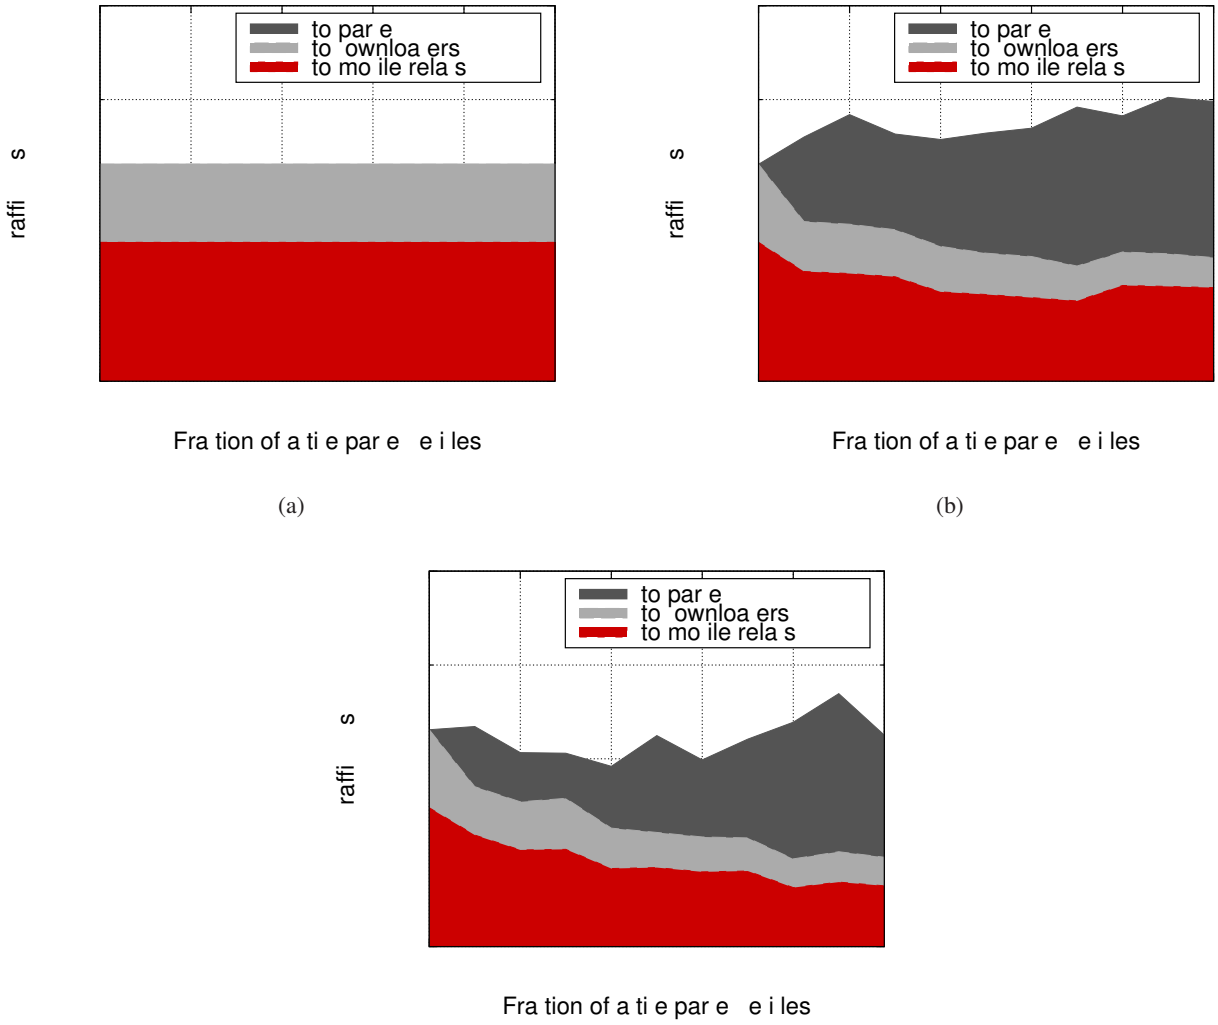

Figure 8. Traffic from RSU for the cases “no park” (a), “no backbone” (b), and “with backbone” (c).

$\Pi = 40$ J per parked vehicle during our observation period – corresponding to 1800 J per night – is enough to achieve the same value of content freshness as the one obtained under $\Pi = \infty$. We stress that an energy consumption of 1800 J corresponds to a very small fraction of battery discharge, as typically a fully-charged car battery has a capacity of 40 Ah, i.e., 1,728,000 J.

Finally, we validate our optimization-based model through realistic simulations. The optimization framework has been integrated in the ns-3 simulations as follows:

(i) vehicles periodically notify a central authority about the route they intend to take;

(ii) based on the current speed over the different road segments, the central authority forecasts the vehicles mobility;

(iii) the central authority uses such a forecast to formulate and solve an instance of the optimization

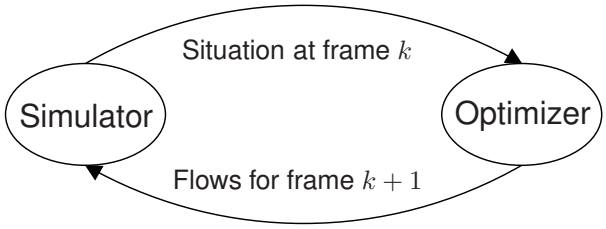

Figure 9. Simulation/optimization integration.

problems introduced in Section III;

(iv) using the data flows provided by the optimal solution ($x_k^v(i,j)$), the central authority suggests to the vehicles (either parked or moving) the optimal data scheduling, i.e., which data to download and from which node.

The integration between the optimizer and the ns-3 simulator shown in Figure 9. Basically, the simulator provides the optimizer with information about the vehicles' positions and amount of downloaded content. Then, the optimzer solves the resulting problem, and tells the simulator which flows are to be scheduled for the next frame. It is important to stress that the model being solved is always the same, i.e., the one described in Section III.

As expected, Figure 11 shows that the freshness obtained through simulation is lower than the one derived under ideal conditions, however the curves exhibit the same qualitative behavior. Interestingly, by carefully looking at the plots, the difference between the upper bound and the simulation results appears to be slightly greater in the "no backbone" case than under the "with backbone" setting. As also confirmed by the plots in Figure 12, this is due to simulations taking into account real-world effects such as message losses (e.g., due to channel fading), and such effects are more likely to occur when transmissions take place between mobile vehicles. We can conclude that the use of parked vehicles, in addition to the other advantages previously highlighted, also represents a way to make message transfers more reliable.

Maximizing the channel idle time. Next, we focus on maximizing the channel idle time (as in (7)), and present the results for three different target values of the freshness, Φ. Such values are set to fractions of the maximum content freshness achieved over the different scenarios (see Figure 4). The average and 10-th percentile of the resulting fraction of idle time are depicted for both the "no backbone" and "with backbone" cases, respectively, in the top and bottom plots of Figure 13. Note that in the "no park" case, only low-medium values of target content freshness are feasible, and that, in any case, the performance is poorer than when parked vehicles are used. Again, a backbone of

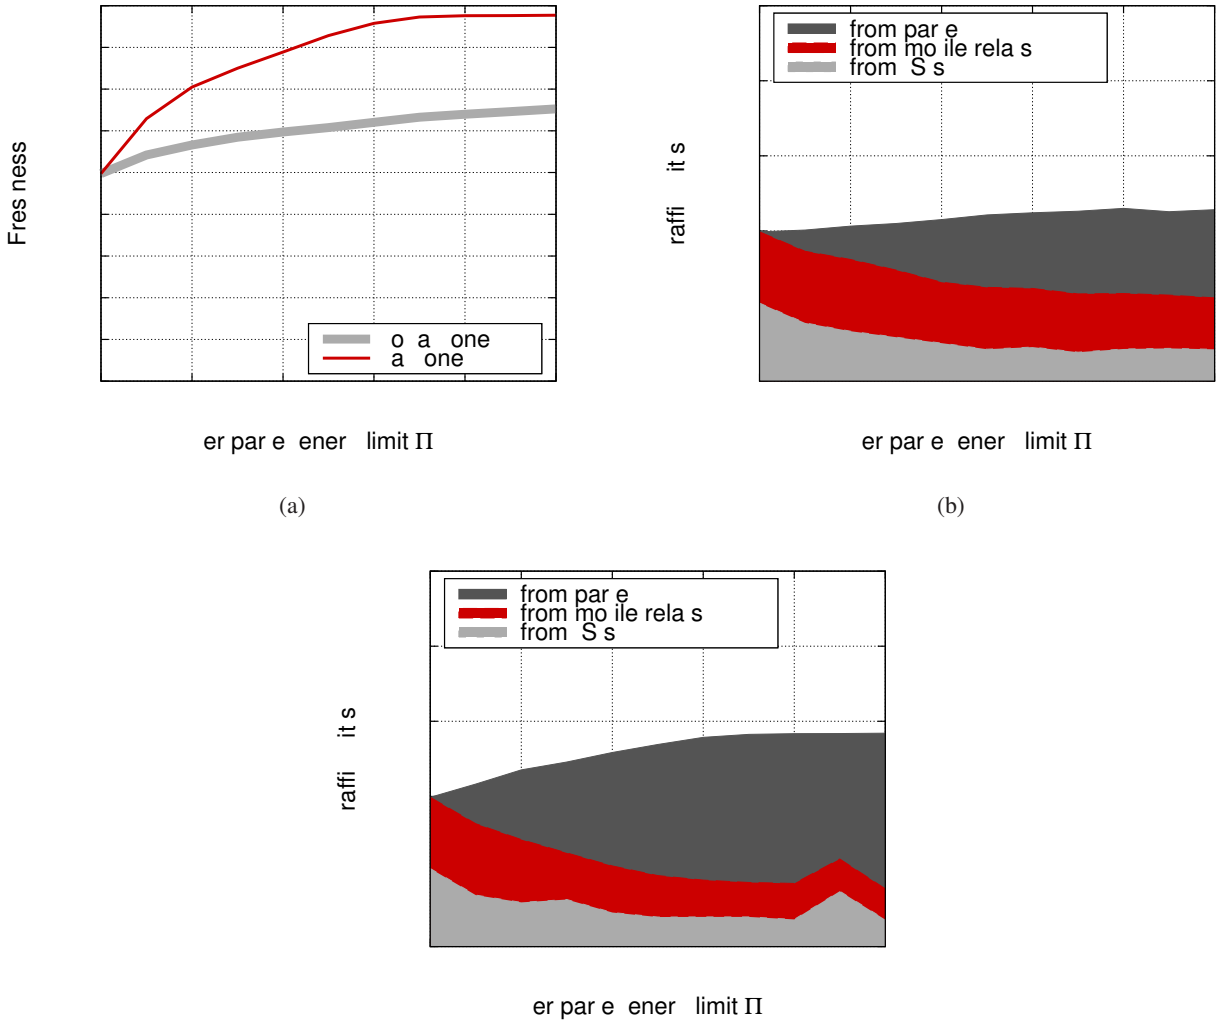

Figure 10. Limiting the use of parked vehicles: freshness (a); traffic to downloaders for the "no-backbone" (b) and "with backbone" (c) scenarios.

parked vehicles is very beneficial, as it results in an increased idle time. Interestingly, the bottom plot underlines that such an effect is particularly evident in the most congested zones of the road topology, suggesting that the goal we aimed at while designing the objective function in (7) has been achieved.

Finally, in order to get some insight on the data delivery delay, Figure 14 shows the average number of hops traversed by the data, from the node storing them to the intended downloader. The top plot, which refers to the case where content freshness is maximized, highlights that the use of parked vehicles gives an advantage also in terms of data delay. Indeed, such vehicles can store content and provide the users with it in place of the RSU, thus speeding up the process. When a backbone is built, the number of parked vehicles that are involved in the delivery process, hence the delay, quickly increases as the

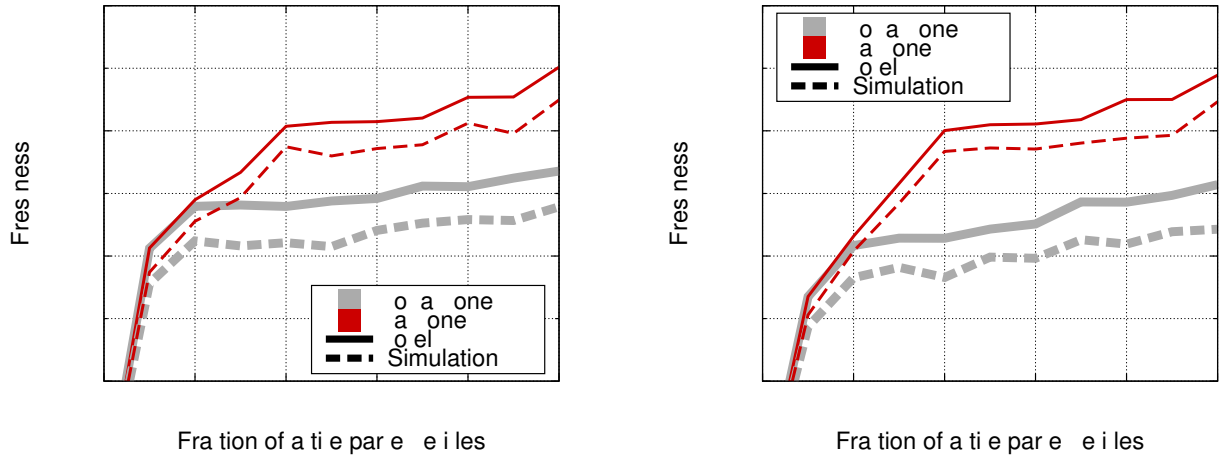

Figure 11. Freshness as a function of the maximum time period parked vehicles can be used as relays, for $R = 0$ (left) and $R = 0.5$ (right). Grey (thick) and red (thin) lines refer to the "no-backbone" and "with backbone" cases, respectively. Solid and dashed lines refer to the results through the model and by simulation, respectively

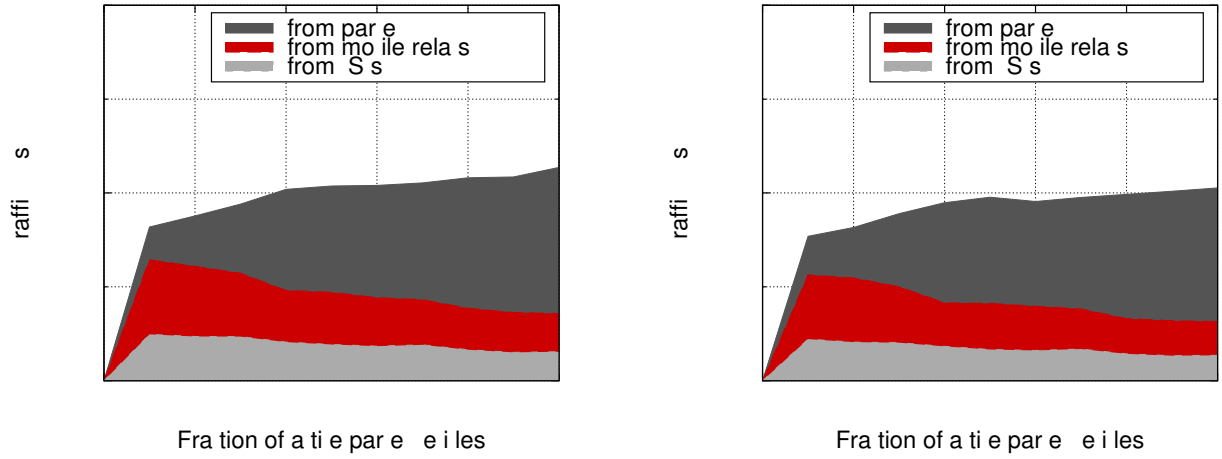

Figure 12. Traffic to downloaders for the "with backbone" case: model (left) and simulation (right).

fraction of parked vehicles that can be used grows. However, looking at Figure 4, it can be noted that this also leads to an increased content freshness. The case where the channel idle time is maximized is depicted in the bottom plot, where half of the parked vehicles can be used. Consistently with the results in the top plot, for such a scenario, the average number of hops traversed by the data is higher in presence of the backbone (but the obtained channel idle time is higher too), while the "no park" case again gives poor performance.

Maximizing fairness in exploiting parked vehicles.

We adopt the fairness objective described in (9) and investigate the amount of traffic that parked

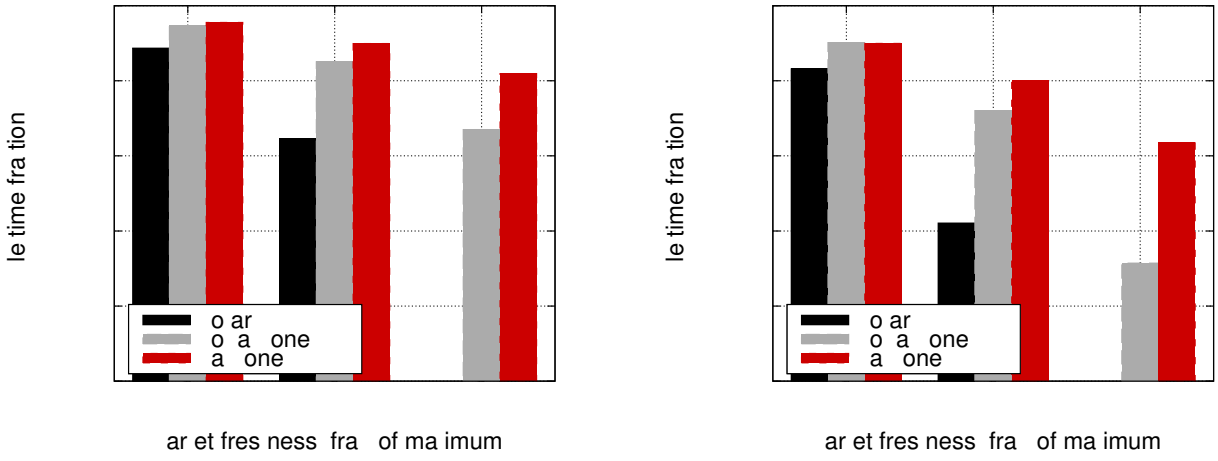

Figure 13. Average (left) and 10-th percentile (right) of the idle time at check points, with and without backbone and for different values of target freshness.

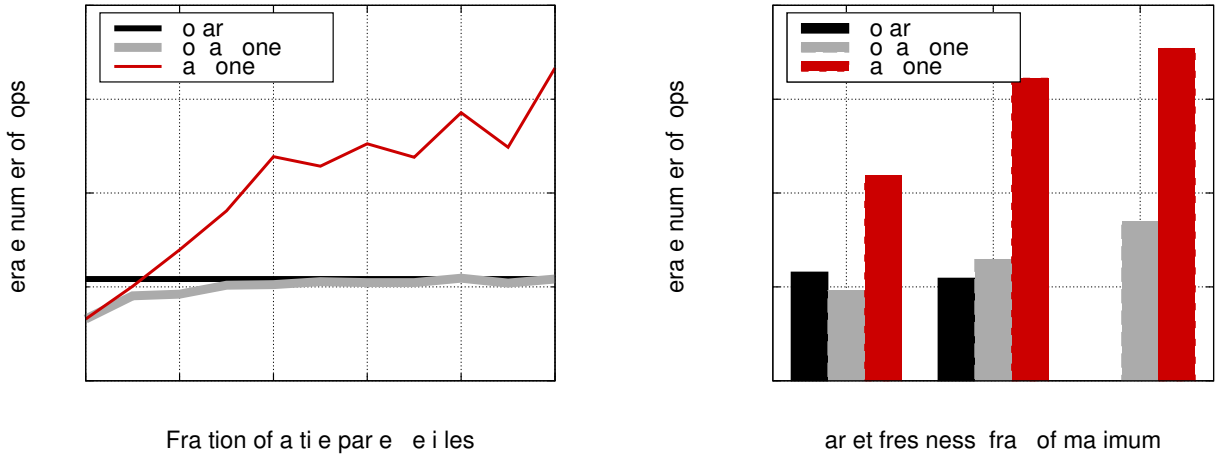

Figure 14. Average number of hops traversed by data traffic that downloaders successfully receive, when freshness (left) and idle time (right) are maximized.

vehicles have to send or receive, for $\Pi_i = \infty$ and as the target value of content freshness, Φ varies. In the derivation of such results, all parked vehicles can be used to relay or deliver data to downloaders.

Figure 15 highlights that the toll taken on parked vehicles increases as the required value of freshness grows. Note, however, that, consistently with the left plot in Figure 4, in the case of "no backbone" the maximum value of freshness achievable is limited to 13. Comparing Figure 15(a) to 15(b), it is also evident that, when a backbone is formed, the amount of traffic that each parked vehicle has to handle is greatly reduced and parked vehicles are used in a much fairer manner with respect to the "no backbone" scenario.

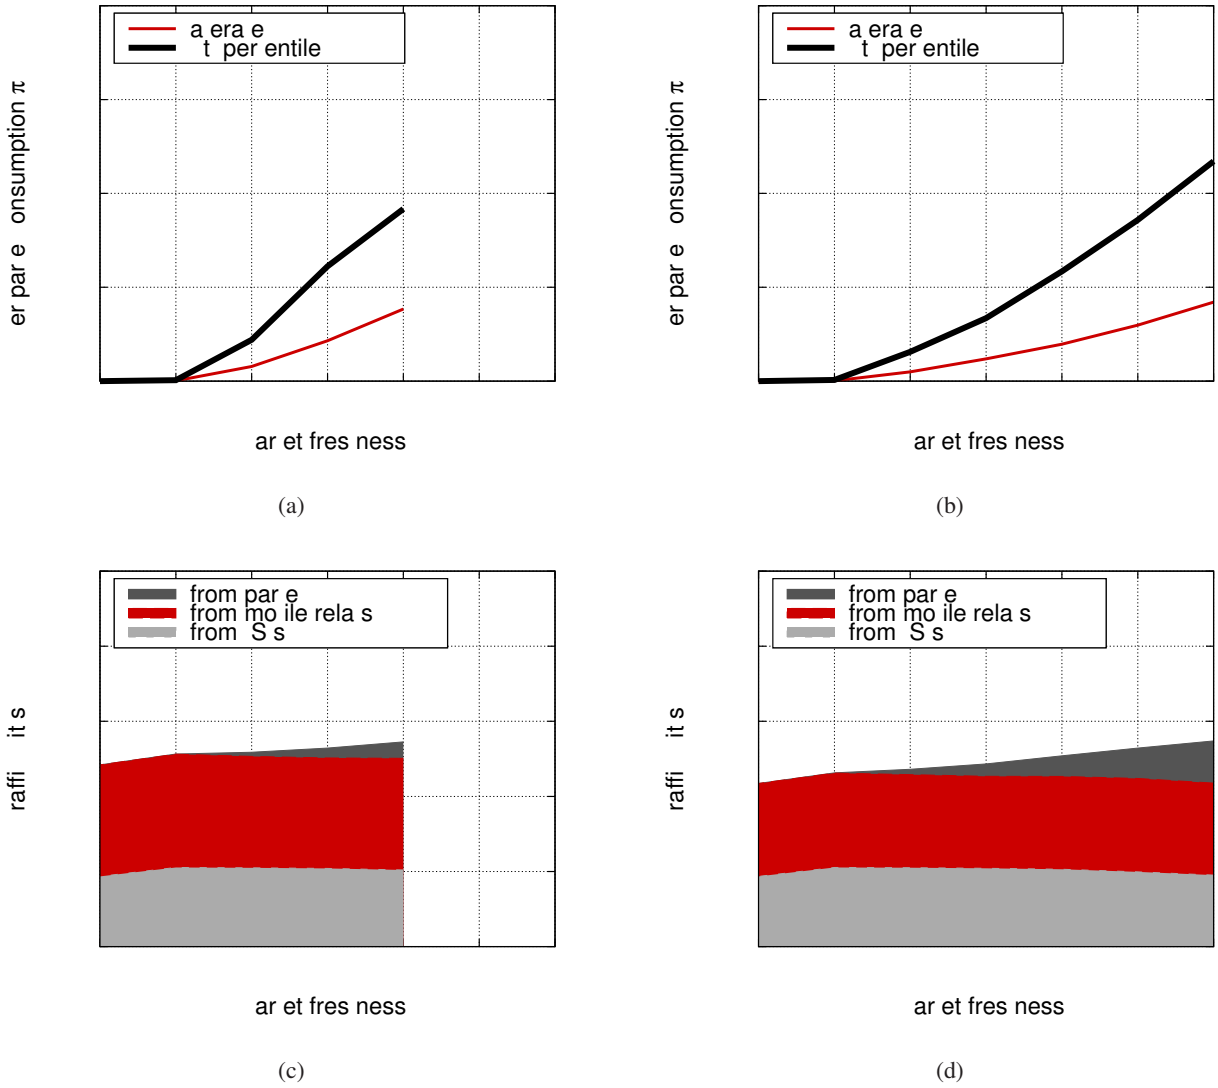

Figure 15. Max-fairness objective: data traffic π handled by parked vehicles for the "no-backbone" (a) and "with backbone" (b) scenarios; traffic delivered to the downloaders for the "no-backbone" (c) and the "with backbone" (d) scenarios.

We further investigate this aspect by comparing the above plots to those in Figure 15(c) and 15(d), which depict the amount of data traffic delivered to downloaders from parked vehicles, other mobile vehicles and RSUs. In the "no backbone" setting (Figure 15(c)), the total amount of traffic delivered to downloaders through parked vehicles remains almost constant as the target freshness increases, hence, from Figure 15(a), it can be deduced that fewer and fewer parked vehicles (i.e., the most suitable ones) are used to send and receive data. In presence of a backbone, instead, the traffic delivered by parked vehicles to downloaders grows as the target fairness increases but the the amount of data traffic remains quite evenly distributed among the available parked vehicles (Figure 15(b)). In other words, allowing the formation of a backbone of parked vehicles allows not only to meet stricter freshness goals, but also

to avoid over-exploiting parked vehicles in absence of explicit constraints on the maximum amount of data they should handle (i.e., with $\Pi_i = \infty$).

VII. Conclusions

The paper studied content downloading in vehicular network with the added twist of enlisting the support of parked cars in order to extend RSU service coverage. Optimization models were proposed with the goal of maximizing the content freshness and the utilization of radio resources, showing the remarkable contribution that parked vehicles can have on vehicular downloading efficiency. The energy toll taken on parked vehicles was accounted for by introducing a limit to the exploitation level of such vehicles as well as by maximizing the fairness in using parked vehicles under a constraint on the target content freshness.

Our results, derived through both numerical optimization and network simulation, proved that leveraging parked vehicles for content downloading is highly beneficial in a variety of scenarios, with different target metrics (e.g., content freshness) and constraints (e.g., maximum per-parked-vehicle usage) in place. Such an approach yields even more benefits when a backbone of parked vehicles can be formed, in terms of both system performance and use of parked vehicles energy resources.

Given such results, directions in which future work on this topic would evolve include the development and testing of backbone-forming protocols built on top of IEEE 802.11p. These may include parked vehicles and, occasionally, slow-moving ones (e.g., in congested areas). Allowing vehicles to exchange data on their battery level and type could also lead to a more efficient backbone creation.

Acknowledgments

This work was supported by the Regione Piemonte through the IoT_|_ToI project, which is part of the second call of POR F.E.S.R. 2007/2013.

References

[1] "Wireless Access in Vehicular Environments," IEEE, Std 802.11p-2010, July 2010.

[2] "Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service," ETSI, TS 102 637-2 V1.1.1, April 2010.

[3] T. L. Willke, P. Tientrakool, and N. F. Maxemchuk, "A Survey of Inter-Vehicle Communication Protocols and Their Applications," *IEEE Communications Surveys and Tutorials*, vol. 11, no. 2, pp. 3–20, 2009.

[4] F. Malandrino, C. Casetti, C.-F. Chiasserini, and M. Fiore, "Content downloading in vehicular networks: What really matters," in *30th IEEE Conference on Computer Communications (INFOCOM 2011), Mini-Conference*. Shanghai, China: IEEE, April 2011, pp. 426–430.

[5] F. Malandrino, C. E. Casetti, C.-F. Chiasserini, C. Sommer, and F. Dressler, "Content Downloading in Vehicular Networks: Bringing Parked Cars Into the Picture," in *23rd IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2012)*. Sydney, Australia: IEEE, September 2012, pp. 1534–1539.

[6] M. Fiore and J. Barcelo-Ordinas, "Cooperative download in urban vehicular networks," in *6th IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS 2009)*. Macau SAR, China: IEEE, October 2009, pp. 20–29.

[7] P. Hui, J. Crowcroft, and E. Yoneki, "Bubble rap: social-based forwarding in delay tolerant networks," in *9th ACM International Symposium on Mobile Ad Hoc Networking and Computing (ACM Mobihoc 2008)*. Hong Kong, China: ACM, 2008, pp. 241–250.

[8] H. Zhu, S. Chang, M. Li, K. Naik, and S. Shen, "Exploiting temporal dependency for opportunistic forwarding in urban vehicular networks," in *30th IEEE Conference on Computer Communications (INFOCOM 2011)*. Shanghai, China: IEEE, April 2011, pp. 2192–2200.

[9] Z. Zheng, Z. Lu, P. Sinha, and S. Kumar, "Maximizing the Contact Opportunity for Vehicular Internet Access," in *29th IEEE Conference on Computer Communications (INFOCOM 2010)*. San Diego, CA: IEEE, March 2010, pp. 1–9.

[10] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal, "Vehicular opportunistic communication under the microscope," in *5th ACM International Conference on Mobile Systems, Applications, and Services (ACM MobiSys 2007)*. San Juan, Puerto Rico: ACM, June 2007, pp. 206–219.

[11] G. Marfia, G. Pau, E. De Sena, E. Giordano, and M. Gerla, "Evaluating vehicle network strategies for downtown Portland: opportunistic infrastructure and the importance of realistic mobility models," in *5th ACM International Conference on Mobile Systems, Applications, and Services (MobiSys 2007), 1st ACM International Workshop on Mobile Opportunistic Networking (MobiOpp 2007*. San Juan, Puerto Rico: ACM, June 2007, pp. 47–51.

[12] Y. Ding, C. Wang, and L. Xiao, "A Static-Node Assisted Adaptive Routing Protocol in Vehicular Networks," in *4th ACM International Workshop on Vehicular Ad Hoc Networks (VANET 2007)*. Montréal, QC, Canada: ACM, September 2007, pp. 59–68.

[13] I. Dietrich and F. Dressler, "On the Lifetime of Wireless Sensor Networks," *ACM Transactions on Sensor Networks (TOSN)*, vol. 5, no. 1, pp. 1–39, February 2009.

[14] N. Liu, M. Liu, W. Lou, G. Chen, and J. Cao, "PVA in VANETs: Stopped cars are not silent," in *30th IEEE Conference on Computer Communications (INFOCOM 2011), Mini-Conference*. Shanghai, China: IEEE, April 2011, pp. 431–435.

[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 2011)*. Houston, TX: IEEE, December 2011.

[16] R. Crepaldi, R. Beavers, B. Ehrat, M. Jaeger, S. Biersteker, and R. Kravets, "LoadingZones: Leveraging Street Parking to Enable Vehicular Internet Access," in *7th ACM International Workshop on Challenged Networks (CHANTS 2012)*, Istanbul, Turkey, August 2012.

[17] F. Bai and B. Krishnamachari, "Spatio-temporal variations of vehicle traffic in VANETs: facts and implications," in *6th ACM International Workshop on Vehicular Inter-Networking (VANET 2009)*. Beijing, China: ACM, September 2009, pp. 43–52.

[18] C. Morency and M. Trépanier, "Characterizing Parking Spaces Using Travel Survey Data," CIRRELT, TR 2008-15, May 2008.

[19] T. Litman, "Parking management: strategies, evaluation and planning," *Victoria Transport Policy Institute*, vol. 25, 2006.

[20] C. Lochert, B. Scheuermann, C. Wewetzer, A. Luebke, and M. Mauve, "Data Aggregation and Roadside Unit Placement for a VANET Traffic Information System," in *5th ACM International Workshop on Vehicular Inter-Networking (VANET 2008)*. San Francisco, CA: ACM, September 2008, pp. 58–65.

[21] V. Badescu, "Dynamic model of a complex system including PV cells, electric battery, electrical motor and water pump," *Elsevier Energy*, vol. 28, no. 12, pp. 1165–1181, October 2003.

[22] R. Crepaldi and R. Kravets, "Governing Energy for Parked Cars," in *10th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2013)*. Banff, AB, Canada: IEEE, March 2013, pp. 87–94.

[23] US DOT. The Connected Vehicle - Next Generation ITS. [Online]. Available: http://www.itsa.org/industryforums/connectedvehicle

[24] 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, January 2011.

[25] "Power Consumption and Energy Efficiency Comparisons of WLAN Products," Atheros Comm., White Paper, April 2004.