This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TVT.2016.2540163

Fair Congestion Control in Vehicular Networks with Beaconing Rate Adaptation at Multiple Transmit Powers

Esteban Egea-Lopez, Pablo Pavon-Mariño

***Abstract*—Cooperative inter-vehicular applications rely on the exchange of broadcast single-hop status messages among vehicles, called beacons. The aggregated load on the wireless channel due to periodic beacons can prevent the transmission of other types of messages, what is called channel congestion due to beaconing activity. In this paper we propose to let vehicles transmit with different transmit powers, each with a particular beaconing rate. The selection of the rate is modeled as a Network Utility Maximization (NUM) problem. Our goal is to maximize the number of beacons delivered at each transmit power according to a well-defined fairness notion while complying with the maximum allowed beaconing load on the channel. The algorithm parameters can be set per vehicle and dynamically changed, which provides enough flexibility to support multiple applications on top of the control scheme, while the NUM model provides a rigorous framework to design a broad family of simple and decentralized algorithms, with proved convergence guarantees to a *fair* allocation solution. Simulation results validate our approach and show that it provides fair rate allocations in realistic multi-hop and dynamic scenarios with packet losses.**

***Index Terms*—Vehicular Communications, congestion control, beaconing rate control, Network Utility Maximization.**

I. Introduction

Inter-vehicle communications based on wireless technologies pave the way for innovative applications in traffic safety, driver-assistance, traffic control and other advanced services which will make up future Intelligent Transportation Systems (ITS) [1]. Communications for Vehicular Ad-Hoc Networks (VANET) have been developed and standardized in the last years. At the moment, a dedicated short range communication (DSRC) bandwidth has been allocated to vehicular communications at 5.9 GHz and both American and European standards [2] have adopted IEEE 802.11p [3] as physical and medium access control layers, based on carrier-sense multiple access with collision avoidance (CSMA/CA). These networks are characterized by a highly dynamic environment where short-life connections between vehicles are expected as well as adverse propagation conditions leading to severe or moderate fading effects [4].

Cooperative inter-vehicular applications usually rely on the exchange of broadcast single-hop status messages among vehicles on a single control channel, which provide detailed information about vehicles position, speed, heading, acceleration and other data of interest [5]. These messages are called *beacons* and are transmitted periodically, at a fixed or variable *beaconing rate*. Beacons provide very rich information about the vehicular environment and so are relatively long messages, between 250 and 800 bytes, even more if security-related overhead is added [6]. In addition, vehicles exchange other messages on the control channel: *service announcements* and *event-driven messages* as a result of certain events. For instance, *emergency* messages are transmitted only when a dangerous situation is detected.

The aggregated load on the wireless channel due to periodic beacons can rise to a point where it can limit or prevent the transmission of other types of messages, what is called *channel congestion due to beaconing activity*. Control schemes are required to prevent this situation and several alternatives are available: adapting either the beaconing rate, the transmission range, the transmission data rate, the carrier sense threshold or a combination of some of them [6]. The practical implementation of the system imposes two strong additional requirements on the control scheme: to be *distributed* and to grant beaconing rates to each vehicle in a *fair* way. Being distributed means that vehicles should control their rates making use only of the signaling information exchanged with their neighbor vehicles and without relying on any centralized infrastructure. Besides, to reduce the signal overhead, the exchanged information should be kept to a minimum. The control scheme should also provide quick and effective adaptation to changes in the environment, such as the channel conditions and the number of vehicles in range. The limits on such capabilities are captured by the *convergence properties* of the algorithm in use.

Fairness must be guaranteed as a safety requirement since beacons are used to provide vehicles with an accurate estimate of the state of their neighbors or mutual awareness [6]. Consequently, the fairness goal implies that no vehicle should be allocated arbitrarily less resources than its neighbors, under the constraints imposed by the available capacity. However, even starting from the previous principle, several notions of fairness can be defined, and in most of them there is a trade-off between fairness and efficiency [7]: more fairness results usually in a less efficient use of the shared resource. But less efficiency is detrimental to safety, since in general, the higher the beaconing rate, the higher the quality of the state information [6]. Thus, using an inadequate notion of fairness implies not simply wasting resources but also has a negative influence on the safety of the users. In summary, in vehicular

E. Egea-Lopez and P. Pavon-Mariño are with the Department of Information Technologies and Communications, Universidad Politécnica de Cartagena (UPCT), Spain, e-mail: {esteban.egea, pablo.pavon}@upct.es

networks it is necessary not only to provide fairness but also to be able to select the *appropriate fairness notion*. This ability is a distinguishing feature of the algorithm put forward in this paper with respect to other proposals [8]–[13].

Several beaconing rate and transmit power control schemes have been proposed in the literature. Although most of them are able to bring the channel load to the desired level, none of them is able to meet all the aforementioned requirements. In particular, all of them consider a very basic approach to fair allocation of beaconing rates, without a formal definition and rigorous convergence support.

In a previous paper [14] we modeled the congestion control as a *Network Utility Maximization (NUM)* rate allocation problem [15], where each vehicle is associated a so-called *utility function*, such that the problem objective becomes the maximization of the sum of utilities of each vehicle. This approach not only allows to design algorithms with proved convergence guarantees, but also, thanks to the work in [16], ensures that they converge to a formally well-defined fair beaconing rate allocation. That is, the particular (concave) shape of the utility function of the vehicles is related to the different notions of fairness induced globally, the so-called (α, ω)-fairness allocations. Therefore, the shape of the utility function is selected in order to enforce a particular type of fairness, such as proportional fairness ($\alpha = 1$), or max-min fairness ($\alpha \to \infty$). In our previous work [14] we assumed a single common power for all the vehicles and focused only on controlling the beaconing rate.

In this paper we apply the NUM methodology with a different approach: First, since standards allow the use of different transmit powers [17], we assume vehicles can use a set of transmit powers and select a particular beaconing rate for each power in the set, transmitting at multiple powers with different rates during the cycle. Networks interfaces for vehicular networks actually allow to set the transmit power for each individual frame. Second, the optimization variable used in the utility function is not simply the beaconing rate, but the beaconing rate used with a power multiplied by the number of neighbors reached with that power. This new variable counts the total number of copies of a beacon that are delivered (in absence of errors) in its neighborhood and so it provides a measure of the degree of dissemination of the state of a given vehicle. Thus, we call this new variable *Beacon Dissemination Rate (BDR)*. Each vehicle seeks to maximize the BDR, which can also be seen as a measure of the awareness its neighbors have of it. From this model we derive a particular distributed algorithm, with guaranteed convergence to a fair allocation, and remarkably flexible since vehicles can independently and dynamically adapt the algorithm parameters to the requirements of a wide range of applications. For instance, it can seamlessly be used to implement prioritized congestion control, in the sense that vehicles with specials needs are allocated higher rates to disseminate their status more frequently. In fact, the use of multiple powers in a cycle, a novel approach compared to previous proposals [8]–[13], supports the parallel execution of multiple applications, with different quality of service requirements, on top of it.

In the remainder of this paper we first review related works in section II. In section III, the problem is formulated as a NUM rate allocation problem where multiple powers are available. First we discuss qualitatively the goals, then we formulate formally the problem and we end proposing a particular algorithm. In section IV, it is validated and compared with other proposals in static scenarios. In section V we extend the comparison and evaluation in different dynamic scenarios. Finally, conclusions and future work are discussed in section VI.

II. Related Work

Transmissions in vehicular networks are broadcast in nature and use a CSMA-based medium access control (MAC) with constant contention window and no acknowledgment or retransmission. ETSI standards define a 10 MHz control channel for vehicular communications at 5.9 GHz [2]. Periodic beaconing over one-hop broadcast communications supports cooperative inter-vehicular applications by disseminating status and environmental information to vehicles on the control channel [5]. The rate of beacons has an influence on the quality of service of the applications [5], [11]. In fact, some applications may require a certain *beaconing reception frequency*, which is dependent on propagation losses, the number of contending nodes and other considerations, although standards [5] specify the required *beaconing generation rate*. A framework for decentralized congestion control (DCC) in the control channel has been published by ETSI [17], which can accommodate a variety of controls such as transmit power, message rate or receiver sensitivity, though the currently suggested mechanisms are very basic, and extensions are being discussed.

Among the aforementioned range of potential controls, there are proposals for pure beaconing [8], [9], [14] or transmit power control [6], [19] and more recently some for joint power and rate control [10], [12], [13]. Regarding pure beaconing rate control, [8], [9] propose rate control algorithms that comply with a generic beaconing rate goal. The former, called LIMERIC, uses a linear control based on continuous feedback (beaconing rate in use) from the local neighbors, whereas the latter, called PULSAR, uses an additive increase multiplicative decrease (AIMD) iteration with binary feedback (congested or not) from one and two-hop neighbors. Both of them, however, show limitations. Fairness is not defined formally by any of them. LIMERIC is shown to converge to a unique rate, that is, the same rate for all the vehicles sharing a link, and the convergence is only proved when all the vehicles are in range of each other, not for multi-hop scenarios. PULSAR requires synchronized updates and piggybacking congestion information from vehicles at a two hops. Authors of LIMERIC propose to combine the LIMERIC rate adaptation mechanism with the PULSAR piggybacking of two-hop congestion information to achieve global fairness [8], called here LIMERIC+PULSAR, but it is not proved neither discussed in detail. Pure transmit power control is discussed in [6], [19]. In the latter, we proposed a transmit power control that linearly adjusts the average carrier-sense range to keep the desired number of neighbors in range.

In this context, there is a related concept called *awareness control*: the techniques to adapt the communications parameters, such as transmit power or beaconing rate, to the requirements of an application. The requirements of the application are usually captured via one of several possible awareness metrics [12] which are directly or indirectly based on the time between successful reception of beacons, or beacon Inter-Reception Time (IRT), which depends among other things, on the probability of reception at a given *target distance*. Joint transmit power and rate control to enforce particular application quality of service requirements has been studied in [10]–[13].

In [10] authors derive analytically a measure called Information Dissemination Rate (IDR) which provides the number of copies of a beacon successfully delivered, according to the simplifying assumptions of the model: a linear network with uniformly placed vehicles and Poisson generation of messages. Then, authors show that the maximal IDR occurs in a certain range of channel occupancy. Therefore, they let the beaconing rate be set by the requirements of a particular tracking application and propose to control the transmission range in order to keep the channel occupancy within the desired levels. The precise mechanism to achieve a desired transmission range or how to cope with the effects of non-deterministic propagation are not discussed. Our optimization variable, BDR, is similar to IDR since in both cases they count the number of delivered copies of a beacon. IDR takes into account packet losses due to MAC and some physical-layer related effects within the limits of the underlying theoretical model, whereas for BDR those effects are included, though implicitly, during the measurement of the number of neighbors. Authors of [12] study with simulations the combinations of beaconing rates and transmit powers which optimize several awareness metrics and propose control strategies. In [13] authors propose to let the vehicles set their beaconing rates and transmit power in order to guarantee their minimum application requirements and use the linear control of LIMERIC+PULSAR in order to equally share the excess of capacity, when available, among all the vehicles. In all the three cases the strategy is similar but with different aspects: in [10] the beaconing rate is set by the tracking application and then the transmission range, via transmit power although not specified how, is set to keep the channel occupancy within some range. In [12] transmit power is set to maximize the IRT at a given common target distance and then the beaconing rate is adapted to comply with the maximum channel load using some beaconing rate control, for example PULSAR, as suggested by the authors. Finally, in [13] each vehicle sets its transmit power and rate to guarantee with a certain probability the IRT at the target distance and then, if available, an additional amount of beaconing rate is assigned via LIMERIC+PULSAR so that the maximum channel load is reached. In all three cases, the application sets the minimum value of rate or power and the control drives the remaining capacity. None of them, however, provides convergence guarantees in multihop scenarios or well-defined fairness notions.

In this paper, we propose a different approach, which partially integrates those strategies, while guaranteeing convergence to different notions of fairness allocations. On the one hand, we assume that a discrete set of powers is available, which is what the standards actually specify [17], and we leverage this extra degree of freedom by letting vehicles allocate beaconing rates at more than one power. Then vehicles maximize the BDR over all powers in a fair way. On the other hand, the constraints ensure that the channel occupancy is kept below the target value, so indirectly maximizing the IDR, according to [10]; but also, the application requirements can be *independently and dynamically* set by each vehicle by establishing a minimum allowed rate for each available power, which in practice supports awareness control as the previously discussed proposals do.

Similar approaches can be found for MANET and MESH networks [18]. However, our mechanism can be implemented with little effort due to the particularities of vehicular networks, whereas in other domains similar procedures are much more problematic to implement in realistic scenarios. Data flows in MANET are mainly multihop end-to-end whereas in vehicular networks they are mainly single-hop broadcast transmissions. Then, controlling congestion in the first case requires to send back to the source information about the congestion state of all the links used by the flow, which is difficult [18], whereas in the latter vehicles simply have to broadcast this feedback to inform all the involved sources of the congestion state of the links they use.

Finally, there are two main differences with our previous work [14]. Allowing the use of multiple power renders a different optimization problem, specially regarding the distributed algorithm that solves it. And the use of BDR as optimization variable provides a more accurate measure of the degree of dissemination of the vehicle state. Simply selecting optimal beaconing rates for each power separately is not valid, since they are coupled by a constraint and in the BDR. We are extending our previous problem, introducing several new degrees of freedom.

III. NUM MODELING OF THE BEACONING RATE CONTROL WITH DIFFERENT POWERS IN VEHICULAR NETWORKS

In this section we model the beaconing congestion control problem as a convex optimization problem. We first discuss the goals, connecting them to the modeling decisions, and provide an informal description of the problem. Afterwards, we formally pose the problem and propose a distributed algorithm for its solution.

A. Goals and design decisions

Let V be a set of vehicles in a vehicular network and assume each vehicle has a common set of available transmit powers, P. Each vehicle $v \in V$ can select a transmit power $p \in P$ for each beacon transmission and so can transmit beacons with a different rate r_{vp} beacons/s for each power. This assumption complies with the current standards [17].

The primary goal of the control mechanism is usually stated as to prevent channel congestion [6], that is, to keep the occupation of the resources below a desired level. However, we consider it rather a requirement, and consequently translate it as a constraint of the optimization problem. Given this requirement, our aim is actually to achieve the best possible usage of the shared resource, that is, to maximize some utility function of the control variables, $U_v(r_{vp})$. To determine it, consider that on the one hand, the quality of state information received by a single vehicle increases with the beaconing rate and, on the other hand, the quality of Cooperative Awareness Service [5] is globally increased with the number of vehicles that receive information about other vehicles. Therefore, our goal is to increase both aspects: since the number of potential receivers is a function of the transmit power used, and there is a discrete set of powers available, we combine them in a variable called Beacon Dissemination Rate (BDR), which is the sum of the beaconing rates at each power weighted by the number of receivers at the corresponding powers, $b(r_{vp}) = \sum_p n_{vp} r_{vp}$. That is, we assume that a vehicle transmits beacons with all the available powers and with a certain beaconing rate for each power which is to be determined by the optimization problem.

Hence, our main goal is to maximize the BDR but we want also to do it in a *fair way*. As shown in [16], we can induce a global notion of fairness by selecting appropriately the shape of the utility function applied to the variable. In particular, utility functions of the form

$$U_v(x_v) = \begin{cases} w_v x_v & \text{if } \alpha = 0 \\ w_v \log x_v & \text{if } \alpha = 1 \\ w_v \frac{x_v^{1-\alpha}}{1-\alpha} & \text{if } \alpha > 0, \alpha \neq 1 \end{cases} \tag{1}$$

result in (α, w)-fair allocations of the variable x_v, see [14] for further details. That is, by setting the α parameter we can globally obtain different fair resource allocations such as proportional ($\alpha = 1$) or max-min ($\alpha \to \infty$) fairness. All together, the goal of our optimization problem is to maximize the sum of the utilities of all the vehicles, $\sum_v U_v(b(r_{vp}))$, where U_v is chosen from eq. (1).

We are also interested in providing means for the applications to set their quality of service requirements. As discussed in section II, such requisites are expressed in terms of demanding, with a certain probability, a minimum beacon IRT at a target distance. Since the probability of reception at the target distance is determined (in absence of interference) by the transmit power, they are readily translated into constraints on the minimum beaconing rate allowed at each transmit power, r_{vp}^{min}. In addition, by letting the minimum rate to be zero a vehicle is not required to transmit beacons at all the available powers if not necessary.

Let us note that alternative goals may have been defined, such as the maximization of throughput or the *reception beaconing rate*, that is, the frequency of correctly received beacons. We focus in this paper solely on the control of the transmission beaconing rates at each power, which is also the customary approach in other proposals [8], [9], because: first, to take into account the reception rate we need to bring into consideration MAC behavior, which in most practical situations is out of effective control by the user. Second, the probability of reception is determined by a particular propagation model which is not known in practice and requires to treat the transmit power as another control variable, which removes, in general, the convexity of the problem, rendering it very difficult to solve in a decentralized way. Third, as pointed

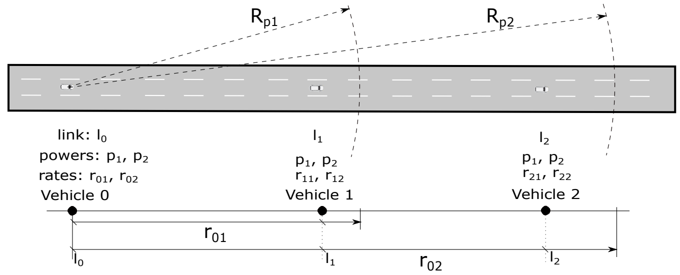

Fig. 1. Three vehicles, with two transmit powers available, transmitting r_{vp} beacons/s at power p respectively. Top: schematic diagram showing ideal transmission ranges of vehicle 0 at each power, R_{p1} and R_{p2}. Bottom: Solid arrows show the spatial area of reception of the beaconing rates at each power for vehicle 0, r_{01} and r_{02}: at power 1, only vehicle 1 is in transmission range of vehicle 0, whereas at power 2, both vehicles 1 and 2 are in range of vehicle 0. Dotted-lines respresent the virtual link of a vehicle, which is loaded by the beacons whose arrows intersect it. To avoid cluttering only the rates and ranges of vehicle 0 are shown.

out also in [8], [9], good performance at reception can be indirectly obtained by keeping the beaconing channel load at a given level. For instance, the IDR is maximized when the fraction of channel capacity used is kept at 0.6 to 0.7 [10].

Finally, we provide an informal description of the problem illustrated in Fig. 1. The key idea is to notice that every vehicle acts as a source, transmitting beacons which use a fraction of the wireless channel of their neighbors; and as a resource, defining a logical channel that is shared in the physical wireless channel with the vehicles in range. That is, we can assign every vehicle a logical resource: a virtual link l_v, which has a given capacity which is shared by all sources that are in range at a given power p_v, including its own beacons. As can be seen in Fig. 1, when vehicle 0 is transmitting beacons with power p_1 at a rate $r_{vp} = r_{01}$ beacons/s they are using the spatially shared physical wireless channel of only vehicle 0 and 1, that is, virtual links l_0 and l_1. Whereas when transmitting with power p_2 the beacons use all the links l_0, l_1 and l_2. And conversely the total load of resource l_0 is the sum of the rates using it at all the corresponding powers, that is, $r_{01} + r_{02} + r_{11} + r_{12} + r_{22}$, whereas r_{21} is not using l_0 because vehicle 2 is not in range of vehicle 0 at power 1.

In the following sections, we formally describe the above considerations. Let us note that, although useful to get an informal idea, we do not use the concept of virtual link to avoid introducing potentially confusing terms. Notation is summarized in Table I.

B. NUM model

Let V be a set of vehicles in a vehicular network and assume each vehicle has a common set of available transmit powers, P. Each vehicle $v \in V$ can select a transmit power $p \in P$ for each beacon transmission and so can transmit beacons with a different rate r_{vp} beacons/sec for each power. Let us assume ideal, deterministic, isotropic radio propagation. Let $n(v)_p$ denote the set of vehicles which have vehicle v in transmission range when they transmit at power p, which also includes vehicle v. Let us note that, since the links are assumed symmetric, this is also the set of vehicles in range of v at power p. Let n_{vp} be the number of elements in $n(v)_p$. The total rate

Notation	Description
r_{vp}	Beaconing rate of vehicle v with power p
$n(v)_p$	Set of vehicles which have vehicle v in transmission range when they transmit at power p, including v itself
n_{vp}	Number of elements of $n(v)_p$
C	Maximum Beaconing Load (MBL) in beacons/s
$b(r_{vp})$	Beacon Dissemination Rate (BDR)
r_{vp}^{min}	Minimum beaconing rate that vehicle v must use with power p
R_{max}	Total maximum beaconing rate
$U_v(x)$	Utility function of vehicle v
π_v	Lagrange multiplier (price) associated with the relaxed constraints
Π_{vp}	Sum of the prices received with power p
α	Fairness parameter
β	Gradient projection stepsize in Algorithm 1
a	Gradient projection stepsize in Algorithm 2
ϵ	Regularization term weight

TABLE I
NOTATION

received by each vehicle is the sum of the rates used by the members of its sets of neighbors for each power and we are interested in limiting this amount to a maximum C beacons/s, that is, a *Maximum Beaconing Load* (MBL), to avoid channel congestion. Let us define the *Beacon Dissemination Rate, (BDR)* as $b(r_{vp}) = \sum_p n_{vp} r_{vp}$. Let r_{vp}^{min} be the minimum beaconing rate that a vehicle v must use at power p. Let $U_v(x)$ be the utility function of vehicle v. Finally, let R_{max} the absolute maximum beaconing rate that every vehicle is allowed to use.

The NUM beaconing rate allocation problem is given by:

$$\max_{r_{vp}} \sum_v U_v(\sum_p n_{vp} r_{vp}) - \epsilon \sum_p r_{vp}^2 \tag{2a}$$

subject to:

$$\sum_p \sum_{v' \in n(v)_p} r_{v'p} \leq C \quad \forall v \in V \tag{2b}$$

$$r_{vp}^{min} \leq r_{vp} \quad \forall p \in P \quad \forall v \in V \tag{2c}$$

$$\sum_p r_{vp} \leq R_{max} \quad \forall v \in V \tag{2d}$$

The objective function (2a) is essentially the sum of the utilities $U_v(b(r_{vp}))$ of each vehicle. We assume vehicle utility functions as the ones in (1) and so we have that a rate allocation to the vehicles is α-fair if and only if, it is the optimum solution of (2). However, the function $\sum_v U_v(b(r_{vp}))$ is concave but not strictly concave which may result in oscillations in the control algorithm. In order to correct it and make it strictly concave we introduce a regularization term in the objective function, $\epsilon \sum_p r_{vp}^2$, with ϵ small enough to avoid significant deviations from our real maximization goal.

Constraints (2b) ensure that the beaconing load at a given vehicle, which is the one generated by the neighboring vehicles at different powers plus its own rates, must be below the MBL, C. Constraints (2c) guarantee a minimum beaconing rate allocated at each power. These constraints are expected to be set by applications to ensure their minimum required quality of service. Finally, constraints (2d) force the sum of vehicle rates at all powers to be under a total maximum value (R_{max}) specified by standards.

C. Dual decomposition

In order to find a decentralized algorithm to solve the problem we use a dual decomposition. We first form the Lagrangian function L of (2a) relaxing the constraints (2b):

$$\begin{aligned} L(r,\pi) &= \sum_v U_v(b(r_{vp})) - \epsilon \sum_p r_{vp}^2 - \\ &- \sum_v \pi_v \left(\sum_p \sum_{v' \in n(v)_p} r_{v'p} - C \right) = \\ &= \sum_v \left(U_v(b(r_{vp})) - \sum_p \sum_{v' \in n(v)_p} r_{v'p}\pi_{v'} - \epsilon \sum_p r_{vp}^2 \right) + \\ &+ C \sum_v \pi_v \end{aligned} \tag{3}$$

where $\pi_v \geq 0$ are the Lagrange multipliers (prices) associated with the relaxed constraints. The Lagrange dual is the maximum value of the Lagrangian function over the domain of the rates:

$$\begin{aligned} g(\pi) = \max_{\substack{r_{vp}^{min} \leq r_{vp} \\ \sum_p r_{vp} \leq R_{max}}} \{ U_v(b(r_{vp})) - \sum_p \sum_{v' \in n(v)_p} r_{v'p}\pi_{v'} - \\ - \epsilon \sum_p r_{vp}^2 \} \end{aligned} \tag{4}$$

Given a set of non-negative prices π, the optimal rate allocation solving the Lagrange dual (4) for each vehicle is $r_{vp}^*(\pi)$.

Since the original problem is convex with linear constraints, it has the strong duality property [20] and the Karush-Kuhn-Tucker (KKT) conditions characterize its optimum solution. Then, it can be shown that there is a set of optimum link prices π^* such that the associated rates given by (4) are the optimal solution of the original problem (2). The problem of finding such optimum prices is called the dual problem, which is defined as:

$$\min_{\pi \geq 0} g(\pi) = \min_{\pi \geq 0} \left\{ \max_{\substack{r_{vp}^{min} \leq r_{vp} \\ \sum_p r_{vp} \leq R_{max}}} L(r,\pi) \right\} \tag{5}$$

In our case, it can be shown that the objective function in (5), called the dual function, is strictly convex and differentiable, since the objective function in (2a) is strictly concave. Thus, the dual problem (5) has a unique set of optimal prices π^* [21]. The classical dual approach for solving the rate allocation problem consists of finding the dual optimal prices π^* using a gradient-based algorithm [21], as a means to (in parallel) obtain the optimal rate allocation r_{vp}^*. Hence, to compute its rates, each vehicle v needs to know just its own utility function U_v and the set of prices $\pi_{v'}$ of its *neighbor vehicles* and use them to solve problem (4), that is, *it can be solved separately by each vehicle*.

To summarize, in order to find the optimal beaconing rate allocation, vehicles have to iteratively exchange their prices π with their one-hop neighbors and use them as input to the optimization problem (4) which can be solved autonomously by each vehicle with its local information. In the next section we provide an algorithm that finds the optimal allocation in a distributed way. Since it is based on a gradient projection, let us notice that given a set of prices π, the gradient of the dual function g evaluated at π is given by:

$$\frac{\partial g}{\partial \pi_v}(\pi) = \sum_p \sum_{v' \in n(v)_p} r_{v'p} - C, \quad \forall v \tag{6}$$

D. Fair Adaptive Beaconing Rate with Multiple Powers for Intervehicular Communications (FABRIC-P)

In this section, we propose FABRIC-P (Fair Adaptive Beaconing Rate with Multiple Powers for Intervehicular Communications), an algorithm that solves problem (5) in a distributed manner. Let us note that there are two main steps: first vehicles compute their gradients (6) and update their prices and, second, solve the *local* optimization problem (4). There are several alternatives to solve the latter: we propose to use a gradient projection algorithm that provides a quick convergence to the optimal solution. In Algorithm 1 we describe the main procedure while in Algorithm 2 we describe the local gradient projection.

Algorithm 1 . Fair Adaptive Beaconing Rate with Powers for Intervehicular Communications (FABRIC-P).

1: At $k = 0$, set initial vehicle prices π_v^0 and rates r_{vp}^0
2: Then, at each time k:
3: Step 1. Each vehicle v receives the prices of neighbor vehicles $\pi_{v'}^k, v' \in n(v)_p$. Then, for each power, vehicle updates n_{vp}^k and stores separately the sum of the prices received: $\Pi_{vp}^k = \sum_{v' \in n(v)_p} \pi_{v'}, \quad \forall p \in P$
4: Step 2. Each vehicle computes π_v^{k+1} according to: $\pi_v^{k+1} = \left[\pi_v^k + \beta\left(\sum_p \sum_{v' \in n(v)_p} r_{v'p} - C\right)\right]^+_{\pi_v \geq 0}$
5: Step 3. Each vehicle computes r_{vp}^{k+1} as the result of execution of Algorithm (2): $r_{vp}^{k+1} = LGP(\Pi_{vp}^k, n_{vp}^k, r_{vp}^k)$

Algorithm 2 . Local Gradient Projection with diminishing step size.

1: **procedure** LGP(Π_p, n_p, r_p)
2: $\quad r_p^1 = r_p$, $i = 1$
3: $\quad$ **repeat**
4: $\quad\quad \nabla L(r_p^i) = n_p(\sum_p n_p r_p^i)^{-\alpha} - 2\epsilon r_p^i - \Pi_p$
5: $\quad\quad r_p^{i+1} = \left[r_p^i + \frac{a}{i}\nabla L(r_p^i)\right]^+_{\substack{r_p^{min} \leq r_p^i \\ \sum_p r_p^i \leq R_{max}}}$
6: $\quad\quad i = i + 1$
7: $\quad$ **until** $\nabla L(r_p^i) = 0$
8: $\quad$ **return** r_p^{i+1}
9: **end procedure**

Where we use the notation $[x]_X^+$ to denote the orthogonal projection with respect to the Euclidean norm of a vector onto the convex set X, that is, $[x]_X^+ = \arg\min_{z \in X} ||z - x||_2$ [21]. The convergence of the algorithm is shown in Appendix A.

We finish this section by discussing and clarifying relevant aspects of the algorithm and its practical implementation. We start with aspects related to the *algorithm and physical meaning of its parameters*.

- The prices, π_v, reflect the congestion state of the wireless channel as measured by a vehicle and can be thought of as the cost to use the vehicle shared channel. Each vehicle measures its own perceived congestion relative to the MBL when the gradient is computed at each iteration k (in Algorithm 1, line 4). The price increases when the channel is congested and the other way round.
- Note that the implementation of this algorithm is *decentralized and uses only collected information from one-hop neighborhood*. Each vehicle broadcasts its price π_v^k to its neighbors in range. It is not necessary to disseminate every price to all the vehicles in the network. Then, each vehicle updates its rate with Algorithm 2 using only the collected prices from its one-hop neighboring vehicles (Π_p) as well as their current rates (r_p) and the locally measured number of neighbors at each power (n_p).
- Algorithm 2 solves locally the maximization of eq. (4) by gradient projection. Line 4 is just the gradient of eq. (4) and is separable in p. As can be seen there, the beaconing rate at a given power p is reduced when the congestion state at this power is high (Π_p). We use a diminishing step with $a > 0$. The gradient update is projected on the constraint set in line 5. All the procedure can be effectively and quickly run on conventional hardware, and to speed the process, it can be stopped once a certain accuracy has been achieved.
- Just by selecting the α parameter *different notions of fair allocations are obtained*. Moreover, our procedure can be used to achieve not only different classes of fairness but also to incorporate heterogeneous utility functions and constraints for different vehicles. And vehicles do not need to know the parameters or utility functions of other vehicles. The criteria for selecting a particular fairness notion are application or even scenario dependent. Further considerations on this are left as future work.
- The parameters β, a, and ϵ do not have any physical meaning and are used just to tune the algorithm behavior. Both β and a basically determine the speed of convergence. High values increase speed, provided they are below the bound derived in Appendix A. Since the algorithm converges any initial price π_v^0 is valid. The values used for our simulations in next sections have been selected by experimentation.

Next we discuss *implementation and practical aspects*.

- The number of neighbors at each power, n_{vp}, is locally measured by each vehicle as it receives beacons from its neighbors. Here we are using the assumption that the links are symmetric and consequently if a vehicle receives a beacon from a neighbor it can also reach it back with the same power. In practice, however, n_{vp} is a random variable because of fading and interference. Our simulations show that such effects tend to overestimate the congestion, as we discuss later. If necessary, it can be compensated by filtering the unreliable links, although we leave a fine tuned proposal as future work.

- Every vehicle updates and stores a single price π_v. However, it broadcasts its value piggybacked in beacons transmitted at different powers. When a vehicle receives a price in a beacon it accumulates it in a local vector (Π_{vp}) indexed by the transmit power used for the beacon, since in Algorithm 2, line 4, it is necessary to subtract the sum of the prices received with different powers. The Π_{vp} also reflect the particular congestion state of the neighborhood due to beaconing with power p.
- The minimum beaconing rate at each power, r_{vp}^{min} is set by a particular application operating on top of the congestion control. The typical procedure would be as follows: according to some propagation model the probability of reception with power p at a target distance d is fixed and so the application sets a minimum r_{vp}^{min} to guarantee the required IRT with a certain probability. In practice it is necessary to send some signaling at all powers in order to let the vehicles know their neighbors at all powers, so $r_{vp}^{min} > 0$ even though it can be a very low value.
- The *procedure is robust against errors* such as packet losses due to fading, collisions or hidden-node interference. We have simulated it with realistic MAC and propagation models and the results show convergence to the close vicinity of the optimal allocation in spite of packet losses.
- To compute the perceived congestion, via eq. (6), our choice is to let vehicles inform others of their current beaconing rate and transmit power used by piggybacking them in the beacon. In our opinion it is the more reliable option with regard to the accuracy of the rate control, since it informs about the actual offered load of the channel in absence of errors, such as fading or interference, though it tends to overestimate the actual load. On the other hand, the algorithm requires each vehicle v to store and send a non-negative real number, its price π_v. So, regarding *overhead*, vehicles should broadcast at most their current beaconing rate, transmit power used and the price, all of them piggybacked in a beacon, that is, three real numbers, which adds little overhead to the current procedures, around 1% for 500-byte beacons.
- Finally, *synchronous or asynchronous* implementations can be considered. In the first case, all the vehicles update their rates at the same instant with the received prices. This is possible in practice for vehicles equipped with a GPS device. In that case, all the neighbor prices are available prior to the beaconing rate update. In the second case, each vehicle may update its rate at different instants. We have used the asynchronous case for our simulations, which show that FABRIC-P also works correctly. Regarding the actual implementation of the multiple beaconing rates, vehicles only need to set up P timers to $1/r_{vp}$ respectively. When they time out, a vehicle sends a new beacon to the MAC layer together with the value of transmit power to be used in its transmission.

Parameter	Value
Data rate (V_t)	6 Mbps
Sensitivity (S)	-92 dBm
Frequency	5.9 GHz
Noise	-110 dBm
SNIR threshold (T)	4 dB
Neighbor Table update time	1 s
Sample period T_s	1 s
Beacon size (b_s)	500 bytes
Maximum Beaconing rate (R_{max})	10 beacons/s
β	10^{-8}
ϵ	10^{-8}
a	100

TABLE II
COMMON PARAMETERS FOR SIMULATIONS

IV. VALIDATION

In this section we test the validity of our algorithm and assumptions in a static scenario where vehicles do not move which allows us an accurate control of the vehicles interactions. It also allows us to compare the results of the algorithm with the exact solution of the maximization problem (2), provided by a numeric solver, JOM [23], based on the Ipopt [24] optimization library. We focus on scenarios where the differences between one or several powers are more evident. We only use 2 transmit powers, low and high, in all the cases, in order to avoid cluttering the figures and results and for the sake of clarity in the discussions. We consider it as a reasonable number for many potential applications.

Simulations setup. We summarize first the simulation parameters that are common to the simulations studies in this and the following section. The simulations have been implemented with the OMNET++ framework and its inetmanet-2.2 extension [25], which implements the 802.11p standard. This library also implements a realistic propagation and interference model for computing the Signal to Interference-plus-Noise Ratio (SINR) and determining the packet reception probabilities, implementing capture effect too.

In our tests, vehicles are located on a straight single lane road and their positions are either *deterministic*, that is equally spaced with distance d m or randomly positioned according to a *Poisson* distribution of average density ρ vehicles/m. Both *free space* and *Nakagami-m* propagation models have been used. In both cases, the path loss exponent has been set to 2.5. This is a slightly lower value than those reported by [22], measured in suburban scenarios. Higher values result in shorter transmission range and so congestion is more unlikely and its effects milder. Thus, our values model a worst case scenario. Nakagami-m shape parameter has been set to $m = 1$, to model severe (Rayleigh) fading conditions. Values reported by [22] suggest even stronger fading. The MBL has been set to 3.6 Mbps, which is 60% of the available data rate of 6 Mbps. Since we use a beacon size of 500 bytes plus 76 bytes of MAC and physical headers, it results in a MBL of $C = 781.25$ beacons/s. Table II summarizes the rest of common parameters. Vehicles collect prices and information from the neighborhood during a sample period T_s and then execute an iteration k of Algorithm 1. All the simulations have been replicated 10 times with different seeds.

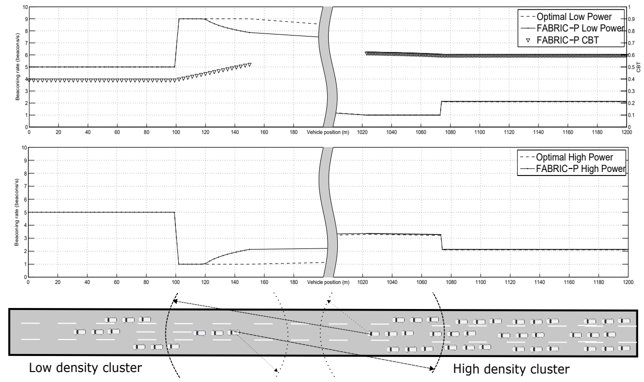

Fig. 2. Beaconing rate versus position in an ideal two-overlapped-cluster scenario for FABRIC-P with $\alpha = 1$ (proportional fairness). Exact optimal value computed with a numeric solver is shown. The part of the X axis without vehicles is not shown. CBT versus position is also shown on the right Y axis in the top subplot. Schematic diagram (not on scale) is shown below, where dashed and dotted arrows and arcs show high and low power transmission ranges respectively.

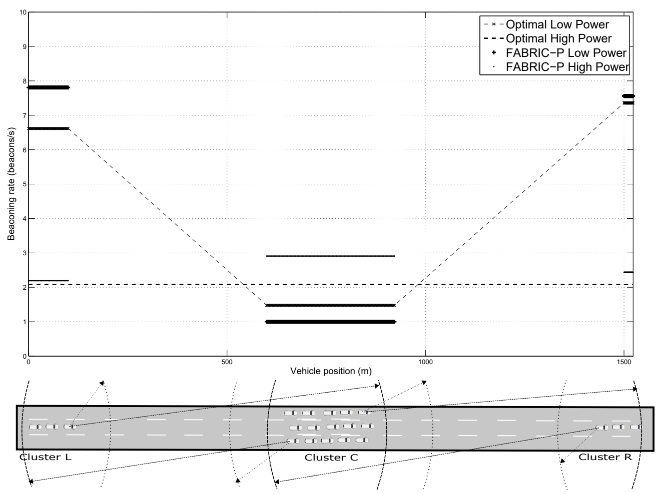

Fig. 3. Beaconing rate versus position in an ideal three-cluster scenario for FABRIC-P with $\alpha = 1$ (proportional fairness). Exact optimal value computed with a numeric solver is shown. Schematic diagram (not on scale) is shown below, where dashed and dotted arrows and arcs show high and low power transmission ranges respectively.

Scenario 1. Multihop ideal scenarios: overlapped clusters and bottleneck cluster.

In these tests we evaluate two scenarios where not all vehicles have the same number of neighbors. An ideal channel is considered and packets are not lost. In the first one, we have *two clusters of vehicles in partial coverage* of each other: on the left, at the origin, there is a low density cluster (A) of 51 vehicles separated 3 m of each other; on the right, at 1023 m of the origin, there is a denser cluster (B) of 181 vehicles separated 1 m. We set up minimum rates to $[1, 1]$ beacons/s and transmit powers to $[100, 1000]$ mW. Since the transmission range is 367.83 and 923.95 m at low and high power respectively, the 17 rightmost vehicles of cluster A are partially in range of the 51 rightmost vehicles of cluster B at high power and reciprocally. The most loaded nodes at each cluster are the last vehicle of A and the first vehicle of cluster B, in range of 102 and 198 vehicles respectively. Except for the overlapped-range nodes, all the rest of the nodes have the same number of neighbors with all the powers in each cluster.

The results are shown in Fig. 2. In this case, for the vehicles not overlapping range in both clusters, the beaconing rate is split between both powers, since they all have the same number of neighbors at each power. For cluster A, there is not congestion and the sum of the rates of each vehicle is equal to R_{max}, that is, 10 beacons/s, whereas in cluster B the sum is 4.22, complying with the MBL constraint. For the overlapped vehicles, in cluster A, the beaconing rate corresponding to high power is decreased while the low power one is increased, keeping the sum constant to 10, whereas in cluster B the opposite occurs. Therefore, from the point of view of the rest of the vehicles in their respective clusters there is no variation in the absolute beaconing rate they are receiving from the vehicles at the edges. However, in the overlapped area, vehicles from cluster A are receiving more beacons from their neighbors at cluster B, while still informing of their presence to B. This is beneficial in some situations, for instance if this scenario models a cluster of vehicles (A) approaching a traffic jam (B), possibly at high speed, it is beneficial to increase the quality of the information received by the high speed vehicles. In all the cases the CBT is kept below the MBL as shown in Fig. 2 top.

In the second scenario we show a typical situation where the use of different powers is appropriate. We have a *central cluster with a very high density of vehicles and two low-density clusters at a certain distance*. This may model again a traffic jam in a highway with clusters of vehicles approaching from both directions. The central cluster is congested. If a single high power is used it acts as a bottleneck for the other two clusters. In particular we have set up a cluster of 51 vehicles from 0 to 100 m (L), a central one of 162 vehicles from 600 to 923 m (C) and another one of 47 vehicles from 1500 to 1520 m (R). At high power L and R are not in range of each other, whereas C is in range of both L and R. The situation is analogous to that of Fig. 1 but with clusters instead of single vehicles.

Results in Fig. 3 show that indeed with high power the beaconing rate of the 3 clusters is limited by congestion and kept at 2.11 beacons/s (in the exact solution). With the use of an additional power, clusters L and R can use a much higher beaconing rate at low power, between 6.5 and 7.5 beacons/s. It basically shows that FABRIC-P in a trivial case like this effectively allocates rates as an intuitive heuristic would: increasing the low power rate when the high power enters congestion. Additionally, the CBT is kept exactly at the MBL limit, 0.6 (not shown in Figure).

With a different fairness ($\alpha = 2$) the results are the same in these particular cases. Whether is preferable to set proportional, max-min or any other α-fairness is a matter of discussion and possibly application or scenario dependent. In any case, the key advantage of FABRIC-P is that it can be configured with the α parameter to achieve any of the fairness notions. Moreover, vehicles can use different fairness parameters simultaneously and they can even be dynamically set. For the moment, a study on the applicability of the fairness alternatives and its dynamical setting is left as future work.

Finally, our results also show that FABRIC-P, without channel access or propagation effects, converges to a solution close to the optimal value in both cases. The solutions in Fig. 2 and 3 have been achieved after 200 steps of the algorithm. Even though the exact optimal allocation has not been achieved at that time, the MBL constraint is already met. The convergence to the exact optimal allocation may be long, specially if the optimal allocation shows sharp edges, but in practice, for randomly positioned scenarios and realistic propagation models, it is not necessary such a long convergence time to observe acceptable results, that is, even though the allocation is not yet optimal the CBT is already below the MBL. The results for the Poisson-like vehicle positioning, shown in Fig. 4, and results in Scenarios 3 and 4 are an example of this.

Scenario 2: Realistic scenario with hidden nodes and packet losses.

In this scenario vehicles are positioned along a 2800 m long line, with a Poisson distribution of average density $\rho = 0.1$ vehicles/m. The propagation models have been set to free space and Nakagami-m with path loss exponent equal to 2.5. In both cases, the transmit power has been set to $[400, 1000]$ mW which results in transmission ranges $[640.43, 923.94]$ m for free space and (an average of) $[568.23, 819.78]$ m for Nakagami-m [19]. And minimum rates have been set to $[1, 1]$ beacons/s. Let us remark that, with this setting, there are interferences due to hidden nodes. In fact, in this scenario we can see the effects of MAC and hidden collisions as well as fading and channel errors in our model, since it has been simulated with an accurate 802.11p and propagation model, including SINR evaluation and capture effect. In this case, we set $\alpha = 2$ instead of proportional fairness.

To better measure the effectiveness of the algorithms in a lossy scenario, we define the *effective delivery ratio at distance* d, $D_v(d) = \frac{c_v(d)}{n_v^S(d)}$, where $c_v(d)$ is the total number of correctly received copies of a beacon up to a distance d of the transmitter v and $n_v^S(d)$ is the total number of copies of a beacon whose power is above the sensitivity up to that distance. That is, $D_v(d)$ indicates the fraction of copies of a broadcast beacon that are correctly received at a distance no greater than d from the transmitter. We then define *effective beaconing rate at distance* d, as $\hat{r}_v(d) = \sum_p r_{vp}\bar{D}_{vp}(d)$, which is the product of the allocated beaconing rate and the average effective delivery rate at a certain distance summed over all powers. That is, $\hat{r}_v$ is an average of the *total beaconing rate correctly received up to a certain distance of the transmitter*.

In Fig. 4 we show the allocated and effective beaconing rate. Confidence intervals at 95% have been computed but, since the maximum radius of the confidence interval does not exceed 0.08 and 0.4 beacons/s for free space and Nakagami-m respectively, they are not shown to avoid cluttering the figure. The exact optimal allocations computed with the numeric solver JOM are also shown.

As can be seen, the optimal allocation is far from obvious in this scenario. As Fig. 4(a) shows, with free space propagation and hidden-node and propagation losses, asynchronous FABRIC-P approximately tracks the optimal allocations after only 90 algorithm steps. More importantly, the CBT is already below the MBL at that time, which shows that it is not necessary to achieve convergence to the optimal allocation to meet the MBL constraint. In fact, in the following sections when the time evolution of the rates for this scenario is provided we show how the algorithm tends to move out of congestion quickly. With strong fading (Nakagami-m), shown in Fig. 4(b) FABRIC-P reduces the high power beaconing rate while increasing the low power one, and smooths the allocation globally. The reason is that there is now a chance of randomly receiving beacons from far away nodes. It has two different effects. On the one hand, the perceived number of neighbors (n_{vp}) is higher for low power, 10% on the average, and due to its weight in the utility function the corresponding beaconing rate is increased. On the other hand, the congestion is overestimated because in the computation of the gradient at step 4 of Algorithm 1 vehicles use the beaconing rate piggybacked in the beacons. However, actually around 12% of the beacons are lost because of fading, and so the physically measured CBT is below the MBL as shown in Fig. 4(b). It results in a less efficient use of the channel. If necessary, this discrepancy between the announced rate and the received one can be corrected in the implementation by filtering unreliable links or using measured CBT as congestion feedback signal, though its evaluation has been left as future work. The opposite allocations of beaconing rates at low and high powers for neighboring nodes, as for instance in Fig. 4(a) for nodes between 600 and 1100 m compared to those between 1100 and 1600 m, might seem odd. But the effective beaconing rate ($\hat{r}_v$) at both reference distances shows that its combination results in a similar received rate all over the range. It also shows the losses due to collisions and fading, which for Nakagami is an average reduction of 15% with respect to the sum of the transmitted rates.

Finally, we confirm that FABRIC-P fulfills its goal. We have been looking at the allocated beaconing rate, but recall that the explicit goal of FABRIC-P is to maximize the BDR in a fair way. In Fig. 5(a) we show the allocated and effective beaconing rate by LIMERIC+PULSAR and compare it with the effective beaconing rate of FABRIC-P. Looking at the effective beaconing rate the results are quite similar, specially in the central region. However, as shown in Fig. 5(b) where we have plotted the ideal BDR (without losses) and the effective BDR (with hidden-node and fading losses) for both cases, FABRIC-P actually maximizes the dissemination of the beaconing rate of a vehicle and does it in a globally fair way, as induced by α, unlike LIMERIC+PULSAR.

Scenario 3: Differences with previous proposals.

We briefly compare FABRIC-P with conventional approaches [8], [9], [14], where only beaconing rate or transmit power are controlled. In Fig. 6 we show the allocations of FABRIC [14], our previous proposal, for the previous scenarios, using only high power (1000 mW). With FABRIC-P we are extending our previous problem, introducing several new degrees of freedom, with the powers and the minimum rates associated to those powers. It allows to introduce new functionality such as the simultaneous use of applications with different requirements over the same congestion control framework. Let us also note that we can recover our previous

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TVT.2016.2540163

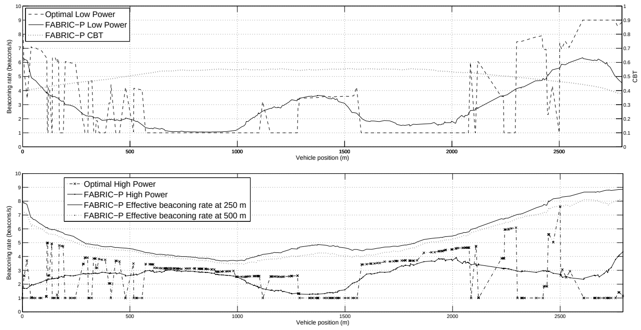

(a) Allocated beaconing rate (r_v), effective beaconing rate ($\hat{r}_v(250)$ and $\hat{r}_v(500)$) and CBT versus vehicle position for free space. Top: Allocated beaconing rate for low power and CBT. Bottom: Allocated beaconing rate for high power and effective beaconing rates.

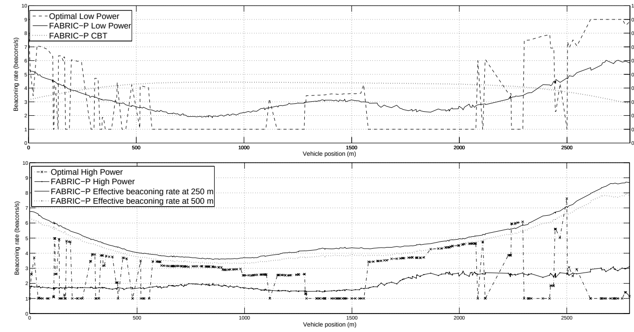

(b) Allocated beaconing rate (r_v), effective beaconing rate ($\hat{r}_v(250)$ and $\hat{r}_v(500)$) and CBT versus vehicle position for Nakagami m=1. Top: Allocated beaconing rate for low power and CBT. Bottom: Allocated beaconing rate for high power and effective beaconing rates.

Fig. 4. Beaconing rate in a realistic multi-hop scenario, with free space and Nakagami m=1, for FABRIC-P asynchronous with $\alpha = 2$. Exact optimal value computed with a numeric solver is shown. Markers show vehicle positions.

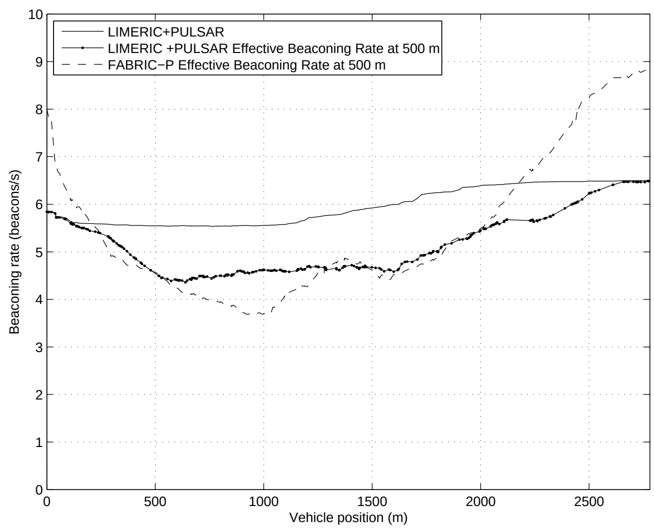

(a) Allocated beaconing rate (r_v) for PULSAR+LIMERIC and effective beaconing rate ($\hat{r}_v(500)$) for FABRIC-P and PULSAR+LIMERIC for free space.

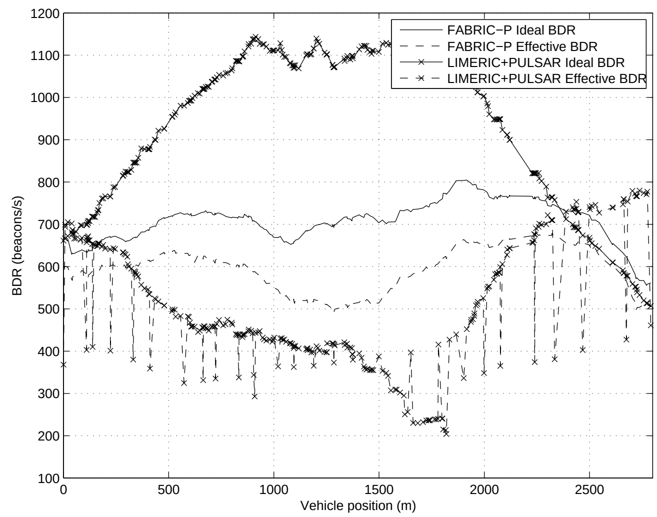

(b) Ideal BDR and effective BDR at 500 m for FABRIC-P and LIMERIC+PULSAR for free space.

Fig. 5. Comparison of beaconing rate and BDR in a realistic multi-hop scenario with free space, for FABRIC-P asynchronous with $\alpha = 2$ and LIMERIC+PULSAR.

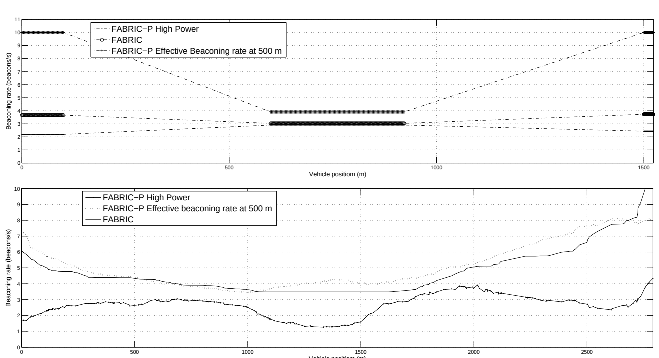

Fig. 6. Comparison of previous scenarios with FABRIC. Beaconing rate versus vehicle position. Top: Comparison of three-cluster scenario. Bottom: Comparison of multi-hop scenario with free space.

problem in [14] as a particular case of this one just by using a single power and setting $\epsilon = 0$, so FABRIC can be also seen as an optimized implementation of FABRIC-P for the single-power case.

In any case, in a moderately congested scenario such as Scenario 2, there are only slight increases in performance: as shown in Fig. 6 bottom, with both approaches up to 500 m the effective beaconing rate is very similar. However, with FABRIC-P the sum of beaconing rates only reaches up to 640.43 m and from then up to 923.24 m just the high power-beaconing rate is received.

However, a highly congested scenario illustrates how the FABRIC-P approach may benefit applications by augmenting the degrees of freedom available: In a conventional approach, a highly congested scenario such as the three-cluster one, shown in Fig 6 top, makes all the vehicles (in middle and edge clusters) reduce the beaconing rate. If a higher one was necessary it would force vehicles to decrease the transmit power as the only solution to increase the beaconing rate, thus reducing the number of neighbors aware of the approaching of vehicles in L and R clusters. On the contrary, FABRIC-P allows to let all vehicles be aware of each other, although with a reduced rate (around 2 or 3 beacons/s up to 923.24 m), while keeping a safer high beaconing rate at short distances (up to 600 m).

V. Dynamic Scenarios

In this section we investigate the FABRIC-P performance in dynamic configurations, where vehicles move. We look mainly at the time evolution of the beaconing rates and CBT. Our goal is to test the ability of FABRIC-P to perform smooth transitions from low to high congestion situations and whether it may support dynamical changes of its parameters.

Scenario 4: Single vehicle and traffic jam.

The first scenario involves (i) a cluster of statically positioned vehicles along a 1500 m road segment, using Poisson positioning with $\rho = 0.14$ vehicles/m, and (ii) a single vehicle approaching the cluster, starting 1320 m away from the first cluster car, and moving at a constant speed of 32 m/s until it passes the cluster. This configuration can model different real scenarios such a highway with a traffic jam in one direction and a single vehicle moving in the opposite direction. Powers and minimum rates have been set as in Scenario 2. The goal of this configuration is to show the dynamics of FABRIC-P when a vehicle switches from none or very few neighbors to a congested state, and back again. Let us note that, with the selected parameters shown in Table II, the channel can accommodate approximately 78 vehicles all in range at the maximum rate. The case of a congested cluster approaching another one is considered in the next configuration. Finally, let us remark again that MAC collisions, hidden node interference and propagation errors are present in these scenarios. Since we are interested in the time evolution of the variables we plot only the results of one replication (all of them show a similar evolution).

In Fig. 7(a) we show the time evolution of the beaconing rate and CBT of the moving vehicle. For free space, the vehicle keeps the initial beaconing rates until it receives beacons from the first neighbors at high power, at t=11 s, and starts increasing the high power beaconing rate and decreasing correspondingly the low power one. As it receives beacons from vehicles at the congested cluster it progressively increases the low power one and decreases the high power one. The latter is increased again at the end, as neighbors at high power become dominant. It shows how FABRIC-P by default behaves in a sensible way, assigning the maximum allowed rate to the power that keeps longer the links with neighbors (high) in absence of congestion. With Nakagami-m the behavior is very similar except that it starts earlier and extends longer since there is a chance to receive beacons from far afar. In both cases, the CBT never exceeds the MBL.

In addition, we provide the results of Fig. 7(b) to show the time evolution of beaconing rates and CBT of the cluster, which are the time evolution of an scenario similar to the one in Fig. 4 Interestingly, FABRIC-P quickly moves rates out of congestion. Moreover, these results confirm that it is not necessary to achieve the optimal allocation to meet the MBL constraint, showing that after only a few steps, 80% of the vehicles experience a load below the MBL. In fact, we have checked that as α increases (not shown in the paper) both beaconing rate and CBT show a steeper reaction to congestion.

Scenario 5: Dynamical settings at intersection.

To conclude, we examine a demanding scenario, where a static cluster of 150 vehicles is set along a 1500 m highway segment oriented north-south (y axis), and it is crossed at the middle position by another highway west-east oriented (x axis). A moving cluster of 100 vehicles, stretching over 500 m, moves from west to east at a constant speed of 32 m/s, starting 2250 m away from the bridge. The moving cluster approaches the static cluster, crosses the intersection, and moves away. We set powers to $[400, 1000]$ mW.

In this scenario we are interested in the capabilities of FABRIC-P to dynamically change its parameters, which provides flexibility in the configuration of the applications on top of it. The availability of multiple powers supports the simultaneous use of several applications. An application would basically select one of the available transmit powers and set the minimum beaconing rate needed to meet its quality of service requirements. How this is done is out of the scope of this paper [12]. As an example, let us assume that there is an intersection safety application which runs in parallel with another one which requires 3.6 beacons/s for the high power. The intersection application requires that in a radius of 200 m of the intersection the vehicles guarantee at least an IRT=0.4 s at target distance of 250 m with a probability of 0.95. Obviously, the higher the power, the lower the beaconing rate requirements to achieve such a reliability. However, the intersection application can also select the low power, 400 mW, which results in a probability of reception at 250 m of 0.9 for Nakagami m=1; and so it sets the minimum beaconing rate for this power to 4 beacons/s in order to satisfy the application IRT requirement [12].

Summarizing, now the moving vehicles set initially the minimum rates to $[1, 3.6]$ beacons/s but when they are within 200 m of the intersection, the minimum rates are switched to $[4, 1]$ beacons/s. The static cluster sets directly the minimum rates according to their distance to the intersection.

In Fig. 8 we show the time evolution of the beaconing rates and CBT for all the vehicles in both clusters. As can be seen in Fig. 8(a) top, vehicles in the moving cluster immediately enforce the minimum rate constraint as they enter the intersection area. This is a consequence of the design of FABRIC-P, the minimum rate constraints (2c) are immediately and locally enforced by Algorithm 2, whereas the MBL constraints are enforced in a distributed way and take more time. In fact, the MBL constraint cannot be met at the intersection with these requirements, as shown in Fig. 8(b). Before the clusters start getting in range of each other, around t=40 s, there is margin for FABRIC-P to increase the beaconing rates above the minimum ones and so it does. These results essentially show that a FABRIC-P vehicle can independently and effectively set its parameters in a dynamical way, and consequently its

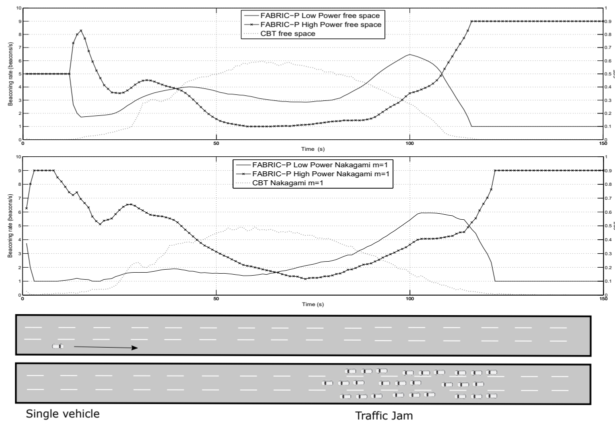

(a) CBT and beaconing rate versus time for a single vehicle approaching the cluster of vehicles at 32 m/s. Top: Free space. Bottom: Nakagami m=1.

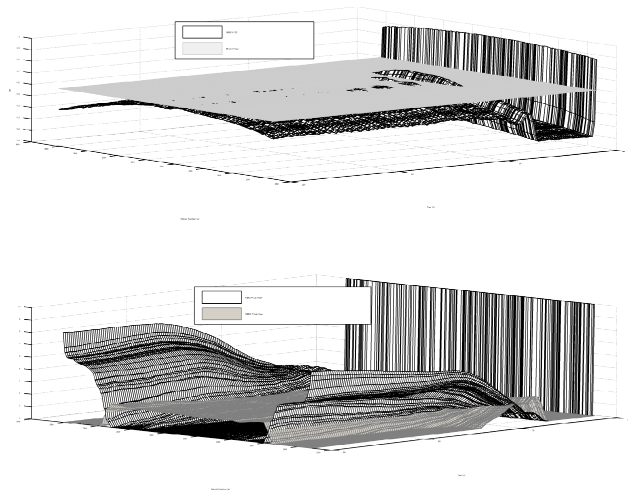

(b) CBT and beaconing rates versus time and vehicle position in the cluster of vehicles with free space propagation. A horizontal plane at z=0.6 has been plotted to mark the MBL.

Fig. 7. Beaconing rate and CBT versus time for a vehicle approaching a cluster of 230 motionless nodes. FABRIC-P asynchronous with $\alpha = 2$.

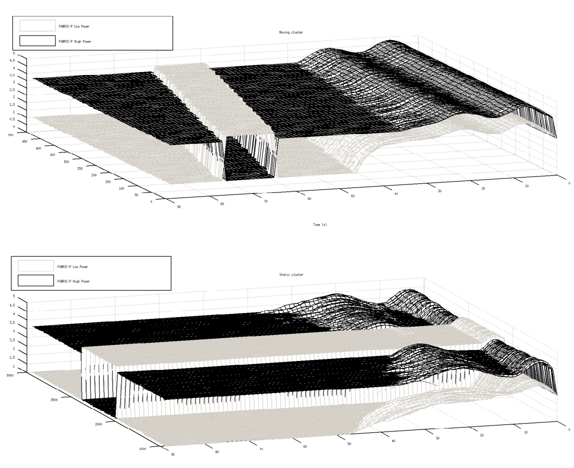

(a) Beaconing rates versus time. Top: Moving cluster. Bottom: Static cluster.

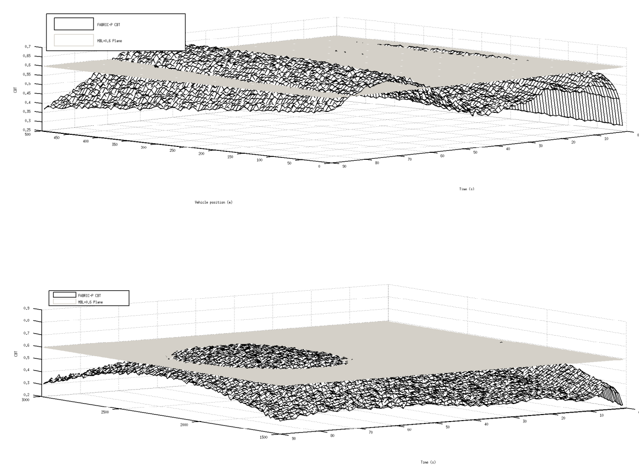

(b) CBT versus time. A horizontal plane at z=0.6 has been plotted to mark the MBL. Top: Moving cluster. Bottom: Static cluster.

Fig. 8. Beaconing rates and CBT versus time for a cluster of 100 nodes crossing an intersection at 32 m/s. Nakagami m=1 propagation. FABRIC-P asynchronous with $\alpha = 2$.

flexibility to support multiple applications on top.

In fact, it behaves in a similar way to other proposals specifically designed to comply with the quality of service of applications. In Fig. 9 we provide a comparison of INTERN [13] with FABRIC-P. INTERN basically assigns the transmit power and rate needed to meet the quality requirements of an application and distributes the excess of channel capacity uniformly via LIMERIC+PULSAR linear assignment. INTERN, with minimum and maximum ΔT_f 1 and 3 beacons/s respectively [13], has been configured to use the high power and 3.6 beacons/s until a vehicle is within 200 m of the intersection, where it switches to low power and 4 beacons/s. For a fair comparison, in Fig. 9 we provide the effective beaconing rate achieved by both of the proposals at 500 m, that is, $\hat{r}_v(500)$, this way, we also provide a view of how losses affect the results shown in Fig. 8. The results show that FABRIC-P supports the safety application with a performance at least as good as INTERN in the intersection, while allocating a higher beaconing rate outside the intersection. Since INTERN assigns the excess of capacity as LIMERIC+PULSAR does, it suffers from the same problem: the allocation is equal for all the vehicles, and therefore it is constrained by the more loaded ones, and the assigned ΔT_f never exceeds 1.5 beacons/s.

VI. Conclusion and Future Work

In this paper we approach the problem of beaconing congestion control in vehicular networks as a NUM rate allocation problem. We assume vehicles can use a set of transmit powers and select a particular beaconing rate for each power in the set, transmitting at multiple powers with different rates during

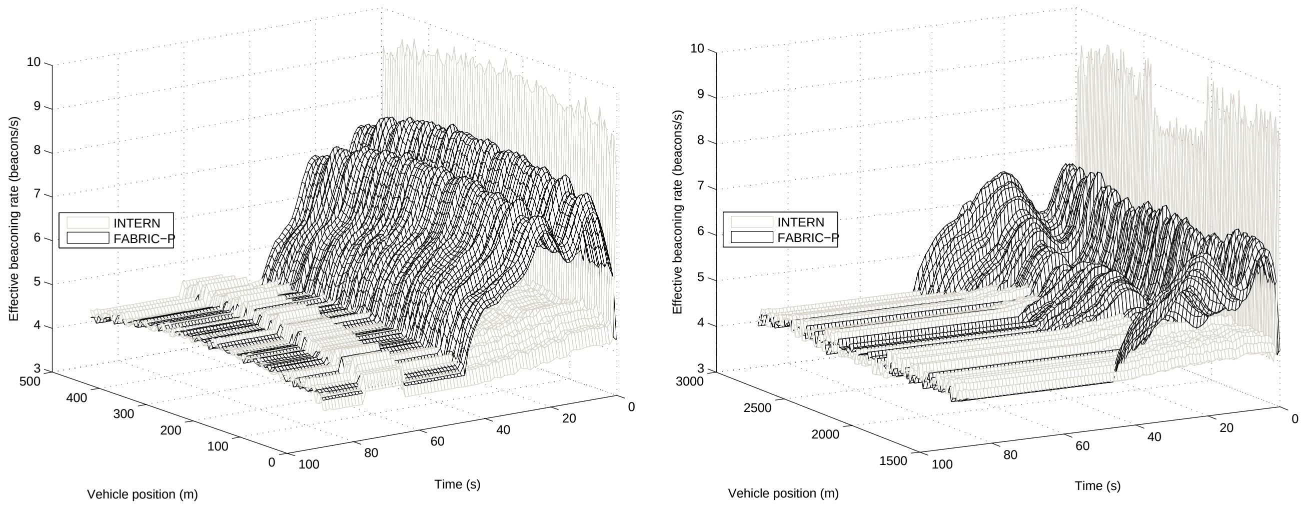

(a) Effective beaconing rates at 500 m ($\hat{r}_v(500)$) versus time for the moving cluster with FABRIC-P and INTERN.

(b) Effective beaconing rates at 500 m ($\hat{r}_v(500)$) versus time for the static cluster with FABRIC-P and INTERN.

Fig. 9. Effective beaconing rates versus time for a cluster of 100 nodes crossing an intersection at 32 m/s. Nakagami m=1 propagation. Comparison of FABRIC-P asynchronous with $\alpha = 2$ and INTERN.

the cycle. The use of multiple powers, in fact, facilitates the execution of multiple applications with different requirements on top of it.

In our NUM model, the utility function of each vehicle is a particular concave function of the BDR, that is, the sum of the beaconing rates over all powers multiplied by the number of neighbors at each power. The shape of the utility function is selected in order to globally enforce a particular type of fairness, controlled by a single parameter, α, which can be independently set by each vehicle.

From this model we propose a particular distributed algorithm, FABRIC-P, which solves the NUM problem, with theoretical and empirical convergence properties and limited signaling overhead. FABRIC-P is remarkably flexible since vehicles can independently and dynamically adapt the algorithm parameters to the requirements of overlying applications.

We have validated FABRIC-P by exhaustive simulations in both static and dynamic scenarios, for different position distributions and propagation models, showing that FABRIC-P effectively controls the congestion while generating fair beaconing rate allocations. Our results have been compared with LIMERIC+PULSAR, a relevant rate allocation scheme in vehicular networks, and INTERN, an awareness control algorithm. They show that FABRIC-P can be used by safety applications that require to dynamically achieve a minimum reliability with equal or better performance. Simulations also confirm that the algorithm is robust against packet losses due to collisions or fading.

There are still a number of practical considerations and implementation alternatives that can be evaluated in order to tune the algorithm. First, the β parameter controls the convergence speed and the amplitude of fluctuations and there is a wide range of possible values meeting the convergence condition to test. Second, also related to the convergence speed, which is expressed in algorithm steps, we have used steps of $T_s = 1$ s in our tests but the interval can be safely be reduced to 500 ms or below, which would reduce the absolute time to convergence. Since the feedback is collected during T_s it is interesting to find how much it can be reduced and how it may affect the stability. Finally, filtering of unreliable links may provide a more accurate measurement of congestion and neighbor estimation in fading scenarios. Even the use of alternative congestion measurements such as the measured CBT can be tested. We intend to carry out an extensive evaluation of these matters in a future work.

The possibility of executing in parallel multiple applications with different quality of service requirements introduces the question of how to manage the priorities and allocate resources to them when the needs are in conflict. We intend to look at this question too in the future. Additionally, from a more general perspective, we have shown how different values of the fairness parameter α result in different allocations, which may be more adequate depending on the intended application or scenario. One of the key advantages of FABRIC-P and our approach is that the fairness allocations can be controlled with this single parameter. Moreover, this approach allows to use different values for each vehicle or even to use totally different utility functions, which can be both dynamically changed. And vehicles do not need to know the functions or values used by other vehicles. Consequently, as a future work, we intend to further explore variations of the discussed problem in the context of vehicular networks. In particular, a comparative application and evaluation of alternative fairness notions and the introduction of heterogeneous utility functions and constraints in the problem.

Appendix A
Proof of Convergence of Algorithm 1

Convergence of the overall scheme is guaranteed as long as the dual iteration is able to find the optimal vehicle prices π^*, that is, the optimal solution of dual problem (5), whose objective function is called the dual function. Since the problem enjoys strong duality, and the objective function of the original problem is strictly concave (thanks to the regularization term), the dual function is everywhere differentiable. Since the original problem has a finite solution, the dual function has a finite minimum $g(\pi^*) > \infty$, and $g(\pi^*)$ equals the optimal utility of the original problem. Finally, in order to guarantee the convergence of Algorithm 1, $\nabla g(\pi)$ must be a Lipschitz continuous function with finite Lipschitz constant H [21][Prop. 3.4]. Therefore we have to prove that

$$\|\nabla g(x) - \nabla g(y)\|_2 \leq H\|x - y\|_2 \quad \forall x, y \geq 0 \tag{7}$$

Now, given any $x, y \geq 0$, using Taylor theorem we have that $\|\nabla g(x) - \nabla g(y)\|_2 \leq \|\nabla^2 g(z)\|_2 \|x - y\|_2$ for some $z = ty + (1-t)x \geq 0, t \in [0,1]$. Therefore, we have to find a bound for the Euclidean norm of the Hessian[1] $\|\nabla^2 g(z)\|_2$. The Euclidean norm is bounded by the product of the maximum column and row sum of the Hessian matrix, that is, $\|\nabla^2 g(z)\|_2^2 \leq \|\nabla^2 g(z)\|_1 \|\nabla^2 g(z)\|_\infty$. And since the Hessian is symmetric, $\|\nabla^2 g(z)\|_1 = \|\nabla^2 g(z)\|_\infty$, and we have that $\|\nabla^2 g(z)\|_2 \leq \|\nabla^2 g(z)\|_1$.

To compute the Hessian we first write the gradient of the dual function (6) in matrix form, $\nabla g(\pi) = \sum_p N^p R^p - C$, where each N^p is a $V \times V$ symmetric neighbor matrix whose components $(N^p)_{vj} = 1$ if vehicle j is in range of vehicle v at power p and $(N^p)_{vj} = 0$ otherwise. Matrices R^p are $V \times 1$ column vectors whose components $(R^p)_v = r_{vp}$ are the beaconing rate at power p for each vehicle. And C is a $V \times 1$ column vector whose components $(C)_v = C$ are the MBL for each vehicle. The Hessian is the $V \times V$ matrix whose components are the partial derivatives of $\nabla g(\pi)$, with respect to $\pi_v, \forall v$. That is, an arbitrary column j of the Hessian is given by

$$\frac{\partial}{\partial \pi_j}\nabla g(\pi) = \sum_p N^p \frac{\partial R^p}{\partial \pi_j} = a_{vj} \quad v = 1 \ldots V \tag{8}$$

where the partial derivative of a matrix is made of the partial derivatives of its components, that is, $\frac{\partial R^p}{\partial \pi_j} = (\frac{\partial r_{vp}}{\partial \pi_j})_v$. Since $\|\nabla^2 g(z)\|_2 \leq \|\nabla^2 g(z)\|_1$ we have to derive an upper bound for the sum of the elements of an arbitary column of the Hessian (8), $\|\nabla^2 g(z)\|_1 = \max_j \sum_v |a_{vj}|$. We know that $r^*_{vp}(\pi)$ are the elements of the unique maximizer of (4) given a vector or prices π. Performing the sum and multiplication of (8) and rearranging the terms, the sum of absolute value of the elements of an arbitrary column becomes

$$\sum_v |a_{vj}| \leq \sum_v \sum_p n_{vp} |\frac{r^*_{vp}}{\partial \pi_j}| \leq \sum_v \sum_p \bar{V} |\frac{r^*_{vp}}{\partial \pi_j}| \tag{9}$$

where the inequality comes from the subadditivity of absolute values and we define $\bar{V}$ as the maximum number of neigbors any vehicle has at any given power.

[1]For simplicity we assume that the Hessian exists everywhere. At points where it may not exist derivatives should be replaced by convex subgradients.

To compute the partial derivatives, we form the Lagrangian function of (4) for a given vehicle v relaxing all the constraints

$$\begin{aligned} Lg(r) = &\frac{(\sum_p n_{vp} r_{vp})^{1-\alpha}}{1-\alpha} - \epsilon \sum_p r_{vp}^2 - \\ &- \sum_v \pi_v \left(\sum_p \sum_{v' \in n(v)_p} r_{v'p} - C \right) - \\ &- \sum_p \lambda_p (r_{vp}^{min} - r_{vp}) - \mu (\sum_p r_{vp} - R^{max}) \end{aligned} \tag{10}$$

From the KKT stationarity condition for (4), all the partial derivatives with respect to each r_{vp} must be equal to zero

$$\begin{aligned} &\frac{\partial Lg(r^*)}{\partial r^*_{vi}} = n_{vi} (\sum_p r^*_{vp})^{-\alpha} - 2\epsilon r^*_{vi} - \Pi_i + \lambda_i - \mu = 0, \\ &i = 1 \ldots P \end{aligned} \tag{11}$$

and since the maximizer are functions of the prices $r^*_{vp} = f(\pi)$ we take the partial derivative with respect to some π_j of the above expression, obtaining

$$\frac{(-\alpha) n_{vi}}{(\sum_p r^*_{vp})^{\alpha+1}} (\sum_p \frac{\partial r^*_{vp}}{\partial \pi_j}) - 2\epsilon \frac{\partial r^*_{vi}}{\partial \pi_j} - 1 = 0, \quad i = 1 \ldots P \tag{12}$$

where from now on we call $K = \frac{\alpha}{(\sum_p r^*_{vp})^{\alpha+1}}$ to simplify notation. For any given vehicle v the P equations (12) define a linear system of equations with P variables $\frac{\partial r^*_i}{\partial \pi_j}$, $\quad i = 1 \ldots P$, where we remove the v index for the sake of clarity. The system written in matrix form is $Bx - 1 = 0$, where x is a $P \times 1$ column vector whose components $x_i = \frac{\partial r^*_i}{\partial \pi_j}$ are the variables, 1 is a $P \times 1$ column vector whose components are 1 and B is a $P \times P$ square matrix whose components are

$$B_{ij} = \begin{cases} -n_i K - 2\epsilon, & i = j \\ -n_i K, & i \neq j \end{cases} \tag{13}$$

The solution of the system is $x = B^{-1}1$ as long as the inverse of B exists. To check the existence of the inverse we compute the determinant of B. Taking advantage of the special structure of B the determinant can be transformed to lower triangular form by first adding all the rows to the first one and then substracting the first column to the rest of the columns. Then the determinant of B is

$$Det(B) = (-2\epsilon)^{P-1} (-2\epsilon - K \sum_p n_p) \tag{14}$$

Since $\epsilon > 0$, $Det(B) \neq 0$ as long as at least one $r^*_p \neq 0$, and the inverse exists. The inverse matrix can be computed directly from the adjugate matrix yielding

$$B^{-1}_{ij} = \begin{cases} \frac{(-2\epsilon)^{P-1}(-2\epsilon - K(\sum_p n_p - n_i))}{|Det(B)|}, & i = j \\ \frac{-n_i K (2\epsilon)^{P-2}}{|Det(B)|}, & i \neq j \end{cases} \tag{15}$$

And finally, the 1-norm of the system solution is $\|x\|_1 = \sum_p |\frac{\partial r^*_p}{\partial \pi_j}| \leq \|B^{-1}\|_1 P$, which we substitute in (9) yielding $\sum_v |a_{vj}| \leq V\bar{V}P\bar{B}$, where $\bar{B} = \frac{1}{|Det(B)|}[| -$

$(2\epsilon)^{P-2}K|(\sum_p |n_p| - n_j) + |(-2\epsilon)^{P-1}(-2\epsilon - K(\sum_p n_p - n_j)||$.

Therefore, we conclude that $\nabla g(\pi)$ is Lipschitz continuous with $H = V\bar{V}P\bar{B}$. From this follows that Algorithm 1 converges to the optimal values provided $0 < \beta < \frac{2}{H}$. Let us note that the bound for β depends on the α and ϵ in use.

Acknowledgment

This research has been supported by the MINECO/FEDER project grant TEC2013-47016-C2-2-R (COINS) and partially supported by the Spanish project grants TEC2014- 53071-C3-1-P (ONOFRE) and TEC2015-71932-REDT (ElasticNetworks)

References

[1] H. Hartenstein and K. P. Laberteaux, "VANET. Vehicular Applications and Inter-Networking Technologies", Wiley, 2010.
[2] ETSI EN 302 663, "Intelligent Transport Systems (ITS); Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band", V1.2.1, July 2013
[3] IEEE Std 802.11-2012, "IEEE Standard for Information technology–Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007), pp. 1-2793, 2012.
[4] G. Karagiannis, O. Altintas, E. Ekici, G. Heijenk, B. Jarupan, K. Lin, and T. Weil, "Vehicular Networking: A Survey and Tutorial on Requirements, Architectures, Challenges, Standards and Solutions," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 4, pp. 584-616, 2011.
[5] ETSI EN 302 637-2, "Intelligent Transportation Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service", v.1.3.1, September 2014.
[6] M. Sepulcre, J. Mittag, P. Santi, H. Hartenstein, and J. Gozalvez, "Congestion and Awareness Control in Cooperative Vehicular Systems," *Proceedings of the IEEE*, vol. 99, no. 7, pp. 1260-1279, 2011.
[7] K, Tian Lan, M. Chiang, A. Sabharwal, "An Axiomatic Theory of Fairness in Network Resource Allocation," *Proceedings IEEE INFOCOM*, pp.1–9, 2010.
[8] J. B. Kenney, G. Bansal and C. E. Rohrs, "LIMERIC: A Linear Adaptive Message Rate Algorithm for DSRC Congestion Control," *IEEE Transactions on Vehicular Technology*, vol. 62, no 9, pp. 4182–4197, 2013.
[9] T. Tielert, D. Jiang, Q. Chen, L. Delgrossi and H. Hartenstein, "Design Methodology and Evaluation of Rate Adaptation Based Congestion Control for Vehicle Safety Communications," *Vehicular Networking Conference (VNC), 2011 IEEE*, pp. 116-123, 2011.
[10] Y. P. Fallah, C. L. Huang, R. Sengupta and H. Krishnan, "Analysis of Information Dissemination in Vehicular Ad-Hoc Networks With Application to Cooperative Vehicle Safety Systems," *IEEE Transactions on Vehicular Technology*, vol. 60, no 1, pp. 233–247, 2011.
[11] C.-ling Huang, Y. P. Fallah, R. Sengupta, and H. Krishnan, "Adaptive Intervehicle Communication Control for Cooperative Safety Systems," *IEEE Network*, vol. 24, pp. 6-13, 2010.
[12] T. Tielert, D. Jiang, H. Hartenstein and L. Delgrossi , "Joint Power/Rate Congestion Control Optimizing Packet Reception in Vehicle Safety Communications," *Proceedings of the tenth ACM international workshop on Vehicular inter-networking, systems, and applications*, pp. 51-60, 2013.
[13] M. Sepulcre, J. Gozalvez, O. Altintas and H. Kremo , "Adaptive Beaconing for Congestion and Awareness Control in Vehicular Networks," *Vehicular Networking Conference (VNC), 2014, IEEE*, pp. 81-88, 2014.
[14] E. Egea-Lopez and P. Pavon Marino, "Distributed and Fair Beaconing Rate Adaptation for Congestion Control in Vehicular Networks," *IEEE Transactions on Mobile Computing*, in press, doi: 10.1109/TMC.2016.2531693
[15] Kelly, Frank, "Charging and rate control for elastic traffic." European transactions on Telecommunications 8.1 (1997): 33-37.
[16] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 556–567, 2000.
[17] ETSI TS 102 687, "Intelligent Transport Systems (ITS); Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range; Access Layer part", v1.1.1, July 2011.
[18] I. Ahmed, A. Mohammed, H. Alnuweiri, "On the fairness of resource allocation in wireless mesh networks: a survey," *Wireless Networks*, vol. 19(6), pp. 1451–1468, 2013.
[19] E. Egea-Lopez, J. J. Alcaraz, J. Vales-Alonso, A Festag and J. Garcia-Haro, "Statistical Beaconing Congestion Control for Vehicular Networks," *IEEE Transactions on Vehicular Technology*, vol. 62, no. 9, pp. 4162–4181, 2013.
[20] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.
[21] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, Inc., Upper Saddle River, NJ, 1989.
[22] L. Cheng, B. E. Henty, D. D. Stancil, F. Bai, and P. Mudalige, "Mobile Vehicle-to-Vehicle Narrow-Band Channel Measurement and Characterization of the 5.9 GHz Dedicated Short Range Communication (DSRC) Frequency Band," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 8, pp. 1501–1516, 2007.
[23] JOM (Java Optimization Modeler), avaliable at http://www.net2plan.com/jom/
[24] A. Wächter and L. T. Biegler,"On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming," *Mathematical Programming*, 106(1), pp. 25–57, 2006
[25] OMNeT++ simulation library, avaliable at http://www.omnetpp.org

Esteban Egea-Lopez received the Telecommunications Engineering degree in 2000, from the Universidad Politecnica de Valencia (UPV), Spain, the Master Degree in Electronics in 2001, from the University of Gavle, Sweden, and Ph.D. in Telecommunications in 2006 from the Universidad Politecnica de Cartagena (UPCT). He is an associate professor with the Department of Information Technologies and Communications at the Universidad Politecnica de Cartagena (UPCT). His research interest is focused on vehicular networks and MAC protocols.

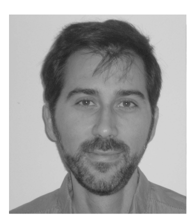

Pablo Pavon-Mariño received the M.Sc. degree in Telecommunication Engineering in 1999 from the University of Vigo, Spain. In 2000, he joined the Universidad Politecnica de Cartagena, Spain, where he is Associate Professor in the Department of Information Technologies and Communications. He received the Ph.D. degree from this University in 2004, and a M.Sc. degree in Mathematics in 2010. His research interests include planning and optimization of communication networks.