HAL
open science

Delivering Video-on-Demand services with IEEE 802.11p to major non-urban roads: A stochastic performance analysis

Thomas Begin, Anthony Busson, Isabelle Guérin-Lassous, Azzedine Boukerche

▶ **To cite this version:**

Thomas Begin, Anthony Busson, Isabelle Guérin-Lassous, Azzedine Boukerche. Delivering Video-on-Demand services with IEEE 802.11p to major non-urban roads: A stochastic performance analysis. Computer Networks, 2020, 182, pp.1-12. 10.1016/j.comnet.2020.107440 . hal-02923337

HAL Id: hal-02923337

https://hal.science/hal-02923337v1

Submitted on 7 Sep 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Delivering Video-on-Demand services with IEEE 802.11p to major non-urban roads: a stochastic performance analysis

Thomas Begin[a,*], Anthony Busson[a], Isabelle Guérin Lassous[a], Azzedine Boukerche[b]

[a] *Univ Lyon, Université Claude Bernard Lyon 1, ENS de Lyon, Inria, CNRS, LIP, France*
[b] *DIVA Lab, University of Ottawa, Canada*

Abstract

In this paper, we introduce a relatively simple and versatile analytical model to compute the performance of an IEEE 802.11-based Video-on-Demand (VoD) service for connected vehicles traveling along a highway or an expressway. We use a probabilistic approach to account for the intrinsic dynamics (mobility) of vehicles, the density of subscriber vehicles, the distance between Road Side Units (RSUs), as well as the video codec bit rate. Our solution delivers estimates for the attained throughput of each subscriber vehicle as well as for their interruption times in the playback of their video. We studied several scenarios to assess the accuracy of our approximate solution comparing its results with those of a discrete-event simulator with realistic vehicular traffic trajectories, video traffic, and networking protocols. Overall, our solution is found to be accurate delivering estimates in good agreement with the simulation results. Through several applications, we illustrate how our analytical solution can help in the design of a VoD service for vehicles for capacity planning purposes, or at improving its performance with the use of a blocking policy.

Keywords: IEEE 802.11p, Video-on-Demand, Performance Evaluation, Stochastic modeling, SUMO, ns-3.

1. Introduction

The market of connected cars is rapidly expanding worldwide, primarily driven by the prospect of safety and security features. Another leverage to this growth is the potential development of infotainment services for the passengers of connected vehicles. A good case in point is video streaming. Connected vehicles subscribe to a Video-on-Demand (VoD) service and, in return, they get access to a library of movies and music in a similar way than air passengers

*Corresponding author

do on commercial flights. Videos can be streamed on a screen embedded in the vehicle, or alternately, on the passenger's own devices (e.g., smartphone, laptop).

A fundamental yet challenging task is to bring connectivity to the moving vehicles. In the architecture for Wireless Access in Vehicular Environment (WAVE), IEEE promotes the use of the 802.11p amendment to the 802.11 standard to implement robust and safe communications from/to the connected vehicles. Unlike solutions based on broadband cellular technologies (e.g. 4G LTE, 5G NR), IEEE 802.11p operates over the unlicensed 5.9 GHz band, which can prove to be a valuable asset. Thus, the list of potential stakeholders for providing a VoD service is not restricted to mobile network operators (MNOs). Furthermore, a VoD provider has then a full control over its networking architecture and can operate without being charged (hefty) fees by MNOs.

In the WAVE architecture, connected vehicles are regularly associated and exchange data with networking devices scattered along the road and called Road Side Units (RSUs). However, sizing the needed amount of network resources, namely RSUs, to provide a good quality of experience for a VoD service is a difficult issue, in part because it is hard to forecast the attained throughput of a vehicle. Indeed, the data rate of the radio channel between a vehicle and a RSU varies with the distance from each other but also strongly depends on the number and location of the other vehicles associated to the same RSU. A simple analytical tool can help a VoD operator design its networking architecture and test various scenarios under different assumptions with regards to the vehicle density (aka concentration), the market penetration rate, the video codec, the number of RSUs, etc.

The contribution of this paper is to propose a relatively simple and versatile solution to the performance computation of an IEEE 802.11-based VoD service for connected vehicles traveling along a highway or an expressway. We use a probabilistic approach to account for the intrinsic dynamics (mobility) of vehicles and to deliver estimates for the attained throughput of each subscriber vehicle as well as for their interruption times in the playback of their video. Numerous discrete-event simulations with realistic traffic trajectories, video traffic, and networking protocols were performed and demonstrate the good accuracy of our approximation. Through several applications, we illustrate how our analytical solution can help in the design of a VoD service for vehicles (capacity planning) and at improving its performance with the use of a blocking policy.

The remainder of this paper is organized as follows. Related works are discussed in the next section. In Section 3, we describe in detail the considered VoD platform. Section 4 presents our proposed simple analytical solution to evaluate the expected attained throughput for each vehicle as well as their interruption time in the playback of their video. The experimental settings regarding vehicle traffic and video traffic are discussed in Section 5. Section 6 is devoted to the numerical validation of our proposed performance model and its solution. We describe in Section 7 possible applications of our analytical solution for the purpose of capacity planning and efficient blocking policies. Section 8 concludes this paper.

2. State of the art

In the very recent paper [1], the authors consider a video streaming architecture in which V2I (Vehicle-to-Infrastructure) and V2V (Vehicle-to-Vehicle) communications are possible. V2V communications are used either for vehicles that experience weak communications with the infrastructure, or for vehicles that freely decide to download the video both from the infrastructure and from neighbor vehicles. In their work, the authors assume that communications are based on the OFDMA (orthogonal frequency-division multiple access) technique. In [2, 3], the authors study video streaming from the resource allocation point of view. In both works, they introduce their own time slot-based MAC (Medium Access Control) protocol. This allows them to provide heuristics to select the quality level of each video as well as to allocate time slots in order to optimize a utility function that expresses the video quality experienced at vehicles. In the same spirit, the authors of [4] design a solution that determines the video quality to use, which data to load in the cache of the RSU, and the radio resources to allocate to each vehicle. As for the two previous works, the time is slotted and slots are allocated to vehicles. In [5], the authors propose a double-buffer system to reduce the delay effect due to frequent handovers in a 5G network. In all these works, the used technologies and the considered link layer imply a share of the radio medium between the vehicles different from the one we consider in our work wherein video frames are transmitted with the 802.11p technology which uses the CSMA/CA approach. Moreover, the objectives of these works differ from ours. For example, the goal of the solution proposed in [5] is to reduce the effect of the frequent handovers due to small distances between RSUs paving the whole roads whereas, in our scenarios, the number of RSUs are limited and vehicles will experience periods without connectivity to the video infrastructure. In [2, 4], the authors aim to optimize resource utilization and the selection of video quality layers for each vehicle in real time whereas the goal of our study is to provide help for dimensioning the networking resources of an IEEE 802.11p-based VoD platform.

Video dissemination is not restricted to video streaming applications. It includes other applications like roadside video emergency (e.g., warning drivers of a dangerous situation happening on the road ahead), video surveillance or entertainment. In general, these applications take advantage of the presence of vehicles that can directly exchange video data when they are within radio communication range. The vehicles then form what is called a vehicular ad hoc network (VANET). Numerous works have studied how to efficiently disseminate video in VANETs. In a recent paper [6], the authors survey the different approaches that have been proposed to enable video streaming in VANETs. The authors analyze the proposed solutions via the QoS (Quality of Service) / QoE (Quality of Experience) that can be provided to this application in these networks. The main approaches investigated in this paper and that can be also found in other papers are: identifying and using video coding adapted to the VANET features (e.g., [7, 8]); efficiently broadcasting videos through a judicious selection of the relay nodes and limiting the transmissions of control frames (e.g.,

[9, 10]); using a routing protocol or multipath routing protocol adapted to the possible dynamics and density / sparsity of the VANET (e.g., [11, 12, 13, 14]). In these studies, video dissemination solutions tend to often make use (partially or exclusively) of V2V (Vehicle-to-Vehicle) communications, unlike our work. Furthermore, these video dissemination studies focus on a given multimedia message or video that has to be delivered to all considered vehicles, whereas, in our work, vehicles request different videos.

Other studies analyze the performance of V2I communications when the infrastructure sends safety or road information messages. This is, for instance, the case in [15, 16, 17]. But in these papers, the used communication technology is LTE and the considered traffic have a profile different from the one of the video traffic. In [18], the authors consider the Wi-Fi technology, but they focus their study on the access procedure corresponding to the steps preceding the possible data exchanges, whereas our study mostly focuses on the performance behavior (exchange of video data) once the access procedure has been done.

Overall, the works that come closest to ours are the ones presented in [19, 20]. In [20], the considered video streaming infrastructure is almost the same, except that the RSU coverage areas overlap whereas, in our work, they do not. Moreover, contrary to our work, in [20], no modeling is proposed. The video streaming infrastructure is evaluated by simulations, with the ns-2 simulator interfaced with the SUMO module. The evaluation is focused on the packet losses due to handovers, whereas we consider different performance evaluation parameters (namely the download rate and the interruption time). In [19], the video streaming infrastructure under study is the same as the one we consider in our work. However, the modeling approach as the validation part are very different. In [19], each vehicle is considered individually in the modeling making the approach hardly scalable when the number of vehicles grows. On the other hand, in the current paper, the proposed model is based on the vehicle density (possibly different for each vehicle type) naturally lending itself to a stochastic modeling. It follows that the performance model can handle virtually any number of vehicles. Furthermore, it can take into account much more vehicles than in [19]. Moreover, in this paper, we validate our proposed model on a more realistic simulator than in [19]. Indeed, we use the SUMO simulator (Simulation of Urban MObility) [21] to simulate realistic vehicular trajectories that are then injected into the ns-3 simulator, whereas in [19], the proposed solution is only validated with ns-3.

3. VoD platfom description

We consider a VoD platform devised to deliver video content to car or bus travelers along a highway or an expressway. Typically, customers would subscribe to an online VoD service so that they can enjoy watching videos while making progress on their journey to their destination. Figure 1 depicts the associated scenario where multiple vehicles access a remotely hosted video library database using radio communication. We now discuss in more details the different components involved in the considered scenario.

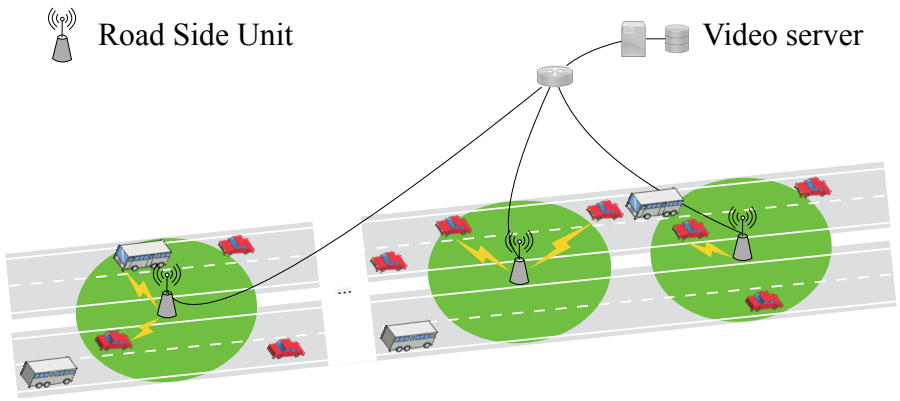

Figure 1: IEEE 802.11p-based infrastructure to deliver VoD service to vehicles on a road. Only vehicles within the green zones can currently download their video from the video server.

3.1. Roads and vehicles

The portions of road involved in the VoD platform can be of different types. It can comprise one, two or more lanes in each direction, and be shared by various classes of vehicles such as cars, buses, and trucks. Each class of vehicles has its maximum speed limit (typically lower for trucks and buses). Furthermore, each vehicle determines its own target speed (below the speed limit of its class) that would be its actual speed if the vehicle was left alone on the road. However, due to the vehicle density on the road, vehicles may have to slowdown, change lanes, and accelerate resulting in varying speeds and overall lower speed values. Finally, in this study, we consider two classes of vehicles: the class of cars and the class of trucks and buses. We let λ (resp. μ) denote the car (resp. bus and truck) density on the road.

3.2. Road side units

RSUs (Road Side Units) are networking devices deployed along the roads to provide connectivity to mobile (subscriber) vehicles. They serve as gateways to the high-speed wired infrastructure hosting the video server and its database, as well as APs (Access Points) for the passing by vehicles. Any vehicle within an RSU signal range can exchange data with the corresponding RSU (though the data rate may vary depending on the current context). Given their cost of deployment, the number of deployed RSUs is limited and RSUs are located so that that their signal range do not overlap. Without loss of generality,we assume that RSUs are deployed at a constant distance from each other along the road. We denote by K the number of RSUs, by I the inter-distance between two successive RSUs, and by L the signal range of each RSU.

3.3. Subscriber vehicles

Subscriber vehicles represent the subset of the vehicles traveling along the road, be it cars, buses or trucks, that subscribe to the VoD service. We use x to denote the penetration rate, i.e., the percentage of vehicles that have subscribed to the VoD service. They are equipped with an IEEE 802.11p interface to enable radio communication with the RSUs. Whenever a subscriber vehicle enters into the signal range of an RSU, it associates with the corresponding RSU so that the video server can start or resume its transmission of video frames to the vehicle. As long as vehicles stay within the signal range of an RSU, they download video frames. If a vehicle downloads video frames faster than it plays them on its media player, then frames are placed in a buffer for later use. On the contrary, if the download speed of a vehicle is less than its speed of play or even null (i.e., the vehicle is out of range of any RSU), the media player fetches video frames, if any, from the buffer and the playback goes on. Otherwise, if the buffer queue is depleted, the media player will simply stall. Because RSU signals do not overlap, it is critical for subscriber vehicles to build up queues of unplayed video frames during the time they stay within an RSU signal range to overcome non-connectivity periods between RSUs. Ideally, at its departure of an RSU, a vehicle should have enough video frames to be played until it reaches the next RSU. Finally, note that no playout delay is added to the video frames. It means that the video is played by the media player as soon as the first frame is received by the subscriber.

3.4. IEEE 802.11p

Radio communications between the RSUs and vehicles operate over IEEE 802.11p. Our proposed architecture does not involve V2V (Vehicle-to-Vehicle) communications, nor priority rules (e.g., EDCA) and that all RSUs transmit over the same radio channel. However, errors in frame transmissions occur typically at a rate of around 5% and 10% due to the high mobility of the vehicles. Because RSU signal ranges do not overlap and VoD traffic mostly consists of download traffic sent from RSUs to the vehicles, there is virtually no risk of interference and colliding frames (RSUs do not compete to access the radio channel) even if they use the same radio channel. Clearly, this latter assumption becomes invalid in urban scenarios where RSUs are typically at much shorter range from each other. Note also that a vehicle can serve as a transitory propagation obstacle for the communication of another vehicle if the former lies on the line of sight of the latter with the RSU. However, in practice, such occurrences are limited due to the high position of RSUs (higher than the road level) and the mobility of vehicles.

The 802.11p standard relies on the CSMA/CA mechanism to control access to the radio channel. In practice, RSUs and vehicles obey to the DCF (Distributed Coordination Function) mechanism to transmit a frame. Assuming a node has a frame to be sent, it first needs to sense the radio channel as idle for a given time period AIFS. Then, the node chooses a random waiting time within its current contention window and postpones its transmission by this value. At

the end of this backoff, the node can transmit its frame if the medium is sensed as idle and automatically waits for an ACK returned from the receiver. Shall the ACK not return, the sender will retransmit the frame using a larger contention window. For more details about DCF, the interested reader can refer to [22].

Note that transmissions of frames can occur at different rates depending on the radio channel assessment (typically estimated by a frame error rate [23, 24]). An 802.11p sender can dynamically adapt its transmission rate within the range of 3 and 27 Mbps. In our study, since vehicles are moving to and away from RSUs, we assume that their transmission rates are expected to start at the lowest value, then gradually increase up to the apex value before returning to the lowest value. More specifically, let denote by Z the total number of zones covering the communication area of an RSU. In each zone, a single transmission rate is used between the RSU and the vehicles located in the zone. Two consecutive zones are associated to two different transmission rates. We use h_j $(j = 1, ..., Z)$ to indicate the length of j-th zone as illustrated by Figure 2 and T_j to indicate the transmission rate used in the j-th zone. Therefore, we have: $\sum_{j=1}^{Z} h_j = 2L$. Note that the exact length of the intervals and their number may depend on the propagation characteristics of the radio channel as well as on the used rate adaptation algorithm.

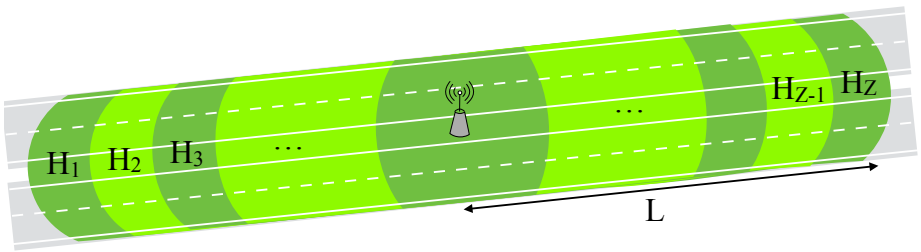

Figure 2: There is a total of Z zones with different transmission rates within an RSU signal range. L denotes the signal range of the RSU while H_j $(j = 1, ..., Z)$ refers to the j-th transmission zone of length h_j and associated to a transmission rate T_j and an achievable throughput A_j.

In order to account for the overheads of 802.11p (e.g., AIFS, Backoff, ACK), we define the achievable throughput in a zone j, denoted by A_j, as the maximum throughput a vehicle in this zone can attain if it had an exclusive access to the radio channel (i.e., not sharing the radio channel with others). By definition, we have: $A_j < T_j$. Note that A_j is easily obtained for any given values of T_j and packet length (e.g., [22]). Table 1 represents what could be the negotiated transmission rate (T_j) and its corresponding achievable throughput (A_j) for a vehicle according to its distance to an RSU using the Ideal WiFi manager of ns-3. Note that this WiFi manager makes use of six possible transmission rates.

Nonetheless, in general, neither T_j, nor A_j are fair estimates for the attained throughput of a vehicle i, denoted by B_i. Indeed, the value of B_i strongly depends on the current location of vehicle i as well as on the other subscriber's.

This relates to the phenomenon known as the performance anomaly [25] where senders with high transmission rates are largely penalized by senders with low transmission rates. Clearly, for any vehicle i located in the j-th zone, we have $B_i \leq A_j < T_j$. However, the intrinsic randomness of DCF combined to the interactions with competing vehicles makes it hard to accurately forecast the actual value of B_i unless using an adapted analytical model (or accurate simulations).

Zone number	h_j (m)	T_j (Mbps)	A_j (Mbps)
H_1	155	3	2.347
H_2	110	6	4.365
H_3	105	9	6.070
H_4	65	12	7.609
H_5	85	18	10.021
H_6	170	27	12.820
H_7	85	18	10.021
H_8	65	12	7.609
H_9	105	9	6.070
H_{10}	110	6	4.365
H_{11}	155	3	2.347

Table 1: Description of the 11 communication zones of IEEE 802.11-based RSU when using the Ideal WiFi manager of ns-3. For each zone j, we indicate its length h_j, transmission rate T_j, and achievable throughput A_j computed for a packet size of 1500 bytes.

3.5. Wired network and videos

Videos are stored on a back-end database server located on a high-speed network. Modern codecs generate video frames with highly variable sizes. In fact, the size of a video frame may be just several bytes as well as thousands of bytes. Nonetheless, each codec usually comes with an average bit rate that, in our scenario, would correspond to the mean throughput (download rate) that a vehicle would need to correctly play the video. Video frames delivered by the server are encapsulated into constant size IP datagrams (of maximum length) before being transferred over the network. Therefore, an IP datagram may contain only a part of a video frame or, alternately, several of them.

Once a subscriber chooses its video content, IP datagram encapsulating video frames are routed first over the high-speed wired infrastructure up to the appropriate RSU and then over the radio channel between the RSU and the subscriber vehicle. Note that this assumes that a table used for routing purpose maintains the current associations between RSUs and vehicles. The video streaming infrastructure can also leverage an edge caching scheme, like the ones proposed in [26, 4], in order to bring the data closer to the subscribers and to reduce frame transmission delays. With such a solution, we can ignore the delay between the video server and the RSUs.

3.6. Performance metrics

Our first metric of interest is the attained throughput for vehicle i, namely B_i. It reflects the rate at which the vehicle downloads video frames from the server. Additionally, we consider a second performance metric that pertains to the interruption time. The latter metric refers to the percentage of time the vehicle is left without being able to watch the video due to a void buffer and a too low (or null) download rate. We refer the interested reader to the work of Hosek et al. [27] for a thorough discussion about video stalling and Quality of Experience (QoE).

Table 2 summarizes the principal notations used in the paper.

Symbol	Description
v	Actual speed of the vehicle under study (m/s)
λ	Car density (veh/km)
μ	Bus and truck density (veh/km)
K	Number of RSUs
I	Distance between consecutive RSUs (m)
L	Signal range (m)
x	Penetration rate (in %)
c	Bit rate of the video codec for the vehicle under study (Mbps)
h_j	Length of the j-th zone of a given RSU (m)
Z	Number of zones for a given RSU
T_j	Transmission rate for a vehicle in the j-th zone (Mbps)
A_j	Achievable throughput for a vehicle in the j-th zone (Mbps)
B_i	Attained throughput for vehicle i (Mbps)
$Q(t)$	Queue length in the buffer of the vehicle under study at time t (s)
N	Number of subscriber vehicles within an RSU signal range
q_j	Proportion of subscriber vehicles that are in the j-th zone
f	Vector of times where the vehicle under study enters or leaves an RSU

Table 2: Principal notation.

4. Analytical solution

Let us first consider the case of a single RSU. We assume that the density of cars (resp. buses and trucks) follows a Poisson point process with rate λ (resp. μ). The relevance of this assumption is discussed in Section 5.

Let N be the number of subscriber vehicles in (the union of) the Z zones of this RSU, i.e., $\cup_{j=1}^{Z} H_j$. Hence, N is a random variable whose probability distribution is given by:

$$\mathsf{P}(N = n) = \frac{[x(\lambda+\mu)\sum_{j=1}^{Z} h_j]^n \exp(-x(\lambda+\mu)\sum_{j=1}^{Z} h_j)}{n!} \quad \text{for } n \geq 0 \qquad (1)$$

For any value of N, the proportion of subscribers located in the j-th zone of the RSU, denoted by q_j, can be calculated as: $q_j = \frac{h_j}{\sum_{k=1}^{Z} h_k}$.

On the other hand, when the number and location of the vehicle subscribers within the RSU signal range are known, simple approximations such as [28, 29, 30, 31] become applicable to estimate the attained throughput of any vehicle i. Let us use n_j to denote the number of vehicles in the j-th zone of a given RSU. Using [28], we can calculate the attained throughput by node i as follows:

$$B_i = \frac{1}{\sum_{j=1}^{Z} \frac{n_j}{A_j}} \tag{2}$$

Note that Eq. 2 returns identical attained throughputs for all nodes, regardless of their own achievable throughput. Therefore, we drop subscript i from B_i which becomes simply known as B. This property is in agreement with the aforementioned performance anomaly of IEEE 802.11 that ensures fairness among the nodes in terms of their number of access to the radio channel rather than their duration.

We propose to extend Eq. 2 to account for the uncertainty with regard to the number and location of the vehicle subscribers.

Using the probability distribution on N, we can approximate the expected value of B for a given vehicle that is connected to an RSU as:

$$\mathsf{E}(B) = \frac{1}{\sum_{n=0}^{+\infty} \mathsf{P}(N = n) \left[\sum_{j=1}^{Z} n.q_j.\frac{1}{A_j} \right]} \tag{3}$$

Note that to avoid singularities in Eq. 3 when the value of density approaches 0, we cap the value of $\mathsf{E}(B)$ to the one found for a single vehicle within an RSU. Besides, when a vehicle is between two RSUs, out of their radio ranges, the expected download rate equals to 0. Note also that Eq. 3 is approximate in the sense that it should include Z Poisson random variables and thus Z infinite sums. The interested reader can find the exact computation of $\mathsf{E}(B)$ in Appendix. In practice, both computations lead to very similar numerical results.

We are now able to calculate the expected amount of data downloaded for any vehicle given its location on the road. Assuming a given vehicle moving at an average speed v has passed K (non-overlapping) RSUs having each a signal range of L, its expected amount of downloaded data is simply obtained as:

$$\sum_{j=1}^{K} \int_0^{\frac{2L}{v}} \mathsf{E}(B)\,\mathrm{d}t = K\mathsf{E}(B).\frac{2L}{v}. \tag{4}$$

Note that if the road is winding, then the lengths of transmission zones h_j vary for each RSU as well as their q_j and $\mathsf{E}(B)$. In this case, the latter quantities differ for each RSU and must be computed separately, thereby invalidating the right-hand side of Eq. 4.

Let us denote by c the bit rate of the codec in use for the vehicle under study. Whenever the vehicle is within an RSU signal range RSU signal range,

it downloads data at speed B while playing (consuming) data at speed c. If $B \geq c$, the video can be played and the vehicle buffer fills up at the rate $B - c$. Otherwise, the vehicle must either resort to the video frames stored in its buffer (if any), or simply interrupt the video playback if the buffer is void.

Let f be a vector whose elements correspond to the times where the studied vehicle either enters or leaves any of the RSUs along the road. We use $f(k)$ to refer to its k-th element. In other words, f contains the arrival and departure times in each RSU for the vehicle. For instance, $f(1)$ (resp. $f(2)$) represents the time upon which the vehicle enters (resp. leaves) the first RSU while $f(3)$ is the time upon which the vehicle enters the second RSU. The amount of data stored in the buffer of the studied vehicle at any time can be obtained iteratively. We use $Q(t)$ to denote the length of the queue (in seconds) in the buffer of the vehicle at time t. Starting with a void buffer at time $f(0)$, its value at time $f(k)$ is easily derived from its value at time $f(k-1)$ as follows. For time $t \in [f(2k-1), f(2k)]$ with $k \geq 1$ (i.e., the vehicle is within an RSU signal range) we have:

$$Q(t) = [Q(f(2k-1)) + (\mathsf{E}(B) - c).(t - f(2k-1))]^+ \tag{5}$$

For time $t \in [f(2k), f(2k+1)]$ with $k \geq 1$ (i.e., the vehicle is out of range of any RSU) we have:

$$Q(t) = [Q(f(2k)) - c.(t - f(2k))]^+ \tag{6}$$

Hence, the total interruption time for the vehicle passing along K RSUs is given by:

$$\sum_{k=1}^{K} \left[\int_{f(2k-1)}^{f(2k)} \mathbb{1}_{Q(t)=0} \, \mathrm{d}t + \int_{f(2k)}^{f(2k+1)]} \mathbb{1}_{Q(t)=0} \, \mathrm{d}t \right] = \sum_{k=1}^{2K} \int_{f(k)}^{f(k+1)} \mathbb{1}_{Q(t)=0} \, \mathrm{d}t \tag{7}$$

The percentage of interruption time in the video playback for the vehicle can then be easily computed.

5. Vehicular mobility and network simulations

5.1. SUMO simulator for vehicle trajectories

SUMO (Simulation of Urban MObility) is an open source vehicular mobility simulator [21] based on micro-simulations. It consists in emulating the acceleration, braking and lane changing of each vehicle, individually and at regular interval. All these decisions are based on different models that depend on the type of road, the target speed of each vehicle, and its regime, which itself depends on the current traffic density (free-flow, car-following regimes, etc.). Let us recall that the target speed is the speed at which the vehicle drives when it is not slowed down by other vehicles. SUMO offers the possibility to simulate

different types of vehicles and roads. In our work, we consider two scenarios: a three-lane highway and a bidirectional expressway. We now describe how SUMO helps us generate realistic vehicle trajectories under various traffic configurations.

Highway and expressway scenarios

The highway scenario is composed of three lanes in the same direction. The expressway scenario has only two lanes, one for each direction. For these two scenarios, a 90 km section is simulated with two classes of vehicles: 90% of cars and 10% of trucks[1]. For each vehicle, the target speed is drawn from a normal distribution whose mean depends on the type of vehicles and roads. All the parameters are given in Table 3. In order to avoid edge effects due to the arrivals and departures of vehicles in the simulated section, only the 30 km at the middle of the expressway/highway are considered for further use in the ns-3 simulations (see Section 6). We run SUMO for different traffic densities ranging from very sparse to very dense. To reflect the fact that not all vehicles subscribe to the VoD service, we consider several possible values for the penetration rate of the VoD service, x.

	Expressway	Highway
Target speed for cars	80 km/h (22.2 m/s)	130 km/h (36.1 m/s)
Target speed for trucks	80 km/h (22.2 m/s)	90 km/h (25.0 m/s)
Proportion of cars	90%	90%
Proportion of trucks	10%	10%
Car length	5 m	5 m
Truck length	16 m	16 m
Number of lanes	2 (one per direction)	3 (same direction)

Table 3: SUMO parameters.

Parameter	Value
Star Wars video properties	Codec H265/HEVC 8.0 - Mean rate: 750 kbps
Harry Potter video properties	Codec H265/HEVC 10.1 - Mean rate: 460 kbps
ns-3 Wi-Fi Manager	Ideal Wi-Fi manager
TCP congeston control	New Reno

Table 4: Video and ns-3 parameters.

Vehicular traffic properties

We study the impact of the density of vehicles on the stream of traffic through the results of several SUMO simulations. For the sake of conciseness, we present

[1] We assume that buses have the same features as trucks.

only the results for the highway scenario as those for the expressway scenario are qualitatively the same.

In Figure 3a, we plot the distribution of the speed for different traffic density (number of vehicles per km simulated by SUMO). For a very low traffic density of 12 veh/km (purple curve), most of the vehicles drive at their target speed (36.1 m/s for the cars and 25 m/s for the trucks). At a traffic density of 54 veh/km (green curve), vehicles need to slow down and the cars target speed cannot be reached. Trucks (and cars behind the trucks) keep an average speed of 25 m/s. At 165 veh/km (blue curve), the speeds further decrease. They are distributed between 0 and 26 m/s. Congestions appear on certain sections of the highway. For 267 veh/km/lane (yellow curve), the traffic is totally jammed and most of the vehicles have speeds less than 10 m/s.

Figure 3b depicts the stream of traffic (expressed as the number of passing vehicles per hour) as a function of the traffic density. Not surprisingly, the stream of vehicles first tends to increase with growing values of traffic density as the total number of vehicles on the highway grows (and no congestion occurs). When the traffic becomes denser, the vehicles need to slow down and the stream of passing vehicles attains its maximum. Afterwards, the stream of passing vehicles begins to decrease as the speed of vehicles significantly drops. The shape of this curve is indeed very classical in transportation theory [32]. Although its precise characteristics depend on a lot of parameters (type of vehicles, highway ramps, speed limit, number of lanes, etc.), its general form remains similar. From a practical point of view, these simulations confirms the ability of SUMO to simulate realistic traffic that will be used in our ns-3 simulation to simulate scenarios with different properties with regard to the traffic density and vehicle speed.

5.2. ns-3 simulator for network communication

The network simulations are performed with the network simulator ns-3 version 3-29. The different parameters used in these simulations are given in Table 4. We use the IEEE 802.11p technology already implemented in ns-3. The chosen physical layer leads to the radio range and transmission rate described in the Table 1. The number of nodes, locations, speeds, and their spatial trajectories are imported from the SUMO simulator. We implemented, on top of the ns-3 IP stack, two video applications, one on the server side and one for the clients (vehicles). A vehicle that wishes to download a video during its journey opens a TCP connection to the central server. This video server is connected to the RSU through a 10 Gbps wired network. When the TCP connection is established, the video application at the server sends the video at its maximum speed until the TCP buffer is full. The effective throughput between the server and its clients is then limited by the TCP congestion control mechanism. Note that the time separating two successive frames is constant (30 frames per second) but frames have various sizes. On the vehicle side, the video application buffers the video frames and plays them at the video speed. The vehicle is thus able to store chunks of the video in advance when it is in the radio range of an RSU, and can play the video when it is disconnected from the

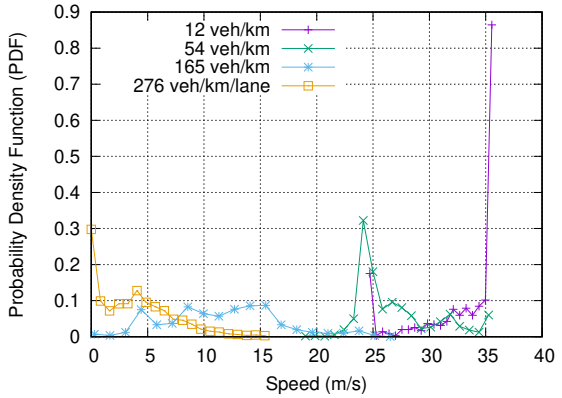

(a) Distribution of the speed attained by vehicles for four levels of traffic density.

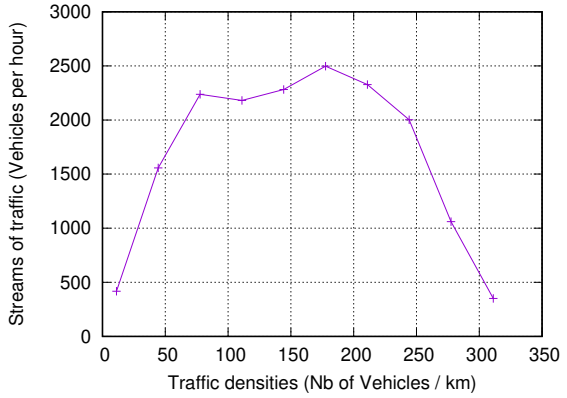

(b) Attained stream of traffic as a function of the traffic density.

Figure 3: Characteristics of a vehicular traffic on a highway as simulated by SUMO.

network. If the disconnection is too long, the buffer may become empty and the video stops until the vehicle reaches the next RSU. The video sent by the server and played by the vehicles correspond to real traces [33]. As the video properties (frame sizes, codec rates) may be very different from one movie to another, we also consider constant video where all the frames have the same size. Note that this latter case corresponds to having a CBR (Constant Bit Rate) flow. For each simulation, we collect the measurements for a unique vehicle that is randomly chosen among all the vehicles running the application. For this vehicle, the player measures the amount of bytes downloaded during the simulation and the disconnection time defined as the amount of time where the player was stopped due to a lack of video frames. To account for the intrinsic variability of the system, all simulation results correspond to the average of 100 ns-3 simulations.

In practice, we generated 10 different SUMO trajectories, and for each of them, we randomly selected 10 subsets of connected vehicles (subscribers) according to the penetration rate. This gives us a total of 100 simulation runs for which we randomly choose a given connected vehicle to compute the sought performance metrics.

6. Model validation

It is well established that when the traffic is sparse (free-flow regime), the time between two successive vehicles can be represented by an i.i.d. exponential random variable [32, 34]. If one assumes that the differences of speed among vehicles do not introduce a significant bias, then the distribution of the vehicle locations can be approximated by a Poisson point process. On the other hand, when the traffic becomes dense (car-following regime), strong correlations appear between vehicle locations and they tend to impede the use of a Poisson point process to represent the distribution of the vehicle locations. However, in our case, we are not considering the locations for all the vehicles but only for the subset of subscriber vehicles. Therefore, unless the penetration rate of the VoD service is close to 1, the density of subscriber vehicle is expected to remain relatively low (even if the traffic is dense) making the use of a Poisson point process representation a viable option. The realistic traffic and network simulations presented in this section aims at determining if the representation of the subscriber vehicles by a Poisson point process may offer an accurate approximation of the VoD system performance under different vehicles traffic conditions.

6.1. Example 1 - Robustness to the density of subscribed vehicles

In our first example, we study the download rate of subscribed vehicles and their percentage of interruption time in the video playback for different density values of subscribed vehicles. Using SUMO and ns-3, we compute the attained download rate and interruption time for the highway scenario (1x3 lanes), or alternately, the expressway scenario (2x1 lane) using the real video trace of Star Wars, or alternately, a CBR flow of equal rate. Note that this leads to a total of 4 combinations for the simulations. The distance between RSUs is set to 3000 meters and the video codec delivers on average a bit rate of 0.75 Mbps. The vehicles density is 35 veh/km for both highway and expressway scenarios. We vary the penetration rate in order to obtain between 0.5 and 5 subscribers per kilometer. Figure 4 shows the corresponding results. Although the specific values returned by the simulators naturally differ with regards to the scenario and video specifically chosen, they tend to exhibit a similar pattern where the download rate (resp. interruption time) is a decreasing (resp. increasing) function of the subscriber density. We observe that our analytical solution, while significantly abstracting from the underlying system, is able to capture the average behavior for each of the two studied performance parameters.

Before moving to our second example, we consider a different value for the video codec while all the other parameters are left unchanged. More precisely,

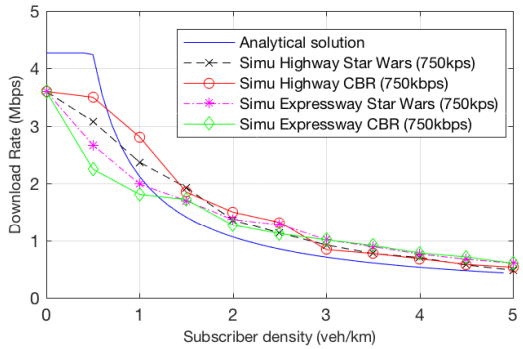

(a) Download rate attained by subscribed vehicles.

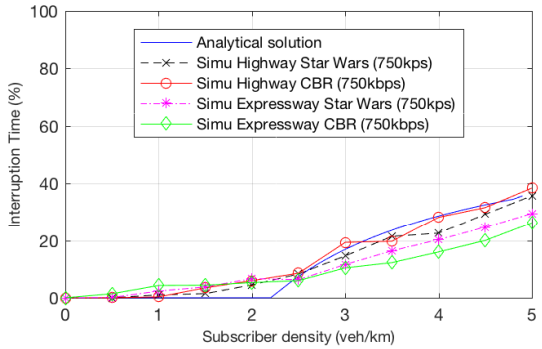

(b) Interruption time experienced by subscribed vehicles.

Figure 4: Example 1 - Robustness of our solution to changes of subscribed vehicle density for a video codec of 0.75 Mbps.

we replace the Star Wars video and its 0.75 Mbps codec by a Harry Potter video encoded at 0.46 Mbps. Figure 5 presents the associated results. As expected, varying the codec barely affects the download rate. On the other hand, comparing Figures 4b and 5b, we observe that downgrading the codec rate significantly postpones the occurrence of interruptions in the video playback to larger levels of vehicle density. While each combination of road scenario and video leads to different values for the download rate and the interruption time, our anaytical solution manages to deliver a fair estimate for any of these combinations.

6.2. Example 2 - Robustness to the distances between RSUs

Our second example studies the robustness of our analytical model to changes of distance between RSUs. In practice, we consider a large set of values for the distance between RSUs ranging from 1500 meters up to 8000 meters. Note that for 1500 meters, the portion of road out-of-range from any RSU is very small since 802.11p-based communications can extend up to 690 meters (i.e. $\sum_{j=1}^{6} h_j$

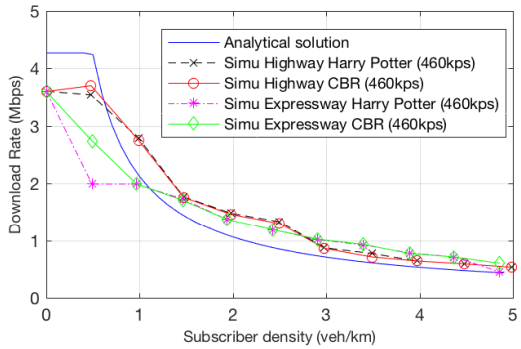

(a) Download rate attained by subscribed vehicles.

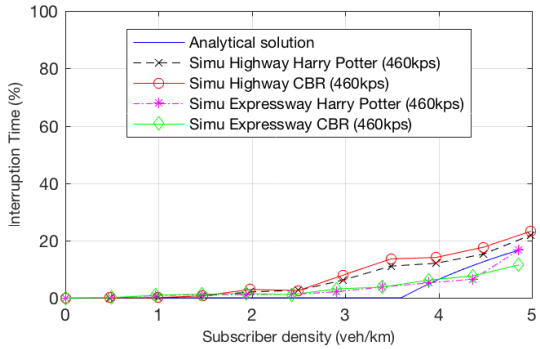

(b) Interruption time experienced by subscribed vehicles.

Figure 5: Example 1 - Robustness of our solution to changes of subscribed vehicle density for a video codec of 0.46 Mbps.

in Table 1). We set the vehicle density to 35 veh/km, the penetration rate to 14.3% leading to 5 subscribers per kilometer in average, and we choose a codec of 0.75 Mbps. We run the SUMO and ns-3 simulators for these values along with our model, and we provide the corresponding results for the interruption time in Figure 6. Not surprisingly, the interruption time grows with increasing distances between RSUs. Although the values returned by the simulator may slightly vary with the specific road scenario and video trace under consideration, our analytical solution delivers a good estimate of the average interruption time experienced by the subscriber vehicles regardless of the specific scenario.

6.3. Example 3 - Robustness to the video codec

In our third example, we study the accuracy of our analytical solution for a large range of codec bit rates. We let the codec bit rate vary from 0.4 to 2.5 Mbps in steps of 0.1 Mbps. Note that because we enforce specific values for the codec bit rate, our simulation experiments include only CBR codec (no real video traces). The distance between RSUs is set to 3000 meters, the subscribed

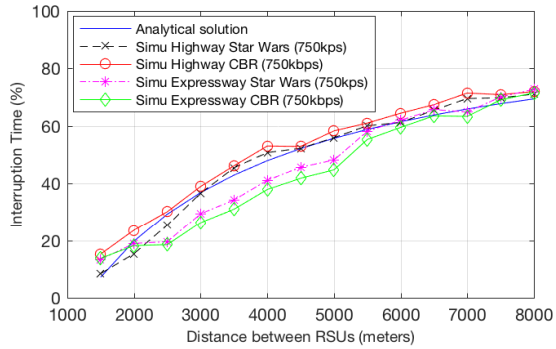

Figure 6: Example 2 - Robustness of our solution to the distance between RSUs.

vehicle density is set to 1.5 veh/km for a total traffic density of 35 veh/km. Figure 7 illustrates the interruption times delivered by the simulators as well as those forecast by our analytical solution. Despite a slight discrepancy in predicting the precise level at which a codec bit rate will cause interruption times, the results returned by the combined simulators SUMO and ns-3 are in good agreement with those provided by the analytical solution.

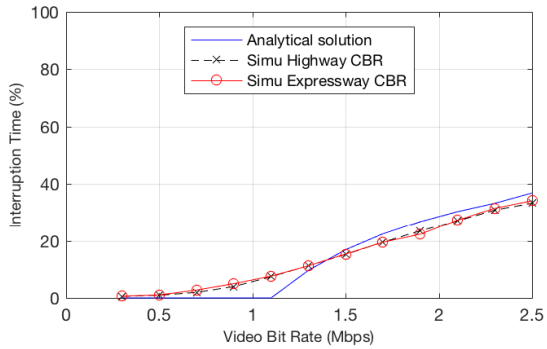

Figure 7: Example 3 - Robustness of our solution to changes of video codec.

6.4. Example 4 - Robustness to the density of all vehicles (subscribed and unsubscribed)

Our last example investigates an interesting property related to our model assumption of not taking into account the whole vehicular traffic density but rather only the subscriber density. If one varies the market penetration rate for the VoD service, namely x, then multiple settings can lead to the same level of density for subscribed vehicles. Indeed, as long as the product $x(\lambda+\mu)$ remains equal (recall that λ and μ denote the car and truck density respectively), the subscriber density holds equal. Nonetheless, in the simulator SUMO, having

different values for λ and μ leads to very different mobility trajectories (with potentially significantly different levels of vehicular traffic congestion). On the other hand, because of its relative simplicity, our analytical solution does not account for these distinctions and does not handle x and $(\lambda + \mu)$ separately but instead their joint product.

To test this assumption, we run our simulators for a wide range of possible values of $(\lambda + \mu)$ going from 10 veh/km up to 180 veh/km. In the same time, to keep the density of subscribed vehicles constant equals to 5 veh/km, the penetration rate varies from 50% to 2.8%. Figure 8 represents the results delivered by the simulators as well as those returned by our analytical solution. The distance between RSUs is set to 3000 meters and the video in use is Star Wars encoded at a bit rate of 0.75 Mbps.

We observe that although the interruption times are not totally insensitive to the specific values for the vehicle density $(\lambda + \mu)$, they are by far mostly determined by the product $x(\lambda + \mu)$ representing the subscriber density. This result is fundamental to our approach as it confirms that, at least for the sake of our purpose, one can solely consider the subscriber vehicle density without taking into account the overall vehicle density (including unsubscribed vehicles).

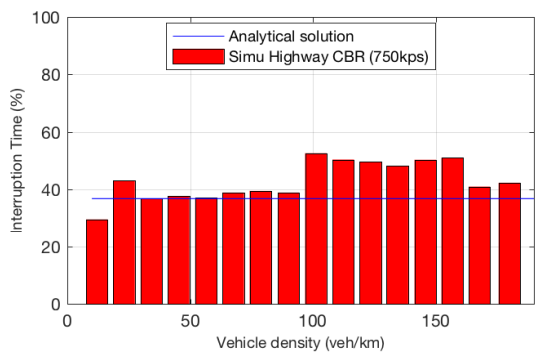

Figure 8: Example 4 - Robustness of our solution to changes in the whole traffic densities (subscribed and unsubscribed vehicles).

7. Capacity planning and blocking policy

We explore a number of possible applications illustrating how our proposed modeling approach can help capacity planning or design efficient blocking policies.

7.1. Application 1 - Joint effect of subscriber vehicle density and distance between RSUs on interruption times

Our first application describes a capacity planning problem in which the VoD provider can decide the distance between RSUs before their deployment. Assuming a video codec with a bit rate of 0.75 Mbps, our analytical solution

can easily and quickly calculate the estimated interruption times in the video playback of subscribed vehicles for a large set of values for the distance between RSUs and for the subscriber density. Figure 9 shows the corresponding results in a contour plot. Note that although the computation of this figure involved thousands of executions of our analytical solution, it was obtained in only a few seconds on a standard machine[2]. This figure pinpoints the combinations of values for the distance between RSUs and for the vehicle density that will not cause interruptions in the video playback. For example, Figure 9 indicates that broadcasted videos will not lag if the RSUs are located at 6000 meters distance from each other as long as that the density of subscribers is below 1 veh/km.

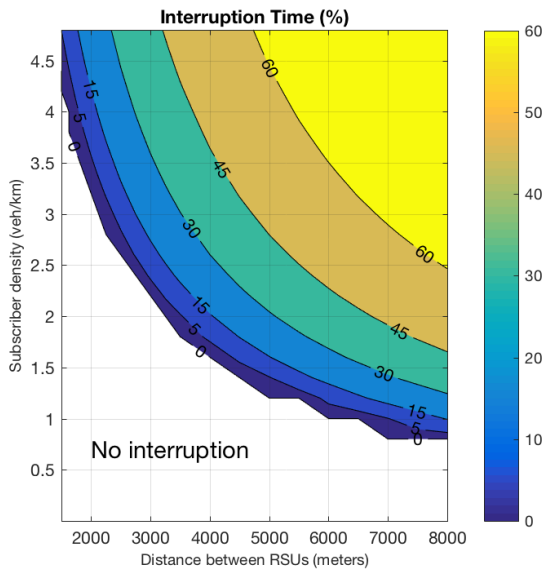

Figure 9: Application 1 - Discovering values of subscriber vehicle density and distance between RSUs causing no interruptions in the playback of a video with a bit rate of 0.75 Mbps.

7.2. Application 2 - Joint effect of video bit rate and distance between RSUs on interruption times

In our second application, we assume that the VoD provider can control both the bit rate of its video codec and the distance between RSUs. On the other hand, the density of subscriber vehicles is set to a constant value of 1.5 veh/km. Using our analytical solution, we compute the estimated interruption times for a wide range of values of the video bit rate and of the distance between RSUs. Figure 10 depicts the obtained results. It indicates, for example, that if RSUs

[2] 2 GHz Intel Core i5 with 2 core processor.

are spaced by approximately 8000 meters, then the video codec should deliver at most a bit rate of 0.4 Mbps.

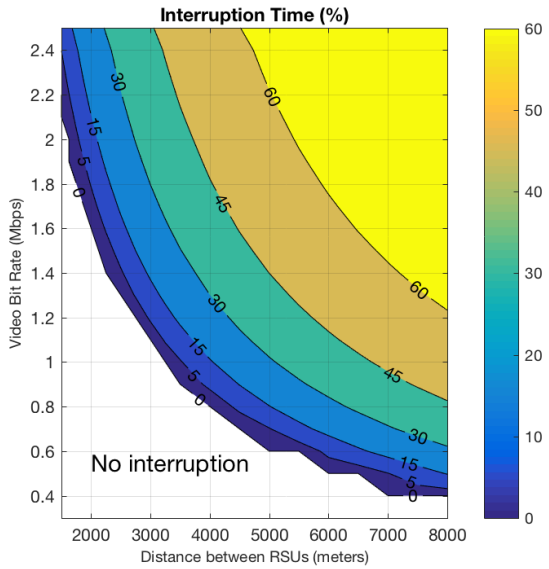

Figure 10: Application 2 - Discovering values of video bit rate and distance between RSUs causing no interruptions in the playback of a video for a subscriber vehicle density of 1.5 veh/km.

To further illustrate how the model can be used to help with the design a VoD streaming system, we resort to a simple example. Let us consider a highway that can carry up to about $1,500$ vehicles per hour at the peak of the day. Based on Figure 3b, we expect the traffic densities to attain at most 50 veh/km. Assuming that the VoD provider forecasts a maximum penetration rate of 3 %, then the density of connected vehicles (subscribers) amounts to approximately 1.5 vehicles per kilometer. If the targeted video quality requires a codec rate of 750 kbps, then our model suggests the provider to choose a distance between RSUs of 4000 meters. This corresponds to the point $(4000, 1.5)$ in Figure 9 that guarantees no disruption time. Note that if this distance between RSUs is not acceptable for the provider (e.g., too costly), Figure 10 shows that a distance between RSU of 5000 meters can also be set (leading to 25% less RSUs) if the quality of the codec is decreased to 600 kbps. This corresponds to the point $(5000, 0.6)$ in Figure 10.

7.3. Application 3 - Blocking or not blocking the lowest transmission rates

In Section 3.4, we discussed the performance anomaly of IEEE 802.11 (e.g., [25]) where vehicles with high transmission rates are penalized by those with low ones since any vehicle associated to the same RSU communicates, more or less, at the same rate. While this phenomenon is typically regarded negatively, it can

be proved useful in our case owing to the regular mobility of our nodes in which any of them will gradually experience all the possible range of transmission rates. Therefore, we investigate if temporarily blocking the transmission of vehicles currently experiencing a low transmission rate may lead to a collective gain for all vehicles.

Using thousands of executions of our analytical solution, we discover the evolution of the interruption time for two different blocking policies considering four possible values for the video bit rate, as well as a large set of values for the subscriber vehicle density and the distance between RSUs. The first policy simply consists of allowing all transmission rates (referred to as the no-blocking policy). The corresponding results are presented in Figure 11a. This figure simply extends the results obtained previously in Figure 9 to three additional video bit rates. Note that, for the sake of clarity, we only represent the frontier lines corresponding to the contour of the onset of interruption times (i.e., 0% of interruption time). As expected, Figure 11a shows that lessening the video bit rate from 2.5 to 0.4 Mbps significantly postpones the onset of interruption times.

Figure 11b illustrates the estimated interruption times if we block data exchange for vehicles currently having a transmission rate less than 9 Mbps. First, we observe that downgrading the video bit rate to lower values enables to further postpone the frontier lines delimiting the onset of interruptions. More interestingly, comparing Figures 11a and 11b shows that implementing a blocking policy can bring significant collective gain, especially when the density of subscriber vehicles is high, or alternately when the distance between RSUs is small. For instance, with a codec of 0.75 Mbps and a distance between RSUs of 3000 meters, the VoD network can accommodate up to a subscriber vehicle density of 2 veh/km under the no-blocking policy as compared to a density of 3 veh/km under the blocking policy. Conversely, in the case where both the density of subscriber vehicles is small and the distance between RSUs is large, the blocking policy may underperform a no-blocking policy. Indeed, under such circumstances, chances are high that each subscriber vehicle is the only node associated with its RSU so that blocking low transmission rates will likely translate into blocking the only vehicle that would otherwise download data. However, with this exception, Figure 11 indicates that the VoD network can largely benefit, without much effort, from implementing an access control policy based on the vehicle transmission rates.

7.4. Application 4 - Determining an adequate blocking policy

Our last application discusses the choice of a cut-off value for the blocking policy. Let us recall that the blocking policy consists in preventing vehicles with the lowest negotiated transmission rates from transmitting. Given the transmission rates of IEEE 802.11p (see Table 1), there are 6 possible ways of parameterizing a blocking policy. Assuming a bit rate for the video codec of 0.75 Mbps, we investigate with our analytical solution the performance associated to each possible cut-off value for a large set of values of distance between

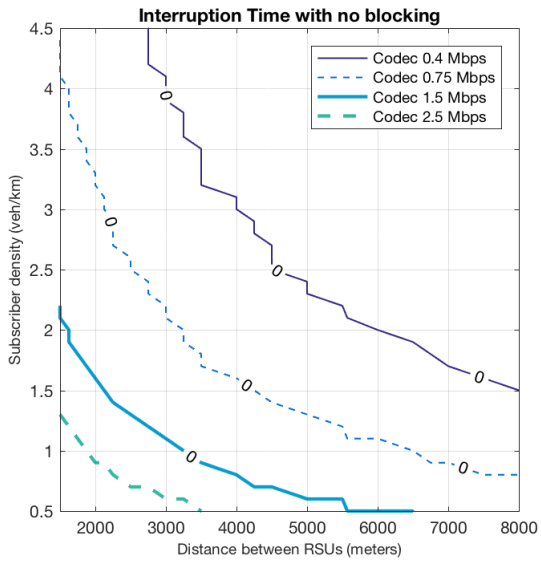

(a) Allowing all transmission rates.

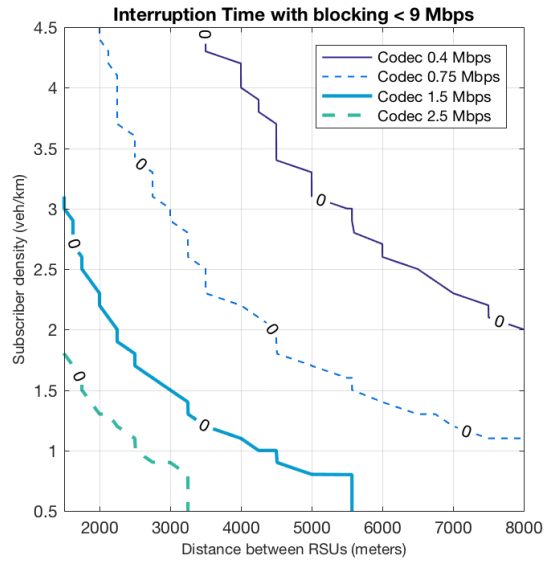

(b) Blocking transmission rates less than 9 Mbps.

Figure 11: Application 3 - Estimated onset of interruption times for various video bit rates, distances between RSUs, and subscriber vehicle densities under the no-blocking and the blocking low transmission rate policies.

RSUs and densities of subscriber vehicles. Figure 12 illustrates the corresponding results. We observe that blocking all transmission rates under 24 Mbps leads to the best performance except when the distance between RSUs exceeds 6700 meters (as discussed formerly because, under these circumstances, subscriber vehicles are likely to be alone on each RSU making any form of blocking virtually useless). However, the policy that consists of blocking all transmission rates except 18 and 24 Mbps appears like an excellent trade-off regardless of the actual distance between RSUs. Hence, this would make it our natural candidate for a real implementation of VoD service broadcasting video at 0.75 Mbps bit rate.

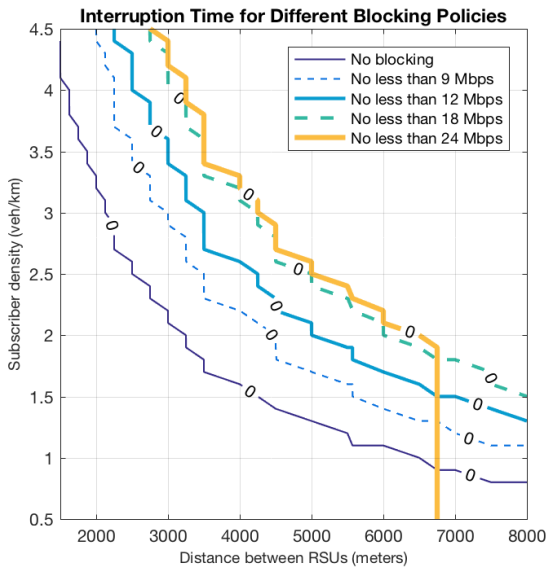

Figure 12: Application 4 - Comparing the onset of interruption times for six different blocking policies given a video codec with bit rate of 0.75 Mbps.

8. Conclusions

We have presented a simple and versatile solution to the performance computation of an IEEE 802.11-based VoD service for connected vehicles traveling along a highway or an expressway. We use a probabilistic approach to account for the intrinsic dynamics (mobility) of vehicles, the density of subscriber vehicles, the distance between RSUs, as well as the video codec bit rate. Our solution delivers estimates for the attained throughput of each subscriber vehicle as well as for their interruption times in the playback of their video.

We studied several scenarios to assess the accuracy of our approximate solution comparing its results with those of a discrete-event simulator with realistic

vehicular traffic trajectories, video traffic, and networking protocols. Overall, our solution is found to be accurate delivering estimates in good agreement with the simulation results. We illustrated the potential of our model for determining an adequate setting of the VoD service parameters using four application examples and derived some qualitative conclusions. For example, we found that by temporarily blocking the transmission of vehicles when they have the lowest transmission rates, all vehicles can benefit from increased performance.

A natural follow-up of this work could be to adapt our modeling framework to the case of urban scenarios wherein vehicle mobility is much harder to predict and the traffic is typically denser. This will be the subject of future work.

9. Acknowledgments

The authors wish to thank the anonymous referees for their thorough and constructive review of an earlier version of this paper.

References

[1] H. Khan, S. Samarakoon, M. Bennis, Enhancing video streaming in vehicular networks via resource slicing, IEEE Transactions on Vehicular Technology.

[2] L. Sun, H. Shan, A. Huang, L. Cai, H. He, Channel allocation for adaptive video streaming in vehicular networks, IEEE Transactions on Vehicular Technology 66 (1) (2017) 734–747. doi:10.1109/TVT.2016.2535659.

[3] D. Bethanabhotla, G. Caire, M. J. Neely, Adaptive video streaming for wireless networks with multiple users and helpers, IEEE Transactions on Communications 63 (1) (2015) 268–285. doi:10.1109/TCOMM.2014.2378774.

[4] Y. Guo, Q. Yang, F. R. Yu, V. C. Leung, Cache-enabled adaptive video streaming over vehicular networks: A dynamic approach, IEEE Transactions on Vehicular Technology 67 (6) (2018) 5445–5459.

[5] J. Qiao, Y. He, X. S. Shen, Improving video streaming quality in 5G enabled vehicular networks, IEEE Wireless Communications 25 (2) (2018) 133–139.

[6] N. Zribi, B. Alaya, T. Moulahi, Video streaming in vehicular Ad Hoc networks: Applications, challenges and techniques, in: 2019 15th International Wireless Communications & Mobile Computing Conference (IWCMC), IEEE, 2019, pp. 1221–1226.

[7] P. Piñol, A. Torres, O. López, M. Martinez, M. P. Malumbres, Evaluating HEVC video delivery in VANET scenarios, in: 2013 IFIP wireless days (WD), IEEE, 2013, pp. 1–6.

[8] A. Vinel, E. Belyaev, B. Bellalta, H. Hu, Live video streaming in vehicular networks, in: International Workshop on Communication Technologies for Vehicles, Springer, 2014, pp. 156–162.

[9] C. Rezende, A. Mammeri, A. Boukerche, A. A. Loureiro, A receiver-based video dissemination solution for vehicular networks with content transmissions decoupled from relay node selection, Ad Hoc Networks 17 (2014) 1–17.

[10] C. Quadros, E. Cerqueira, A. Santos, J. Lim, M. Gerla, Beacon-less video streaming management for VANETs based on QoE and link-quality, in: 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM), IEEE, 2015, pp. 191–198.

[11] C. Rezende, A. Boukerche, H. S. Ramos, A. A. Loureiro, A reactive and scalable unicast solution for video streaming over VANETs, IEEE Transactions on Computers 64 (3) (2014) 614–626.

[12] X. Zhu, J. Bi, M. Li, H. Liu, Adaptive video streaming uploading with moving prediction in VANETs scenarios, in: 2015 IEEE International Conference on Networking, Architecture and Storage (NAS), IEEE, 2015, pp. 39–44.

[13] R. Wang, M. Almulla, C. Rezende, A. Boukerche, Video streaming over vehicular networks by a multiple path solution with error correction, in: 2014 IEEE International Conference on Communications (ICC), IEEE, 2014, pp. 580–585.

[14] A. Bazzi, B. M. Masini, A. Zanella, G. Pasolini, Vehicle-to-vehicle and vehicle-to-roadside multi-hop communications for vehicular sensor networks: Simulations and field trial, in: 2013 IEEE International Conference on Communications Workshops (ICC), IEEE, 2013, pp. 515–520.

[15] P. Luoto, M. Bennis, P. Pirinen, S. Samarakoon, K. Horneman, M. Latva-aho, System level performance evaluation of LTE-V2X network, in: European Wireless 2016; 22th European Wireless Conference, VDE, 2016, pp. 1–5.

[16] P. Luoto, M. Bennis, P. Pirinen, S. Samarakoon, K. Horneman, M. Latva-Aho, Vehicle clustering for improving enhanced LTE-V2X network performance, in: 2017 European Conference on Networks and Communications (EuCNC), IEEE, 2017, pp. 1–5.

[17] H. Soleimani, T. Begin, A. Boukerche, Safety message generation rate adaptation in LTE-based vehicular networks, Computer Networks 128 (2017) 186–196.

[18] W. Xu, W. Shi, F. Lyu, H. Zhou, N. Cheng, X. Shen, Throughput analysis of vehicular internet access via roadside WiFi hotspot, IEEE Transactions on Vehicular Technology 68 (4) (2019) 3980–3991.

[19] T. Begin, A. Busson, I. Guérin Lassous, A. Boukerche, Performance analysis of video on demand in an IEEE 802.11 p-based vehicular network, Computer Communications 146 (2019) 174–185.

[20] S. Reñé, C. Ganán, J. Caubet, J. Alins, J. Mata, J. L. Munoz, Analysis of video streaming performance in vehicular networks, in: The First International Conference on Advanced Communications and Computation, INFOCOMP, Vol. 10, 2011, p. 2011.

[21] D. Krajzewicz, J. Erdmann, M. Behrisch, L. Bieker, Recent development and applications of SUMO - Simulation of Urban MObility, International Journal On Advances in Systems and Measurements 5 (3&4) (2012) 128–138.
URL `http://elib.dlr.de/80483/`

[22] A. Duda, Understanding the performance of 802.11 networks, in: PIMRC, Proceedings IEEE, Vol. 8, 2008, pp. 2008–1.

[23] R. Grünblatt, I. Guérin Lassous, O. Simonin, Simulation and performance evaluation of the Intel rate adaptation algorithm, in: Proceedings of the 22nd International ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems, 2019, pp. 27–34.

[24] S. Derek, Minstrel, `https://sourceforge.net/p/madwifi/svn/HEAD/tree/madwifi/trunk/ath_rate/minstrel/minstrel.txt`.

[25] M. Heusse, F. Rousseau, G. Berger-Sabbatel, A. Duda, Performance anomaly of 802.11 b, in: INFOCOM, Proceedings IEEE, Vol. 2, IEEE, 2003, pp. 836–843.

[26] Z. Su, Y. Hui, Q. Xu, T. Yang, J. Liu, Y. Jia, An edge caching scheme to distribute content in vehicular networks, IEEE Transactions on Vehicular Technology 67 (6) (2018) 5346–5356.

[27] J. Hosek, D. Kovac, D. Uhlir, O. Galinina, S. Andreev, Y. Koucheryavy, M. Ries, On predicting video quality expectations of mobile users, in: 2015 7th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT), IEEE, 2015, pp. 110–115.

[28] M. Amer, A. Busson, I. Guérin Lassous, Association optimization in Wi-Fi networks: Use of an access-based fairness, in: MSWIM, Proceedings ACM, ACM, 2016, pp. 119–126.

[29] L. Li, M. Pal, Y. R. Yang, Proportional fairness in multi-rate wireless LANs, in: INFOCOM, Proceedings IEEE, IEEE, 2008, pp. 1004–1012.

[30] O. B. Karimi, J. Liu, J. Rexford, Optimal collaborative access point association in wireless networks, in: INFOCOM, Proceedings IEEE, IEEE, 2014, pp. 1141–1149.

[31] W. Li, S. Wang, Y. Cui, X. Cheng, R. Xin, M. A. Al-Rodhaan, A. Al-Dhelaan, AP association for proportional fairness in multirate WLANs, IEEE/ACM Transactions on Networking 22 (1) (2014) 191–202.

[32] D. Gazis, Traffic Theory, International Series in Operations Research & Management Science, Springer US, 2002.

[33] Video trace library, `http://trace.eas.asu.edu/`, arizona State University.

[34] S. Yousefi, E. Altman, R. El-Azouzi, M. Fathy, Connectivity in vehicular Ad Hoc networks in presence wireless mobile base-stations, in: 2007 7th International Conference on ITS Telecommunications, 2007, pp. 1–6. `doi:10.1109/ITST.2007.4295839`.

Appendix: Palm calculus

We consider two point processes R and V. $R = \{R_i\}_{i\in\mathbb{Z}}$ represents the RSU. The distance between two consecutive points of R is supposed to be constant equals to I. The closest point of R from the origin is assumed to be uniformly distributed in $[-\frac{I}{2}, \frac{I}{2}]$. The Poisson point process describing the vehicles is denoted V. Its intensity is $\lambda + \mu$ (cars and trucks). The number of vehicles in a given interval X is denoted $V(X)$. As defined in Section 3, each RSU has a communication range L with $L < I$ that can be decomposed in Z transmission zones H_j. The zones H_j with $1 \leq j \leq Z$ corresponds to a transmission rate T_j and an achievable throughput A_j. By convenience, we assume that these zones are defined for an RSU located at the origin. For an RSU at location R_i, the different zones H_j are shifted. Formally, we define the different zones of a point/RSU R_i as $H_j^i = H_j \oplus R_i$ where $\oplus$ is the Minkowski sum. The number of vehicles that can download at transmission rate T_j from the RSU located at R_i is then defined as $V_j^i = V(H_j^i)$. With these definitions, we can compute the mean instantaneous throughput of a typical vehicle. As in Section 4, our throughput estimation relies on the formula given in [28] that gives the mean throughput for a set of stations associated to a same AP. The typical vehicle is defined as the vehicle located at 0 under Palm expectation and with regard to the process V. For our model, it leads to:

$$\mathbb{E}_V^0[B] = \mathbb{E}_V^0\left[\frac{1}{\sum_{j=1}^{Z}\frac{V_j^\tau}{A_j}}\mathbb{1}_{0\in R\oplus[-L,L]}\right] \tag{8}$$

where B is the mean throughput, and τ is the index of the point of the process R that is the closest from the origin. Note that τ is a random variable that takes only two possible values: 0 or -1. $R \oplus [-L, L]$ is the union of the radio zones of all RSUs.

To compute this quantity, we condition on the different zones where the typical point (typical vehicle) can be located. We get:

$$\begin{aligned}
\mathbb{E}_V^0[B] &= \sum_{n=1}^{Z} \mathbb{E}_V^0 \left[\frac{1}{\sum_{j=1}^{Z} \frac{V_j^\tau}{A_j}} \mathbb{1}_{0 \in H_n^\tau} \right] && (9) \\
&= \sum_{n=1}^{Z} \mathbb{E}_V^0 \left[\frac{1}{\sum_{j=1}^{Z} \frac{V_j^\tau}{A_j}} \middle| 0 \in H_n^\tau \right] \mathbb{P}_V^0 \left(0 \in H_n^\tau \right) && (10) \\
&= \sum_{n=1}^{Z} \mathbb{E} \left[\frac{1}{\sum_{j=1}^{Z} \frac{V_j^\tau + \mathbb{1}_{j=n}}{A_j}} \right] \mathbb{P}_V^0 \left(0 \in H_n^\tau \right) && (11)
\end{aligned}$$

Eq. 11 is obtained by applying the Slyvniack's theorem. The random variable V_j^τ follows a discrete Poisson distribution with parameter $\lambda_j = (\lambda + \mu)\nu(H_j^\tau)$, where $\nu(.)$ is the Lebesgue measure in $\mathbb{R}$. $\nu(H_j^\tau)$ is then simply the size of the jth transmission zone.

The process R has the same distribution under Palm measure w.r.t. the process V as under the classical expectation. The distance from 0 to the first RSU on the positive axis is uniformly distributed in $[0, L]$. Finally, we get:

$$\begin{aligned}
\mathbb{E}_V^0[B] &= \sum_{n=1}^{Z} \sum_{i_1=0}^{+\infty} \cdots \sum_{i_Z=0}^{+\infty} \frac{1}{\sum_{j=1}^{Z} \frac{i_j + \mathbb{1}_{j=n}}{A_j}} && (12) \\
\times & \quad \frac{\lambda_1^{i_1}}{i_1!} \quad \cdots \frac{\lambda_Z^{i_Z}}{i_Z!} e^{-(\lambda_1 + \cdots + \lambda_Z)} \frac{\nu(H_n^0)}{I} && (13)
\end{aligned}$$