Video on Demand in IEEE 802.11p-based Vehicular Networks: Analysis and Dimensioning

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

► To cite this version:

Thomas Begin, Anthony Busson, Isabelle Guérin-Lassous, Azzedine Boukerche. Video on Demand in IEEE 802.11p-based Vehicular Networks: Analysis and Dimensioning. MSWiM'18 - 21st ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems, Oct 2018, Montreal, Canada. 10.1145/3242102.3242109 . hal-01858019

HAL Id: hal-01858019

<https://hal.science/hal-01858019>

Submitted on 18 Sep 2018

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.

Video on Demand in IEEE 802.11p-based Vehicular Networks: Analysis and Dimensioning

Thomas Begin
Univ Lyon, Université Claude
Bernard Lyon 1, ENS de Lyon,
Inria, CNRS, LIP, France

Anthony Busson
Univ Lyon, Université Claude
Bernard Lyon 1, ENS de Lyon,
Inria, CNRS, LIP, France

Isabelle Guérin Lassous
Univ Lyon, Université Claude
Bernard Lyon 1, ENS de Lyon,
Inria, CNRS, LIP, France

Azzedine Boukerche
DIVA Lab, University of Ottawa
Canada

ABSTRACT

We consider a VoD (Video on-Demand) platform designed for vehicles traveling on a highway or other major roadway. Typically, cars or buses would subscribe to this delivery service so that their passengers get access to a catalog of movies and series stored on a back-end server. Videos are delivered through IEEE 802.11p Road Side Units deployed along the highway. In this paper, we propose a simple analytical and yet accurate solution to estimate (at the speed of a click) two key performance parameters for a VoD platform: (i) the total amount of data downloaded by a vehicle over its journey and (ii) the total “interruption time”, which corresponds to the time a vehicle spends with the playback of its video interrupted because of an empty buffer. After validating its accuracy against a set of simulations run with ns-3, we show an example of application of our analytical solution for the sizing of an IEEE 802.11p-based VoD platform.

1. INTRODUCTION

With more than 1.25 billion vehicles, cars and buses represent the most popular mode/means of traveling in most countries and cities, be it for work or for leisure. These possibly lengthy and frequent journeys can grow tiresome for passengers. On the other hand, in less than a decade, Netflix has become the most-watched VoD (Video on-Demand) service in the US and in many other countries with 125 million subscribers worldwide surpassing the cable TV on the USA market. This success was made possible through the availability of a large network infrastructure connecting residential units and homes with high bandwidth links. Then, one can wonder if it would be possible to provide a VoD service for vehicles traveling on a highway and other major roadway using a network infrastructure.

Any architecture for VoD must deal with unicast streams (as opposed to broadcast traffic) unless the architecture focuses exclusively on live streaming where multiple subscribers may watch the same content at the same time. Said differently, in our case, each (subscriber) vehicle receives specific content based on its requests. This virtually precludes any architecture involving V2V (Vehicle-to-Vehicle) communications where some vehicles relay video frames for the others. However, several options are possible for connecting the vehicles and transmit them, over the radio channel, their requested content. In this paper, we opt for IEEE 802.11p but broadband cellular network technology such as 4G and 5G is a possible alternative.

In this paper, we study an application-oriented platform designed to stream videos to vehicles passing on a highway. More specifically, we study what could be its network infrastructure using IEEE 802.11p for the delivery of video frames to the passing vehicles. Our contributions are the following. First, we devise a conceptually simple and yet accurate analytical solution to estimate performance parameters of this platform. Second, we illustrate how the proposed analytical solution can help at dimensioning the networking resources of an IEEE 802.11p-based VoD platform. Finally, we present a means of improving the collective experience of vehicles by temporarily blocking access to the radio channel for some vehicles.

The remainder of the paper is organized as follows. Section 2 presents a brief overview of the related works. In Section 3 we detail our proposed analytical solution to estimate performance parameters of the considered VoD platform. Section 4 presents numerical results illustrating the accuracy of our solution. In Section 5 we show an example of application of our analytical solution for the sizing of an IEEE 802.11p-based VoD platform. Section 6 concludes this paper.

2. RELATED WORKS

The performance modeling of IEEE 802.11 has attracted much attention since the pioneering work by Bianchi [5]. Numerous extensions have been proposed to take into account unsaturated traffic (e.g., [17]) or imperfect radio channel conditions (e.g., [6]). More recently, several contributions

such as [8, 19] have specifically addressed the case of IEEE 802.11p amendment and they consider EDCA (Enhanced Distributed Channel Access) mechanisms where frames have different priorities depending on their criticality. Typically, these works rely on mono-dimensional or two-dimensional Markov chains whose solutions may not necessarily be trivial. Aside from these fine-grained models, more macroscopic models for IEEE 802.11 that deliver estimates for the mean throughput attained by a node (for a given level of priority) have been proposed [3, 12, 11, 13]. In this paper, we too consider a general and versatile model that eases its application to the case of a VoD platform.

Video dissemination over vehicular networks, which consists in distributing the same video to a set of vehicles, has drawn significant interests. The video in question may pertain to roadside video emergency (e.g., warning drivers of animals crossing the roadways), video surveillance or entertainment. For example, in [15], Quadros et al. have proposed to combine a geographical dissemination routing protocol with local decisions in order to forward video frame/sequence which does not depend only on distance but also on video QoE (Quality of Experience) parameters. In a separate work [18], Xing et al. consider optimal paths to disseminate videos among a set of vehicles. Given the roadmap of a vehicle and the location of vehicles and RSUs, the authors cast the problem as an optimization problem that aims to minimize the delivery delay of multimedia messages (e.g., video) and proposed a solution to compute the best dissemination strategy. These video dissemination protocols may be improved through network coding techniques (e.g., [10]) that enable to overcome packets losses at different vehicles thanks to a limited number of coded packets. Note that, unlike our work, video dissemination studies focus on a given multimedia message or video which has to be delivered to all considered vehicles. Furthermore, video dissemination solutions tend to often make use (partially or exclusively) of V2V (Vehicle-to-Vehicle) communications.

To our knowledge, few studies have explored the problem of delivering a Video on Demand service (where vehicles request different videos) over wireless vehicular networks. In [16, 4], the authors have tackled this issue 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 the allocation of time slots in order to optimize a utility function that expresses the video quality experienced at vehicles.

Overall, the work that comes closest to ours is that of Sun et al. in [16]. The authors consider a similar scenario to the ones considered in this paper where videos are distributed to a set of vehicles, and where RSUs do not entirely cover the highway or road under study. For this scenario, Sun et al. propose a joint resource allocation and video streaming algorithm based on a multiuser game. Note that their work differs from ours for two main reasons: (i) their approach assumes a TDMA scheme and does not rely on existing standardized MAC protocols (e.g., IEEE 802.11p), and (ii) they aim to optimize resources and the selection of video quality layers for each vehicle in real time whereas the goal of our

Figure 1: Considered VoD architecture for vehicles along a highway.

study is to provide an help for dimensioning the networking resources of an IEEE 802.11p-based VoD platform.

3. SYSTEM CONSIDERED AND ITS ANALYTICAL SOLUTION

3.1 System description

We describe the networking aspects of a VoD (Video on Demand) platform designed for vehicles traveling on a highway or other major roadway. Typically, cars or buses would subscribe to this delivery service so that their passengers get access to a catalogue of movies and series stored on a back-end server. Along the highway, RSUs (Road Side Units) are deployed and serve as APs (Access Points) to the passing vehicles using IEEE 802.11p [2]. Figure 1 illustrates the outline of the considered VoD platform. We now discuss each of the components involved in a possible architecture designed to provide the VoD service on a highway.

RSUs

RSUs are network equipments that are deployed along a road. Each RSU is linked through a high-speed wired link to a back-end video server that stores, possibly in an external database server, all available videos. Additionally, RSUs have another wireless interface that serves as an access point to passing vehicles. We denote by N the total number of RSUs along the studied section of highway and by L the length of the signal range of an RSU. We assume that RSUs are placed uniformly along the highway with distance I from each other. In practice, to limit the number of RSUs and so the cost of deployment, we expect I to be larger than L . Therefore, the coverage of consecutive RSUs will not overlap and portions of the highway will be left without network access.

Vehicles

Assuming that only vehicles that have subscribed to the VoD service interact with the RSUs, we can restrict our description to the subset of these subscribers. Each (subscriber) vehicle i is characterized by its velocity v_i . Vehicles are equipped with IEEE 802.11p interfaces to enable communication with RSUs. Provided a vehicle is within an RSU signal range, it attempts to download the content of its video as much as possible. Downloaded data corresponds to video frames. Whenever a vehicle downloads frames at a faster rate than it plays frames (on its media player), it stores the surplus in a buffer building up a queue with received but not yet played frames. We assume that there is no upper limit on the buffer size. As soon as a vehicle leaves an RSU

signal range and ends up in a portion without network access, it keeps playing frames drawn from the buffer until the queue exhausts. A vehicle will experience an interruption of the video playback unless it reaches the next RSU before its buffer becomes entirely depleted.

IEEE 802.11p

The communications between the RSUs and the passing vehicles operate on IEEE 802.11p [2]. Note that the described architecture involves no V2V (Vehicle-to-Vehicle) communications. This setting is motivated by the user-specific content of the communications since each vehicle picks the video of its choice. It is also worth noting that by doing so the vast majority of the traffic is sent from the RSUs towards the vehicles. It follows that chances of interferences (colliding frames) are close to zero since RSUs do not overlap. For the same reason, there is no MAC contention between RSUs.

The length for transmitting a frame depends on the size of the frame and on the transmission rate negotiated for the radio communication between the RSU and the passing vehicle. Indeed, 802.11p allows 8 different transmission rates ranging from 3 Mbps to 27 Mbps. In practice, an RSU regularly assesses the quality of the radio channel towards each of its destination vehicles. Based on these quality assessments it selects the highest transmission rate that still maintains a low probability of errors. Since the RSUs are fixed in space and operating in an outdoor environment, we assume that the transmission rates selected by the WiFi manager, which is responsible for selecting the transmission rate based on the assessed quality of the radio channel, may be expressed as a function of the distance separating the RSU to the destination vehicle. Figure 2 reports the transmission rates chosen by an RSU as a function of the distance to its destination vehicle using the Ideal WiFi manager of ns-3 [14]. Note that this WiFi manager skips two possible transmission rates so that six possible transmission rates (instead of eight) appear in Figure 2.

Because vehicles associated with an RSU typically have different transmission rates, their communications are exposed to the performance anomaly [9]. All the vehicles associated to the same RSU obtain (approximately) the same throughput despite having different transmission rates. This is because DCF of 802.11p provides equal treatment in the channel access.

Note that IEEE 802.11p provides multiple independent channels so that multiple communications on different channels occur at the same time without collision. Without loss of generality, we study the case of only one radio channel but all the following discussion can easily be extended to the case of multiple channels. Note also that, in the described VoD platform, all IEEE 802.11p frames used for the video transmission have the same level of priority.

Videos

Movies and series available at the back-end server are compressed using a codec that produces videos encoded at a variable bit rate. We denote by c_i the mean bit rate of a given video i averaged over its whole duration. Video frames are transmitted from the video server to the destination ve-

Figure 2: Transmission rates and their corresponding achievable throughputs for packets of 1500 bytes as a function of the distance between an RSU and its destination vehicle. Note that the two y-axes have different scales.

hicles over TCP (Transmission Control Protocol) (or over new Transport Layer protocols such as QUIC (Quick UDP Internet Connections)). We assume that RSUs always have frames waiting to be transmitted to each of their associated vehicles. In practice, when a vehicle enters into the signal range of the first RSU along the highway, it associates with it. In return, the RSU updates its forwarding table and so is the forwarding table of the video server. Hence, the video server gets a path to the vehicle in question. It then opens a TCP connection towards the vehicle and can start the transfer of the video. When the vehicle leaves the RSU signal range, the TCP connection may be closed (if a timeout due to inactivity occurs) causing the data transfer to stop. As soon as the vehicle enters into the signal range of the next RSU, Link layer/routing information will be updated on this RSU, and the data transfer will resume.

Performance parameters of interest

To evaluate the behavior of the described VoD platform, we study two performance parameters that can be computed for each vehicle.

First, we consider the total amount of data downloaded by a vehicle over its journey. It is simply computed as the sum of data received from each RSU.

The second parameter is referred to as the total “interruption time”. It corresponds to the time a vehicle spends with the playback of its video interrupted because of an empty buffer. Recall that a vehicle must buffer enough data so that it can keep playing its video during its stay in the non-covered areas between RSUs.

3.2 Analytical solution

At the level of each RSU

Let us consider a given RSU along with its associated vehicles at a given time. To begin with, it is worth noting that the current state of this two-way highway scenario can be viewed as an one-way highway scenario. It suffices to remove vehicles in one direction and to reposition them in the other direction at the same distance from the RSU. Note

Figure 3: Substituting a two-way highway scenario by a one-way highway scenario.

that this equivalence holds because of the symmetry around an RSU. Figure 3 illustrates this process. Although our analytical solution works for two-way highways, having all vehicles moving in the same direction eases its description. Therefore, in the remainder of the paper, we will consider, whenever we study one RSU that its associated vehicles are all moving in the same direction.

Let n be the current number of vehicles associated with the considered RSU. Given their distance to the RSU, each vehicle negotiates and determines its transmission rate (cf. Section 3.1). We denote by T_i ($i = 1, \dots, n$) the transmission rate of the i -th vehicle. Because IEEE 802.11p involves transmission overheads (e.g., AIFS, Backoff, ACK), we introduce the definition of achievable throughput that corresponds to the throughput a vehicle would attain if left alone with its RSU (no competing vehicles). We use A_i to indicate the achievable throughput of the i -th vehicle. Note that achievable throughputs are easily derived from the transmission rates given the size of frames (e.g., [7]). Figure 2 indicates the found values for A_i given T_i and frames of length 1500 bytes.

Having the achievable throughput of each vehicle, we aim at estimating their attained throughput, which takes into account the sharing of the radio channel between vehicles. To do that we rely on the formula proposed by Amer et al. [3] which is derived from previous works [12, 11, 13] and aims at ensuring fairness in the number of channel accesses among the vehicles (rather than the duration in which those accesses occur). Denoting the attained throughput of vehicle i by B_i , and assuming that RSUs have always frames waiting to be sent to their associated vehicles, we have:

$$B_i = \frac{1}{\sum_{j=1}^n \frac{1}{A_j}} \quad \text{for } i = 1, \dots, n. \quad (1)$$

Note also that Equation 1 implies that all the vehicles associated to the same RSU receive the same attained throughput, which is in line with the performance anomaly (cf. Section 3.1). To account for the time dependency of B_i , we use $B_i(t)$ to denote the attained throughput of vehicle i at time t .

To compute how much data a vehicle i downloads from the RSU, we need to introduce a vector whose elements correspond to the time where either the current number of vehicles associated to this RSU changes or the (negotiated) transmission rate of any of these vehicles, including vehicle i itself, changes. Let e be this vector, $|e|$ its length, and e_k its k -th element. Denoting by D_i the amount of data

downloaded by vehicle i , it follows that:

$$D_i = \sum_{k=1}^{|e|} B_i(e_{k-1}) \cdot (e_k - e_{k-1}). \quad (2)$$

At the level of the N RSUs

The derivation of the total amount of data downloaded by vehicle i along the N RSUs is straightforward. It suffices to sum up the amount of data downloaded on each RSU.

The computation of the interrupted time for a vehicle is more complicated. To begin with, we need to keep track of the current number of video frames queued in each vehicle buffer and waiting to be played. Let us consider a vehicle i . We denote by f the vector comprising all the times where the vehicle transmission rate varies, including 0's when the vehicle is not associated with an RSU. Recall that these times may correspond to a change of the distance between this vehicle and its associated RSU but may as well correspond to a change of the distance between the RSU and another of its associated vehicles. The vector f is a sub-vector of e that contains only times where there is a change impacting the vehicle i . We use f_k to refer to the k -th element of f . It follows that during the interval $[f_{k-1}, f_k]$, the amount of data received by vehicle i is equal to:

$$B_i(f_{k-1}) \cdot (f_k - f_{k-1}). \quad (3)$$

During this time interval, provided vehicle i has enough data to be played, the playback of its video plays, on average, an amount of data equal to:

$$c_i \cdot (f_k - f_{k-1}). \quad (4)$$

If the quantity computed in Equation 3 is larger than that of Equation 4, then the buffer of vehicle i fills over the duration $f_k - f_{k-1}$. Otherwise, it depletes. We use $Q_i(t)$ to denote the current length (in time units) of the queue buffered on vehicle i at time t . $Q_i(t)$ can be computed as follows. Posing $Q_i(f_0) = 0$ (buffer is initially empty), we have:

$$Q_i(t) = (Q_i(f_{k-1}) + (B_i(f_{k-1}) - c_i)(t - f_{k-1}))^+ \quad \text{for any time } t \text{ in } [f_{k-1}, f_k] \quad (5)$$

The interrupted times experienced by vehicle i over each period $[f_{k-1}, f_k]$ can now be evaluated. Let $J_i(k)$ denote the interrupted times experienced by vehicle i over the period $[f_{k-1}, f_k]$. Note that the playback of the video of vehicle i is interrupted whenever its queue length comes to 0. Hence, we can derive $J_i(k)$ as follows:

$$J_i(k) = \int_{f_{k-1}}^{f_k} \mathbb{1}_{Q_i(t)=0} dt \quad (6)$$

Finally, the total interruption time for vehicle i along the N RSUs is simply obtained by summing up all the $J_i(k)$ over k . Table 1 summarizes the main notation used in this paper.

4. PERFORMANCE AND ACCURACY OF THE PROPOSED SOLUTION

We explore a number of scenarios to study the accuracy of our proposed analytical solution. The scenarios considered include the following ranges of parameter values: the number of vehicles varies from 1 to 20 (reflecting low, average

N	Number of RSUs
I	Distance between two successive RSUs (meters)
L	Signal range of RSUs (meters)
v_i	Velocity of the i -th vehicle (meters/sec)
c	Bit rate of the codec (Mbps)
n	Current number of vehicles associated to a given RSU
T_i	Transmission rate of the i -th vehicle in a given RSU (Mbps)
A_i	Achievable throughput for the i -th vehicle in a given RSU (Mbps)
B_i	Attained throughput for the i -th vehicle in a given RSU (Mbps)
e	Vector of times where the transmission rate of a vehicle are susceptible to change within an RSU signal range
f	Vector of times where the transmission rate of vehicle are susceptible to change during its whole journey
D_i	Amount of data downloaded by the i -th vehicle in a given RSU
$Q_i(t)$	Queue length in the buffer of vehicle i at time t (sec)
$J_i(k)$	Interrupted time in the playback of the video of the i -th vehicle over the period $[f_k, f_{k-1}]$

Table 1: Main notation.

and high traffic densities), the distance between RSUs lies in between 1400 to 8000 meters (representing close and distant RSUs), and the velocity of the vehicles varies from 14 to 36 meters/sec (corresponding to slow and fast vehicles, viz. approx. 50 and 130 km/h respectively). Because a long portion of a highway can be viewed as the concatenation of shorter portions, we keep the number of RSUs set to 3 through this validation section. With these settings, on any scenario, each vehicle undergoes connected periods, in which its attained throughput depends on the number and positions of the other vehicles, and disconnected periods whose lengths depend on its velocity.

We now discuss the videos being played by vehicles in our scenarios. We consider that all vehicles are watching videos encoded at the same average bit rate and we use c to denote this bit rate. Although our approach can handle scenarios in which vehicles are watching videos encoded at different codecs, having the same value of c for all vehicles will ease the understanding of the interplay between the two considered performance parameters, namely amount of downloaded data and interruption times.

4.1 Simulation details

We implemented the proposed model in Matlab. Its execution is very fast. Performance parameters are returned at the speed of a click. Besides the proposed analytical solution, we perform realistic simulations using ns-3 [14]. The simulator implements a wireless physical layer (with possible transmission errors), the IEEE 802.11p standard, and the TCP-IP stack. At the MAC (Medium Access Control) layer, we use the Ideal WiFi manager of ns-3 that determines the IEEE 802.11p physical transmission rate (ranging from 3 to 27 Mbps) according to the SNR (Signal-to-noise Ratio).

Figure 2 depicts the behavior of this manager as a function of the distance between an RSU and a vehicle. The y-axis on the left represents the 802.11p physical transmission rates while that on the right expresses the corresponding achievable throughput taking into account the different overheads and times to access the medium (AIFS, Backoff, ACK, etc.). On top of the TCP-IP stack, we coded a video player which is instantiated on each vehicle, and an application running on the back-end video server that distributes videos to the vehicles. Note that we use a 1 Gbps wired LAN to interconnect the back-end server with the RSUs.

Each video stored at the back-end server is described by a trace file that indicates the video frame sizes and the times they are played. As mentioned earlier we consider Star Wars Episode 4 where the frame sizes vary from 13 bytes to 30 kB as well as a CBR video where the frame sizes are constant (3125 bytes). Therefore, depending on the frame sizes, an IP packet may contain only a part of a single video frame or several frames. The CBR video trace file has been set to obtain the same video bit rate as Star Wars, which is approximately $c = 0.75$ Mbps. This allows us to assess the impact of the frame sizes on the performance parameters of the VoD platform.

The transfer of video frames between the back-end server and the passing vehicles is performed using the Transport protocol TCP. This guarantees an appropriate sending rate while ensuring reliability (by retransmitting lost packets). Note that the back-end server opens only one TCP connection per vehicle. Depending on the current position of a vehicle (within or out of the signal range of an RSU), its TCP connection may be active or inactive. If active, the server reads the video frame by frame and attempts to transfer them as fast as possible. If inactive, the transfer is automatically stopped and whenever the associate vehicle re-enters in an RSU signal range, its TCP connection will be re-opened and its data transfer will resume.

On each vehicle, the video player collects and stores the received data into a buffer and plays the video frames according to the video trace. The video begins when the first video frame is received.

Note that we repeated each simulated experience 10 times with different seeds. Simulated results depicted in the following figures correspond to their mean values along with their confidence interval at a degree of 95%. The ns-3 code is made available [1] with the scripts that generate the simulations and process ns-3 output.

4.2 Scenario 1

In our first scenario, we consider a set of 5 vehicles, each with a velocity of 30 meters/sec and spaced from each other by 200 meters. Therefore, the 5 vehicles are spread over 1000 meters and, because the signal range of an RSU goes up to 700 meters on each direction (see Figure 2), the number of vehicles connected to a given RSU may vary from 0 to 5. We consider several possible distances between adjacent RSUs ranging from 1500 to 8000 meters leading to disconnected periods of 3 and 220 seconds, respectively. We run the simulator ns-3 for both videos (CBR and Star Wars) as well as our analytical solution on this scenario using a value

Figure 4: Scenario 1: Average amount of data downloaded by vehicles as a function of the distance between RSUs.

Figure 5: Scenario 1: Average interruption time experienced by vehicles as a function of the distance between RSUs.

of $c = 0.75$ Mbps and we obtain in return the associated performance parameters.

Figure 4 reports the average amount of downloaded data (in MB) computed over all vehicles. First, as expected, we observe that this value holds constant regardless of the actual distance between RSUs. Note that the simulator results may slightly vary though due to changed circumstances regarding the update of forwarding tables and time-out of TCP connections. Second, regardless of the specific video used by the simulator, the average amount of downloaded data are much alike. Indeed, although frame sizes differ, the mean number of access performed by each vehicle and its mean attained throughput remain roughly constant. Third, the analytical solution delivers accurate results with a relative error close to 10%. More precisely, the analytical solution exhibits a constant negative bias. This discrepancy results from the delay in establishing a connexion (forwarding ta-

ble, TCP setup) when a vehicle enters into the signal range of an RSU. Indeed this delay, which exists only in the simulator, postpones the association of distant vehicles with low data rates to the RSU and hence lessens the effect of the IEEE 802.11 performance anomaly (which was discussed in Section 3.1). We will further investigate this phenomenon in the next section.

In Figure 5, we represent the average interruption time experienced by vehicles for different distances between adjacent RSUs. We observe that up to a distance of 2000 meters, vehicles move along the highway with virtually no interruption in their video playback. For larger distances between RSUs, the average interruption time increases steadily. Besides, unlike the average data downloaded, there is a difference, though small, when passing from the CBR video to Star Wars. Finally, the predictions delivered by the analytical solution are close to the values delivered by the simulation with typical relative errors lying under 5%.

4.3 Scenario 2

Our second scenario studies the accuracy of the analytical solution for different vehicle density. To do that we let the number of vehicles vary from 1 up to 20 while the distance between consecutive vehicles is set to $\frac{1000}{(\text{number of vehicles})}$ meters. The distance between RSUs is kept to 2000 meters and the velocity of vehicles is set to 30 meters/sec.

Figure 6: Scenario 2: Average amount of data downloaded by vehicles as a function of the number of vehicles.

Figure 6 reports the associate results for the average amount of downloaded data computed over all vehicles. As expected, the amount of downloaded data decreases quickly as the number of vehicles grows. Our analytical solution is able to accurately capture this trend and provides accurate estimate of the average amount of downloaded data.

The results for the average interruption time experienced by vehicles are given in Figure 7. First, we observe that depending on the video being used, the simulator results may fairly differ. However, in any case, virtually no interruption times occur up to a number of vehicles of 4. Second, we observe that the analytical solution (which only considers the

Figure 7: Scenario 2: Average interruption time experienced by vehicles as a function of the number of vehicles.

average video bit rate) tends to slightly overestimate the cut-off value of the number of vehicles for the interruption time. Nonetheless, its estimates appear to be fair and in agreement with those delivered by the simulator.

4.4 Scenario 3

In our third scenario, we evaluate the accuracy of our analytical solution for different vehicle velocities. We let their velocities vary from 14 up to 36 meters/sec (approx. 50 and 130 km/h, respectively) while maintaining the distance between RSUs to 4500 meters and the number of vehicles to 5 (spaced by 200 meters).

Figure 8: Scenario 3: Average amount of data downloaded by vehicles as a function of the velocity of vehicles.

Figure 8 reports the corresponding results for the average amount of downloaded data. As expected, we observe that as their velocity grows, vehicles tend to download less. Similarly to Scenario 1, the results returned by our analytical solution tend to slightly underestimate those delivered by

Figure 9: Scenario 3: Average interruption time experienced by vehicles as a function of the velocity of vehicles.

the simulator. However, the mean relative error stays within 10%.

In Figure 9 we represent the found values for the interruption times. We notice that, depending on the video in use and the vehicle velocity, our analytical solution may over- or underestimate the interruption times. This underlies its high sensitivity to the actual video used. Nonetheless, the proposed analytical solution manages to deliver reasonably accurate estimates for the average interruption time.

5. DIMENSIONING

To illustrate the application of our analytical solution, we study the problem of dimensioning an IEEE 802.11p-based VoD platform.

Our scenario analyzes the joint effect of the distance between RSUs and number of vehicles on the performance of a VoD platform. We use the same setting than in Scenario 1 except that the vehicle velocity is kept constant to 30 meters/sec while we let the distance between RSUs vary from 1500 meters up to 8000 meters.

Figure 10 reports the values found by our analytical solution for the average interruption time. We observe that, for this setting of the VoD platform, both the number of vehicles and the distance between RSUs significantly affect the value of the average interruption time. In practice, depending on the actual distance between RSUs, the maximum number of vehicles while maintaining no interruption in their video playback varies from 1 up to 7. We believe that such figures may be of interest for dimensioning purposes of a VoD platform.

Note that the predicted interruption time can also be used to size the length of the video to buffer before its playback starts. For example, if 10 seconds of video are buffered before their playback starts, then for a distance between RSUs of 1500 meters, the VoD platform can accommodate up to 10 vehicles (instead of 7).

Figure 10: Dimensioning: Average interruption time experienced by vehicles against the number of vehicles and their velocity.

6. CONCLUSIONS

In this paper, we study the network infrastructure of an IEEE 802.11p-based VoD platform that aims at delivering video content to traveling vehicles. Videos are delivered through IEEE 802.11p Road Side Units deployed along the highway. We propose a conceptually simple analytical solution to estimate (at the speed of a click) two key performance parameters for a VoD platform: (i) the total amount of data downloaded by a vehicle over its journey and (ii) the total “interruption time”, which corresponds to the time a vehicle spends with the playback of its video interrupted because of an empty buffer. We run realistic simulations on several scenarios to evaluate the accuracy of the analytical solution. Despite its simplicity, the analytical solution provides fair estimates for the two considered performance parameters. Then, we illustrate how the proposed analytical solution can help at dimensioning the networking resources of an IEEE 802.11p-based VoD platform. Future works will aim at considering more realistic scenarios (e.g. by introducing mobility models designed for vehicles).

7. REFERENCES

- [1] Ns-3 simulations: Video on demand in an IEEE 802.11p network. http://www.anthonibusson.fr/index.php?option=com_content&view=article&id=18:ns-3-vod&catid=2:uncategorised.
- [2] IEEE standard for information technology-local and metropolitan area networks-specific requirements-part 11: Wireless LAN medium access control and physical layer specifications amendment 6: Wireless access in vehicular environments. *IEEE Standard 802.11p*, 2010.
- [3] M. Amer, A. Busson, and I. Guérin Lassous. Association optimization in Wi-Fi networks: Use of an access-based fairness. In *MSWIM, Proceedings ACM*, pages 119–126. ACM, 2016.
- [4] D. Bethanabhotla, G. Caire, and M. J. Neely. Adaptive video streaming for wireless networks with multiple users and helpers. *IEEE Transactions on Communications*, 63(1):268–285, Jan 2015.
- [5] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, March 2000.
- [6] F. Daneshgaran, M. Laddomada, F. Mesiti, and M. Mondin. Unsaturated throughput analysis of IEEE 802.11 in presence of non ideal transmission channel and capture effects. *IEEE Transactions on Wireless Communications*, 7(4):1276–1286, April 2008.
- [7] A. Duda. Understanding the performance of 802.11 networks. In *PIMRC, Proceedings IEEE*, volume 8, pages 2008–1, 2008.
- [8] C. Han, M. Dianati, R. Tafazolli, R. Kernchen, and X. Shen. Analytical study of the IEEE 802.11p MAC sublayer in vehicular networks. *IEEE Transactions on Intelligent Transportation Systems*, 13(2):873–886, June 2012.
- [9] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance anomaly of 802.11 b. In *INFOCOM, Proceedings IEEE*, volume 2, pages 836–843. IEEE, 2003.
- [10] F. Jamil, A. Javaid, T. Umer, and M. H. Rehmani. A comprehensive survey of network coding in vehicular ad-hoc networks. *Wireless Networks*, 23(8):2395–2414, Nov 2017.
- [11] O. B. Karimi, J. Liu, and J. Rexford. Optimal collaborative access point association in wireless networks. In *INFOCOM, Proceedings IEEE*, pages 1141–1149. IEEE, 2014.
- [12] L. Li, M. Pal, and Y. R. Yang. Proportional fairness in multi-rate wireless LANs. In *INFOCOM, Proceedings IEEE*, pages 1004–1012. IEEE, 2008.
- [13] W. Li, S. Wang, Y. Cui, X. Cheng, R. Xin, M. A. Al-Rodhaan, and A. Al-Dhelaan. AP association for proportional fairness in multirate WLANs. *IEEE/ACM Transactions on Networking*, 22(1):191–202, 2014.
- [14] *The Network Simulator ns-3*. <https://www.nsnam.org/>.
- [15] C. Quadros, A. Santos, M. Gerla, and E. Cerqueira. QoE-driven dissemination of real-time videos over vehicular networks. *Computer Communications*, 91-92:133 – 147, 2016.
- [16] L. Sun, H. Shan, A. Huang, L. Cai, and H. He. Channel allocation for adaptive video streaming in vehicular networks. *IEEE Transactions on Vehicular Technology*, 66(1):734–747, Jan 2017.
- [17] O. Tickoo and B. Sikdar. Modeling queueing and channel access delay in unsaturated IEEE 802.11 random access MAC based wireless networks. *IEEE/ACM Transactions on Networking*, 16(4):878–891, Aug. 2008.
- [18] M. Xing, J. He, and L. Cai. Utility maximization for multimedia data dissemination in large-scale vanets. *IEEE Transactions on Mobile Computing*, 16(4):1188–1198, April 2017.
- [19] J. Zheng and Q. Wu. Performance modeling and analysis of the IEEE 802.11p EDCA mechanism for VANET. *IEEE Transactions on Vehicular Technology*, 65(4):2673–2687, April 2016.