

The Impact of Transmission Power on the Safety-Related Performance of IEEE 802.11p

Martin Klapez, Carlo Augusto Grazia, and Maurizio Casoni

Abstract—In this paper, we report and analyze results from a field test campaign aimed at assessing the real-world performance of IEEE 802.11p for safety-related applications. Regulator bodies in the EU and the US recommend the use of domestic-level transmission powers for general usage while granting the opportunity of employing higher powers for safety purposes. While the latter are generally expected to result in higher network performance, we want to quantitatively assess what they translate to in terms of metrics relevant from an application point of view, such as, goodput, round-trip times, jitter, and losses. These are studied and associated with conditions out of the network control at urban and suburban speeds, that is, in absence or in presence of different degrees of congestion, and in complete or partial line-of-sight. In general, we find evidence of an inverse correlation between the degree of congestion and the benefits granted by higher transmission powers. Up to the tested speed, IEEE 802.11p appears able to provide safety guarantees even at relatively long distances, as long as an appropriate transmission power is employed. At the same time, it must be acknowledged that higher power alone cannot overcome the significant dips in performance resulting from highly congested environments and non-line-of-sight scenarios.

Index Terms—Application Performance, DSRC, Field Trials, IEEE 802.11p, Safety, V2X

I. INTRODUCTION

The IEEE 802.11p Wireless Access for Vehicular Environments (WAVE) is a collection of amendments pertaining to vehicular settings made to the IEEE 802.11-2012 standard [1] (now superseded by the IEEE 802.11-2016 standard [2]). The latter is the basis for the so-called Dedicated Short-Range Communications (DSRC), which enables vehicle-to-vehicle and vehicle-to-infrastructure communications. DSRC was originally designed to enable collision-prevention applications [3], but the interest in Vehicle-To-Everything (V2X) transmissions quickly spread to encompass consumer-centric services like entertainment, traffic data, navigation, and also autonomous driving. Traditionally more focused on the latter uses, Cellular-V2X (C-V2X) and, lately, 5G New Radio C-V2X (NR C-V2X), have emerged as competing technologies with their own selling points as, for instance, the leverage of existing Long-Term Evolution (LTE) infrastructures, the coverage of larger areas by single radio items, and the addressing of some of the IEEE 802.11p's shortcomings. However, 802.11-based DSRC is reliable, almost patent-free, easy and cheap to implement, already available, and even lately integrated into the vehicles of some manufacturers [4]. It is also being evolved

into a new backward-compatible standard, IEEE 802.11bd, which also aims to address some of the shortcomings reported for IEEE 802.11p [5]. Both DSRC and C-V2X have backers with significant industrial weight [6], [4], and both are being considered by regulatory organizations [7], [8].

This work tries to evaluate the user-level performance of a network based on IEEE 802.11p deployed to provide safety-related V2X services in urban and suburban environments. Our goal is to understand the safety guarantees that can be expected during real-world operations in these scenarios. With respect to the many related works in the literature, none of the performance figures presented here derive from analytical evaluations or simulations but, instead, all result from extensive field trials that have been conducted in mobility against distance, network congestion, and transmission powers, with these variables being alone and combined. In analyzing the outcomes, we observe the behavior of several metrics that pertain to the network, transport, and application layers, i.e., TCP Goodput, UDP Goodput, RTT, Jitter, and Datagram/Packet Loss, shedding light on the application-level performance guarantees that can be assured under the tested conditions, with particular focus on the impact of transmission power on performance and in what terms this translates on specific metrics that can affect end-users.

The paper is organized as follows. Section II briefly overviews related works. Section III summarizes the IEEE 802.11p standard, while IV describes the field setup. Section V presents and analyzes the results, and Section VI extracts final considerations, concluding the article.

II. RELATED WORKS

This section briefly overviews representative works in the literature that contain performance evaluations of IEEE 802.11p through field tests. In [9], the authors try to improve the IEEE 802.11p channel estimation's robustness with a dynamic equalization scheme, identifying several performance trends related to the channel access. Software-Defined Radio (SDR) is employed in works such as [10], where the authors introduce a framework to link simulations with real experiments and validate it with field trials by investigating L2 frame delivery ratio versus distance. L1 measures are collected and analyzed in [11] during field tests with four trucks in platooning formation. Likewise, platooning trials are presented in [12], [13] with application-specific metrics. This is often the case when the works are focused on specialized scenarios. For instance, a Virtual Traffic Light algorithm is field-tested in [14]. Green Light Optimal Speed Advisory systems are also

M. Klapez, C.A. Grazia, and M. Casoni are with the Department of Engineering *Enzo Ferrari*, University of Modena and Reggio Emilia, via Pietro Vivarelli, 10, 41125, Modena, Italy (e-mail: {martin.klapez, carloaugusto.grazia, maurizio.casoni}@unimore.it).

Device	Arch.	Network card (NIC)	Working freq.	Radio Chains	Set [Max] NIC Output Power	Antenna [Gain]	EIRP (Est.)
MikroTik RB912UAG-5HPnD	MIPS	AR9342 SoC	5.825 GHz	2x2	+18 (+15, +15) / +27 (+24, +24) [+30] dBm	External, S151FL-5-RMM-2450S, omnidirectional [+5 dBi]	+23 dBm / +32 dBm

TABLE I: Technical specifications of test devices

often scrutinized in field tests [15]. [16] focus on latencies and Received Signal Strength (RSS), while [17] studies RSS and Packet Delivery Ratio (PDR) at intersections. PDR on the 5.9 GHz band is analyzed in detail within [18], which considers an urban environment where network devices operate with a +17 dBm transmission power. The relation of the latter with RSS and PDR is characterized in [19] considering various scenarios. The PDR of specific messages is instead recorded with ITS-G5 prototypes and reported in [20], [21]. [22] reports field tests that, in terms of road shape, path length, and investigated metrics, are similar to our preliminary mobility tests presented in [23], although they are conducted at lower speeds. In [24], we presented a methodology to study packet collision interference due to the hidden terminal problem. In [25], authors test the number of received frames when employing IEEE 802.11p in a Disruption Tolerant Network on a highway. [26] is also based on a highway setup and tests L1 and L2 metrics with the addition of L4 data rate at 80 km/h and 120 km/h, considering shadowing effects and other issues. [27] reports packet loss and latencies resulting from a field trial performed at speeds up to 100 km/h. Up to the same speed, authors in [28] analyze the performance in terms of L1 metrics resulting from an extensive array of road trials, including tests of safety-related scenarios such as intersection collision warning and precrash sensing. 100 km/h is also the maximum speed employed in [29] with field trials that test packet error rate versus distance and speed.

III. IEEE 802.11P IN A NUTSHELL

V2X DSRC operations have been generally defined to work in the 5.9 GHz band. IEEE 802.11p Wireless Access for Vehicular Environments (WAVE) [30], based on the ASTM E2213 v2 standard [31] (the updated v3 is available at [32]), is an amendment to the familiar IEEE 802.11 standard (WiFi). With respect to specifications made for domestic or other environments, this amendment is designed to meet the challenges of high-speed mobile radio environments, such as strong Doppler shifts, mutable multipath conditions, and fast connection establishment. Indeed, the pressure for faster access to the physical medium has been the original primary motivation for the amendment. The new access rules that have been integrated are called “Outside the Context of a Basic Service Set (BSS)” (OCB). OCB allows unicast, multicast, and broadcast data communications without any MAC sublayer setup. It guarantees, for safety-related applications, that no coexistent BSS will operate on the 5.9 GHz band. IEEE 802.11p employs the Orthogonal Frequency Division Multiplexing (OFDM) protocol derived from the IEEE 802.11a standard. DSRC applications are expected to mainly use the 10 MHz channel, supporting transmissions with PHY data

rates up to 12 Mbit/s. Further details on the subject can be found in [3]. Our tests have been performed using 16 points Quadrature Amplitude Modulation (16-QAM) as the modulation technique with a Forward Error Correction (FEC) coding rate of 1/2. IEEE 802.11p specifies the PHY and MAC layer operation while other standards define the upper layers, namely the IEEE 802.2 for the LLC layer, the IEEE 1609 family, and the SAE application standards. In alternative to IEEE 1609, both the the ETSI ITS-G5 [33] and the US DSRC [3] architectures contemplate the use of the regular TCP/UDP-IP stack for the network and transport layers, as we did here. The use of TCP/UDP-IP simplifies both the testbed deployment and the reproducibility of the experiments on general-purpose devices. For the same reason, we use IPv4 instead of IPv6 as defined by [33], [3]. Safety messages as defined by ETSI divide in Cooperative Awareness Messages (CAMs) [34] and Decentralized Environmental Notification Messages (DENMs) [35]. CAMs are periodic messages designed to exchange status among nearby vehicles with a default frequency of 10 Hz, while DENMs are occasional warnings repeatedly issued to communicate road hazards.

IV. FIELD SETUP

A. Test Devices

We set the packet size and the packet interval rate following the ETSI standard [34] in order to model CAM messages. The network equipment embedded in a vehicle is generally referred to as an OnBoard Unit (OBUs), while an infrastructural instance is called a RoadSide Unit (RSU). For both OBUs and RSU we used the MikroTik RB912UAG-5HPnD in Table I (MTik from now on), general-purpose devices that allow the use of open-source software and the tuning of Equivalent Isotropic Radiated Power (EIRP) to appropriate levels. In fact, we wanted to employ an EIRP that, while abiding by the law limits, can be set to operate as close as possible to the specifications for safety-critical applications. It appears that standard commercial devices limit the transmission power up to between +22 dBm and +24 dBm, which is consistent with the maximum power allowed by the Mask C ETSI regulations for regular applications. However, both the ETSI in the EU and the FCC in the US approve and recommend the use of much higher power values for safety-critical applications. ETSI (ITS-G5) allows a peak power of +33 dBm [36], while the US FCC (DSRC) pushes the limit even higher up to +44.8 dBm [37], depending on the channel. We installed on MTiks the Linux kernel version 4.4.153 with the ath9k driver, in order to have access to the open-source implementation of IEEE 802.11p that enables, with compatible network cards, the 5.9 GHz band with 10 MHz wide channels. Furthermore, it also gives us access to the OCB mode support in Linux, to the wireless

Metrics Investigated	EIRP (Est.)	Congestion
TCP & UDP Goodput (Mbit/s)	+32 OBUs, +32 RSU (dBm)	None (0 CAM/s)
RTT & Jitter Goodput (ms)	+23 OBUs, +32 RSU (dBm)	15 vehicles (150 CAM/s)
Datagrams Lost (%)	+23 OBUs, +23 RSU (dBm)	30 vehicles (300 CAM/s)

TABLE II: Test Settings

network card driver, and to the configuration tool `iw` [38], which allowed us to tune the transmission power. The most notable feature of these devices, in fact, is the relatively high transmission power of the integrated wireless card, which can pass up to 30 dBm into the antennas, although they operate by default at +27 dBm. As external antennas, we used two isotropic Laird Connectivity S151FL-5-RMM-2450S, with a gain of 5 dBi. As the feeder cables are very short (≈ 20 cm), with negligible losses, the maximum total EIRP we employed is between +32 dBm and the ETSI upper bound of +33 dBm. MTik hardware reportedly supports frequencies of up to 5.875 GHz, although it appears that these are by default software-locked to a maximum of 5.825 GHz. Therefore, considering the IEEE 802.11p 5.9 GHz range of between 5.850 and 5.925 GHz, the maximum nominal frequency falls short by only 25 MHz with respect to the minimum standardized frequency.

B. Test Setup

In our tests, we considered fixed the frequency for both V2V and V2I communications, i.e., modeling a situation in which both V2V and V2I traffic would travel on the same frequency, as this tends towards a worst-case analysis and adhere to the fact that it is unknown how each vehicle manufacturer or infrastructure tenant will implement V2X services. we employed two cars with very similar power and acceleration (≈ 0.1 s) as our vehicles. The total distance traveled by each vehicle in every test run amounts to 1.5 km on a straight road, with two areas where vehicles suddenly lose line-of-sight (LOS) for ~ 50 m altogether. The first 500 m have been used in each direction to accelerate the vehicles to the test speeds; for the central 500 m, which is the area for which the measures are reported, the cars' velocity has been kept constant, while the last 500 m have been used to safely decelerate and bring the vehicles to a stop. We used two cars traveling in opposite directions to safely reach higher relative speeds than the maximum velocities individually attainable by our test vehicles. One served as RSU, while the other as a generic OBU. We ensured that results were consistent with those attainable with a single vehicle as OBU and a roadside-mounted RSU by also performing control experiments. In the tests with congestion, two devices mounted at the side of the road were used to generate CAM streams and signal the starting and end points for the measurements, as they were set to the 500 m and 1 km marks. Roadside devices were installed at heights of approximately 3 m and 2.5 m., as different vehicles have different heights, and RSU manufacturers can recommend different mounting heights for their devices. The equipment on vehicles was mounted on their roofs, at the height of approximately 2 m for the RSU and 1.5 m for the OBU. The antennas detailed in the previous subsection were

set up pointing upwards in order for their blind spot to be in a perpendicular direction with respect to those the vehicles travel to.

Table II lists the test settings. Every combination has been put to the test. Therefore, each test is associated with a combination of three parameters, namely, metric, EIRP, and congestion amount. The tests presented here were all performed at a GPS speed of 100 km/h (110 km/h indicated by the vehicles' dashboards). Lower and higher speeds have also been tested. For the former, 100 km/h can be safely considered as a lower bound in terms of performance. The latter are not discussed here as they are out of the scope of this paper. The test type reflects the data collected during the run (UDP Goodput, Datagrams Lost, and Jitter were collected together). The estimated EIRPs have been organized in three sets. The first (OBUs +32 dBm, RSU +32 dBm) models the reference scenario in which all devices operate on powers defined for safety-related applications. The second (OBUs +23 dBm, RSU +32 dBm) represents a setting in which RSUs are set on powers defined for safety operations, while OBUs function on non-safety specifications. The third (OBUs +23 dBm, RSU +23 dBm) has been added to represent operativity in non-safety scenarios. The three sets of EIRPs will be referred to throughout the paper as *safety EIRP* or +32 dBm, *hybrid EIRP* or (+32,+23) dBm, and *domestic EIRP* or +23 dBm.

In addition to the tests with no congestion, two tiers of CAM streams have been introduced in separate runs, for an equivalent of 150 additional CAM/s (15 vehicles) and 300 additional CAM/s (30 vehicles) broadcasted on the network at the same time. The `netcat` tool has been used to run multiple CAM instances on each device in order to model a greater number of nearby vehicles generating V2V traffic with respect to the number of available devices. Given that few congesting sources may not be able to fully capture the totality of PHY-layer effects that are present when more devices are employed, we investigated the magnitude of the issue by using a higher number of different IEEE 802.11p apparatus in our possession. As these can only reach EIRPs of about +23 dBm, they have been placed closer together than the devices listed in Table I. From our evaluations, the variation in performance resulting from using n physical devices versus n `netcat` instances in a lower number of devices amounts to a difference between 1% and 3%.

V. SAFETY-RELATED PERFORMANCE

A. Goodput

Figure 1 plots the goodput value distributions for TCP. Green is associated with safety EIRP, blue with hybrid EIRP, and red with domestic EIRP. The box-and-whisker plots are composed of the boxes, that group values between the first and

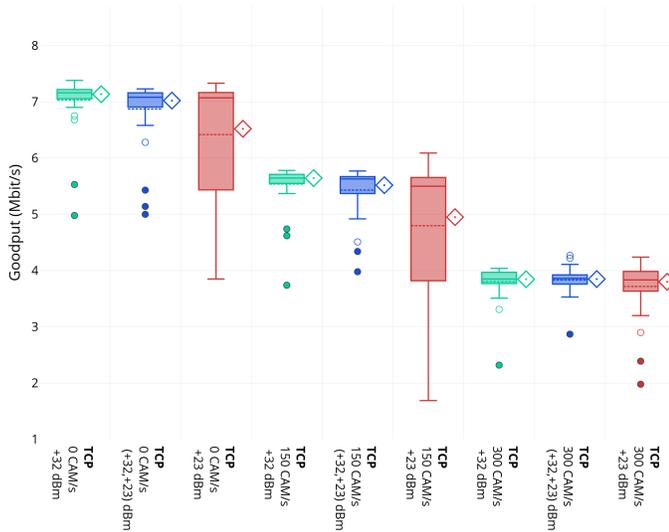


Fig. 1: TCP Goodput, Distributions

the third quartiles, and the whiskers, which show the upper fence on top and the lower fence at the bottom. The full line inside the boxes is the median value, while the dashed line is the mean. There are also points outside the box-and-whisker plots. The full circles are outliers, while the empty ones are suspected outliers. The diamonds with a dot in the middle represent the means computed without considering the outliers, that is, considering LOS only. What follows may not be immediately apparent, but the chart indicates that the magnitude of the performance drop under occasional non-LOS scenarios is generally greater at safety EIRP and, as such, the minimum values are lower than those found with hybrid EIRP. This does not apply when considering domestic EIRP that, instead, exhibits the lowest minimum values regardless of the amount of congestion. Therefore, starting with a comparable but slightly lower average goodput correlates with a *lower* drop in the minimum performance that can be experienced due to the sudden presence of obstacles. While it may be coincidental, this may also be due to the congestion control algorithm in TCP [39] (Cubic in our case), As it will be shown, in fact, with UDP, that does not attempt to detect congestion and regulate its sending rate accordingly, this possible phenomenon does not appear.

Another information that can be extracted from the chart is a direct correlation between transmission power and goodput variability. Higher transmission power results in more stable data rates, regardless of the amount of congestion present at any specific moment. It is also worth noting that, with no congestion or with 150 CAM/s, the +23 dBm whiskers include values that with higher powers are regarded as outliers. This is due to the fact that lower transmission power means shorter radio coverage and, consequently, lower goodput at the edges of the covered area, as the median remains strikingly similar. Indeed, domestic EIRP is able to guarantee optimal performance only within 350 m, hybrid EIRP within 450 m, and safety EIRP within more than 500 m, as illustrated by Figure 2. The latter shows the TCP goodput in the first 100 m of the test area. As it can be seen, lower powers translate

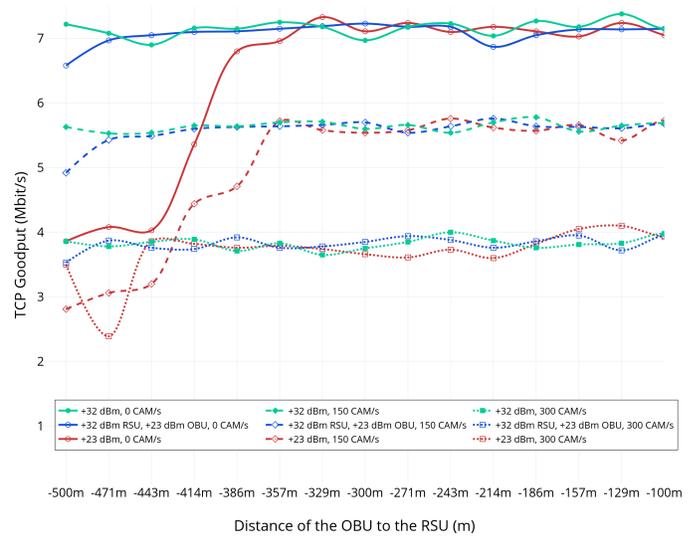


Fig. 2: TCP Goodput in the initial 100 m of the test area

to lower performance as the OBU moves away from the RSU, as expected. However, Figure 2 also attests that, with safety EIRP, RSUs may be likely deployed farther than anticipated.

The final note regards the average values. While, in general, with lower power one can expect a lower goodput, the difference becomes very narrow with high amounts of congestion, where domestic EIRP behaves almost as good as safety EIRP. This is because, in this case, higher EIRPs are held back by reaching higher goodput in the first place. Therefore, it can be concluded that higher EIRPs earn better stability in the presence of issues such as distance and non-LOS, but does not intrinsically guarantee higher goodput. All EIRPs result in stable goodput values in the presence of any level of congestion, which is instead a source of goodput degradation regardless of the EIRP.

Figure 3 plots the goodput value distributions for UDP. As with TCP, we always measured the actual goodput at the receiver. In general, the considerations made for TCP hold also

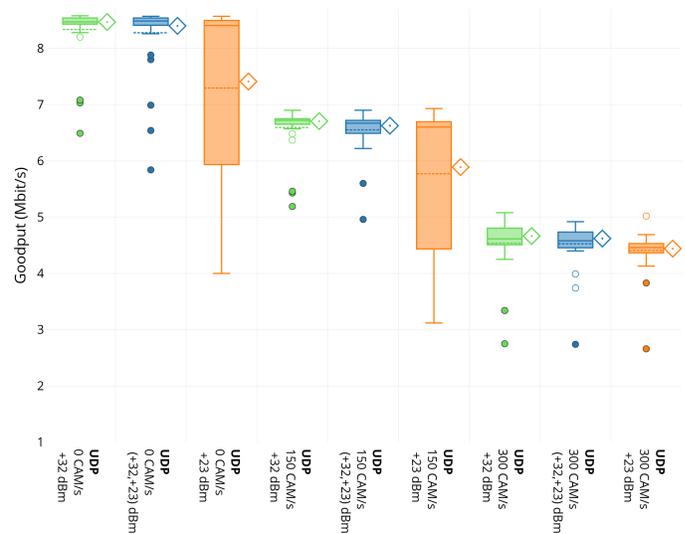


Fig. 3: UDP Goodput, Distributions

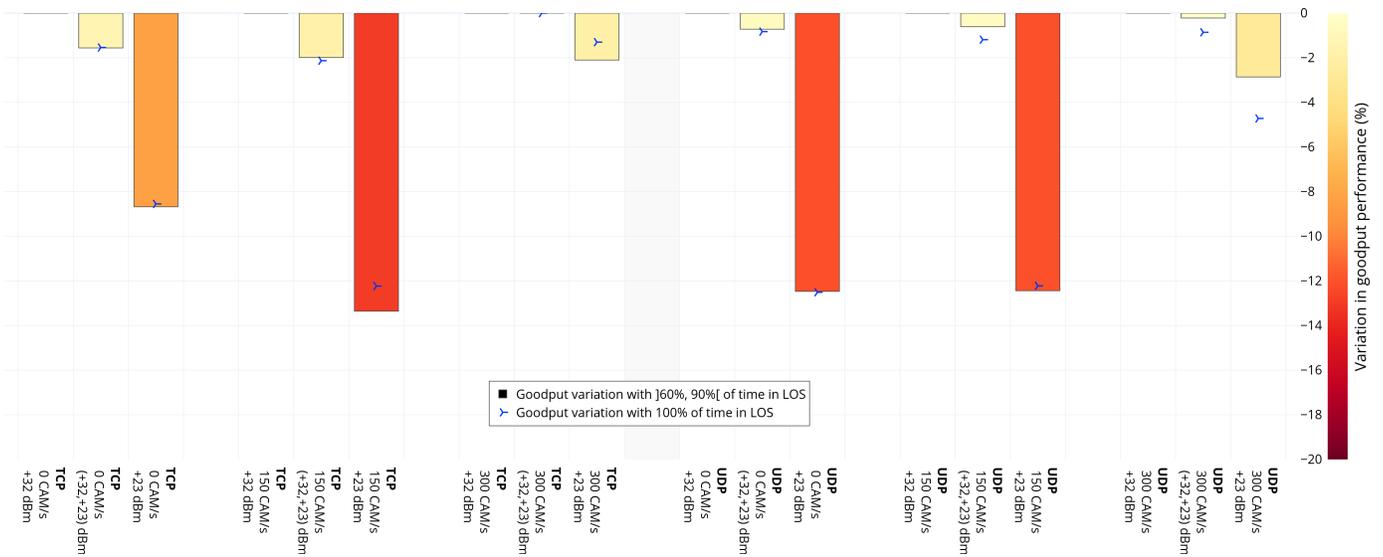


Fig. 4: Impact of Transmission Power on Goodput

here. Regardless of congestion, there is a direct correlation between transmission power and goodput variability, or, in other words, data rate stability. +23 dBm whiskers include values that with higher powers are regarded as outliers, and the same coverage issues present with TCP are also present when employing UDP. The difference of the averages among transmission powers still becomes very narrow when congestion is high, although here the discrepancy between safety or hybrid EIRP and domestic EIRP is more noticeable even with 300 CAM/s. Higher EIRP earn better stability in the presence of issues such as distance and non-LOS, but does not intrinsically guarantee higher goodput, that is, in turn, more affected by congestion. However, in one thing there is a different trend, that is, the minimum values of hybrid EIRP are always lower than those of safety EIRP. In this case, starting with a comparable but slightly lower average goodput correlates with a *greater* drop in the minimum performance that can be experienced due to the sudden presence of obstacles. As here there is no congestion control performed by the transport protocol, this makes sense if put into context with the speculation above about the reason for the opposite behavior experienced with TCP. However, our data can only suggest that such behavior may exist, we cannot either transform the speculation in a claim, or dismiss it.

Figure 4 sums up the observations. The diagram is divided into two main parts, left for TCP and right for UDP, and plots the variation in goodput performance that resulted from the employment of different transmission powers. The baseline used to compute the variations is always the test with safety EIRP. It should be noted that the y-scale has been reduced for readability, i.e., it goes from 0% (no change in goodput) to -20% (maximum worsening in goodput). Moreover, the diagram contains two overlaid plots; the yellow/red bar chart lay out the averages, while the scatter points in blue report the values obtained by filtering out all samples where devices were either out-of-coverage or in non-LOS. In the bar chart, with respect to the baseline, a yellow color means limited change, a

red color means significant change, and a dark red color means very high change. From both sides of the chart is immediately apparent that the impact on performance of domestic EIRP is significantly larger than that of hybrid EIRP. Furthermore, Figure 4 confirms our previous observations where we claimed that congestion may have a greater effect on goodput than transmission power. Indeed, in the 300 CAM/s columns is it clear that the impact on performance of even +23 dBm is very limited, about 2% for TCP and 3-4% for UDP. In the other cases, only hybrid EIRP is able to maintain such a restricted burden on the goodput, while domestic EIRP often exceeds 10%.

B. Latencies

Figure 5 shows the distributions for RTT and jitter over the lower horizontal axis, and datagrams lost as a percentage on the total over the upper horizontal axis. The turquoise box-and-whisker plots show the RTT, while the yellow violin plots represent the jitter. The latter plot type was chosen not because the jitter distribution data is multimodal (it is, but mostly due to outliers) but mainly to better differentiate jitter from RTT in the diagrams. The full lines in the boxes are the medians, while the dashed lines are the means for both RTT and jitter. The pink semi-transparent diamonds display datagrams lost. Where their opacity increases, it means that there are multiple samples for that loss percentage. The full pink vertical lines serve as indicators for the average loss values.

Figure 5 attests that RTTs stay within a very acceptable range most of the times. With no congestion, the RTT remains within 2 ms regardless of the transmission power, with the safety EIRP giving the lowest figures on average. With 150 CAM/s, the RTTs increase but stay on average within the 3 ms mark, with peaks that never reach 5 ms even at domestic EIRP. With respect to the case without congestion, both the box and the whiskers are larger, meaning that the variability of results is higher. This, as visible from the yellow violin plots underneath, is also attested by the jitter. With 300 CAM/s, the RTT is

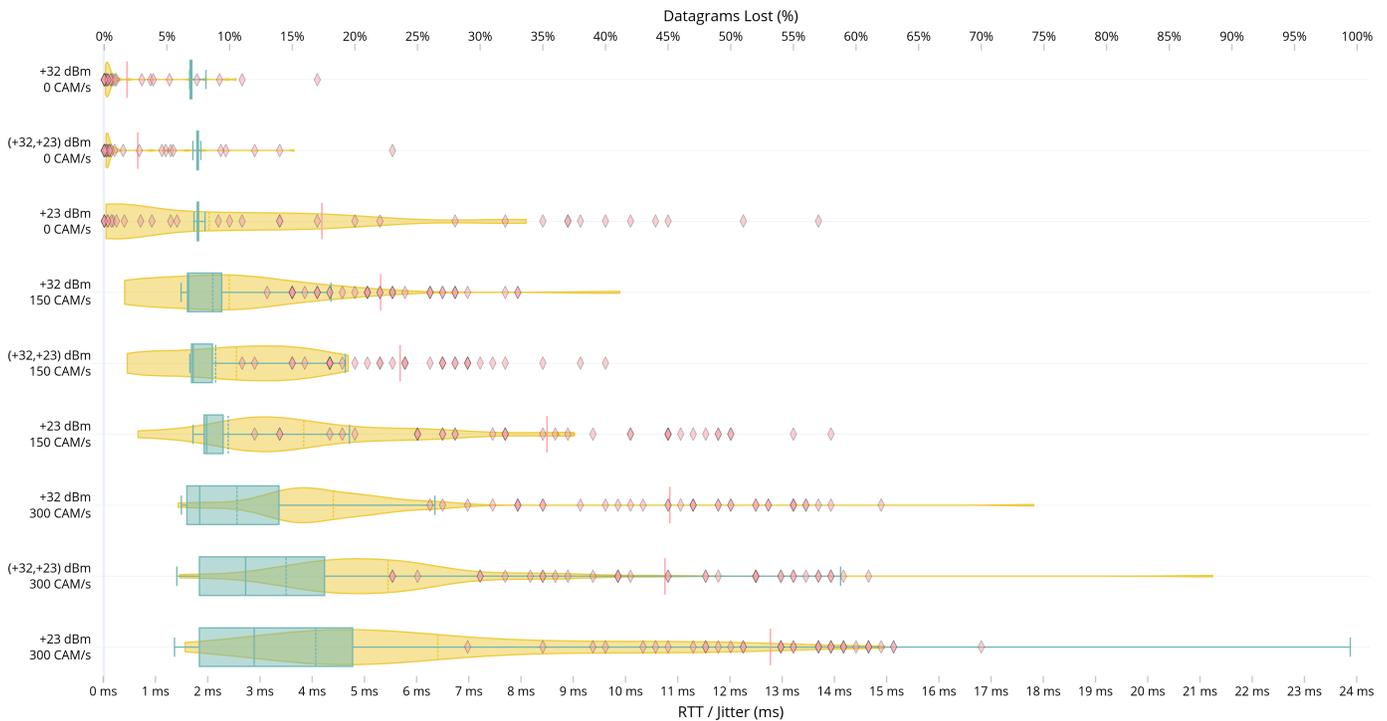


Fig. 5: RTT, Jitter, and Datagrams Lost, Distributions

visibly higher. Only safety EIRP seems able to guarantee worst-case performance within the 10 ms mark, while hybrid EIRP reaches 14 ms and domestic EIRP almost 24 ms. However, those are maximum values, and can be misleading, as the median line shows that the majority of figures group around values between ~ 2 ms and ~ 3 ms. Indeed, the means are higher than the medians but not by much, between ~ 2.5 ms and ~ 4 ms for +32 dBm and +23 dBm, respectively. We can conclude that RTTs are, again, more affected by congestion than transmission power. However, a proper safety EIRP guarantees lower latencies on average *and* significantly lower maximum latencies when under congestion.

Jitter is less linear in its behavior than RTT. Without congestion, there is almost no difference between safety EIRP and hybrid EIRP, with maximum values of ~ 2.5 ms and ~ 3.5 ms, respectively. Domestic EIRP, instead, result in much higher variability and in a significantly higher mean, that moves from ~ 0.1 ms to more than 2 ms. The same can be said for the maximum value, which lands at ~ 8.5 ms. The addition of congestion cause jitter in the network also with higher powers. While only a few values exceed the 6 ms mark, from Figure 5 we can extract an inverse correlation between jitter and transmission power, with the amplitude of the violin plots generally shifting right as the EIRP gets lower. Regarding the means, they are significantly different only between domestic EIRP and either hybrid EIRP or safety EIRP, as in the latter two the average falls within the margins of error. With 300 CAM/s congesting the channel, the maximum jitter is significantly higher, being within ~ 15 ms and ~ 21.5 ms. While the maximum values for higher powers are clearly outliers, the distinction is more subtle with +23 dBm, as the distribution of the other samples spans a larger

area. Interestingly, the maximum amplitude is virtually the same between hybrid EIRP and domestic EIRP. This is also apparent from the 150 CAM/s case, suggesting that jitter may be more dependent than RTT and goodput to the transmission power employed by the OBU's, instead of that of the RSU. With maximum congestion, averages are different in all three cases, with a clear inverse correlation between transmission power and means.

C. Datagram Losses

From Figure 5, it can be observed how the majority of datagrams lost samples concentrate around zero when the network is not congested, regardless of the transmission power. Yet, the effect of employing lower powers is visible, as higher samples start to appear and the mean moves right. Domestic EIRP causes a much more significant deviation from the optimal given by +32 dBm than hybrid EIRP. As expected, increasing the congestion also increases the losses that, however, almost always stay within the 50% threshold with 150 CAM/s, with only +23 dBm causing in two instances samples with losses between 50% and 60 %, as it is also the case with no congestion. Again, the trend of averages seen with 0 CAM/s repeats itself; safety EIRP and hybrid EIRP stay within 20% and 25% of datagrams lost, while domestic EIRP reaches a value of almost 35%. The same applies, in a smaller measure, with 300 CAM/s, where losses almost reach the 70% mark. Yet, the mean is higher than 50% only with +23 dBm, while +32 dBm resulted in an average value that is even slightly worse than the hybrid EIRP, although the difference is so small that is probably within reasonable margins of error. Between the $\sim 45\%$ of the higher EIRPs and the $\sim 53\%$ of the domestic EIRP, the difference is significant but not vast

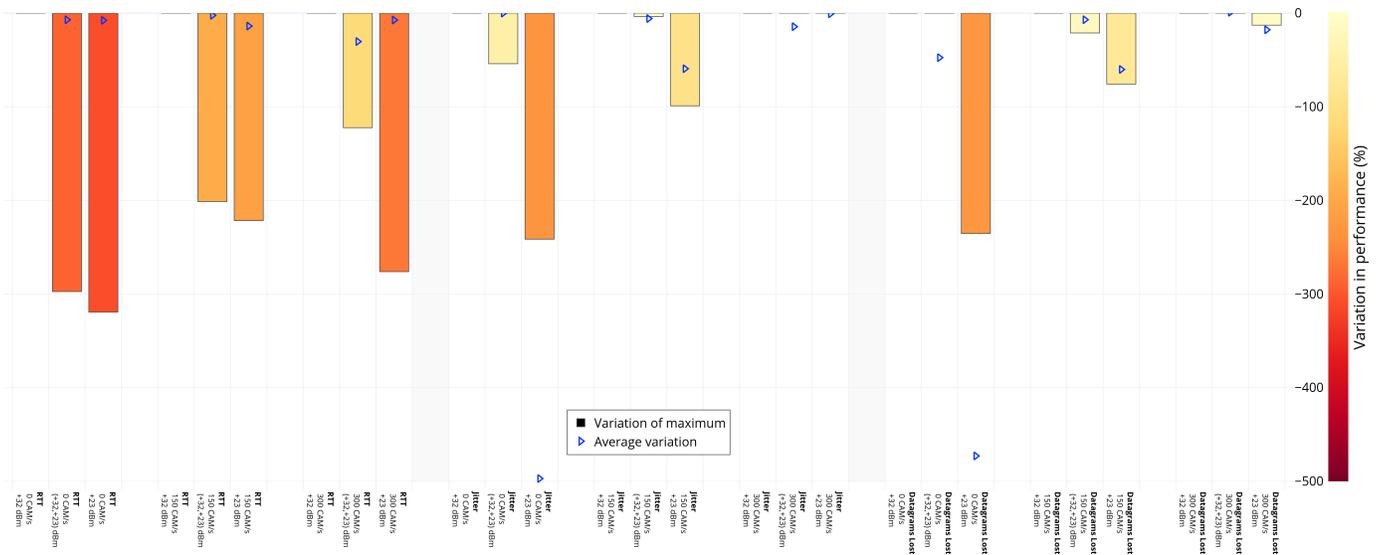


Fig. 6: Impact of Transmission Power on Latencies and Datagram Losses

as in the cases with lower congestion. Yet, we can say that, on average, at least one packet over two will be delivered when at least the RSU is operating at +32 dBm, while we cannot make the same claim with +23 dBm. Concluding, we can affirm that, unlike the previous metrics, the amount of datagrams lost are similarly dependent on both congestion and transmission power.

It is important to note that all these losses are counted on a UDP stream that tries to saturate the available bandwidth, which is unrealistic to expect in a safety-only transmission scenario. However, datagrams lost become relevant for mixed scenarios where non-safety transmissions may concurrently occur [40] or for specific mission-critical safety-related applications where RSUs may have to transmit data streams, for instance, video streams from/to first responders during an emergency.

Figure 6 summarizes the impact of transmission power on latencies and datagram losses. The diagram is divided into three main parts; from left to right, it shows RTT, jitter, and datagrams lost, respectively, plotting the variation in performance that resulted from the employment of different transmission powers. Again, the baseline used to compute the variations is always the test with safety EIRP. In this case, while the colors' meaning remains the same as in Figure 4, the y-scale goes from 0% to -500%. Here, it is clear that the most impacted metric is RTT, for which the considerations above about goodput do not hold. RTT, in fact, is severely impaired with respect to safety EIRP even with maximum congestion. The chart also suggests an inverse relation between the worsening caused by the hybrid EIRP and the degree of congestion. This seems to hold for jitter and datagrams lost too that, instead, indicate a leveling in the performance difference with respect to the baseline with maximum congestion, as it was the case for TCP and UDP goodput. This is not surprising, as the latter metric is extracted together with jitter and datagram losses.

VI. CONCLUSIONS

This paper presented the outcomes of IEEE 802.11p field trials aimed at assessing safety-related application-level performance. We focused on speeds targeting urban and suburban environments, with the goal of understanding the safety guarantees that can be expected during real-world operations in these scenarios and the impact that different transmission powers may have on performance. With the possible exception of round-trip times, we found an inverse correlation between the degree of congestion and the benefits granted by higher transmission powers. However, while in cases of very high congestion the averages may remain similar as the transmission power decreases, we also found a direct correlation between the latter and the performance variability. In other words, all metrics' stability is unequivocally impacted by the transmission power, under any level of network congestion. While this finding is secondary when considering goodput for safety-related applications, is instead meaningful when considering latencies and packet losses, metrics much more important in such scenarios. In conclusion, up to the tested speed, IEEE 802.11p appears able to provide safety guarantees even at relatively long distances, as long as an appropriate transmission power is employed. This principle applies in particular to delays and losses, where transmission power directly affects the performance level and the performance window that safety-related applications may be guaranteed to fall in. At the same time, it must be acknowledged that higher power alone cannot overcome the significant dips in performance resulting from highly congested environments and non-line-of-sight environments. Possible future works include investigating the performance of IEEE 802.11bd and Cellular-V2X, and more in-depth analyses of the relation between congestion and non-LOS.

REFERENCES

- [1] "IEEE Standard for Information technology-Telecommunications and information exchange between systems Local and metropolitan area

- networks-Specific requirements-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” *IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007)*, Mar 2012.
- [2] “IEEE Standard for Information technology—Telecommunications and information exchange between systems-Local and metropolitan area networks—Specific requirements-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” *IEEE Std 802.11-2016 (Revision of IEEE Std 802.11-2012)*, Dec 2016.
 - [3] J. B. Kenney, “Dedicated Short-Range Communications (DSRC) Standards in the United States,” *Proceedings of the IEEE*, vol. 99, no. 7, pp. 1162–1182, July 2011.
 - [4] “Toyota Comment 4.26.19 FINAL to: ET Docket No. 13-49, GN Docket No. 18-357,” FCC Filing, 2019.
 - [5] G. Naik, B. Choudhury, and J.-M. Park, “Ieee 802.11bd 5g nr v2x: Evolution of radio access technologies for v2x communications,” *IEEE Access*, vol. 7, pp. 70 169–70 184, 2019.
 - [6] H. Seo, K. Lee, S. Yasukawa, Y. Peng, and P. Sartori, “LTE evolution for vehicle-to-everything services,” *IEEE Communications Magazine*, vol. 54, no. 6, pp. 22–28, June 2016.
 - [7] “Request for Comments: V2X Communications,” US Department of Transportation, DOT-OST-2018-0210-0001, 2018.
 - [8] K. Kiela, V. Barzdenas, M. Jurgo, V. Macaitis, J. Rafanavicius, A. Vasjanov, L. Kladovcikov, and R. Navickas, “Review of V2X-IoT Standards and Frameworks for ITS Applications,” *Applied Sciences*, vol. 10, no. 12, p. 4314, Jun 2020. [Online]. Available: <http://dx.doi.org/10.3390/app10124314>
 - [9] J. A. Fernandez, K. Borries, L. Cheng, B. V. K. Vijaya Kumar, D. D. Stancil, and F. Bai, “Performance of the 802.11p Physical Layer in Vehicle-to-Vehicle Environments,” *IEEE Transactions on Vehicular Technology*, vol. 61, no. 1, pp. 3–14, Jan 2012.
 - [10] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, “Performance Assessment of IEEE 802.11p with an Open Source SDR-Based Prototype,” *IEEE Transactions on Mobile Computing*, vol. 17, no. 5, 2018.
 - [11] K. Karlsson, J. Carlsson, M. Larsson, and C. Bergenheim, “Evaluation of the V2V channel and diversity potential for platooning trucks,” in *2016 10th European Conference on Antennas and Propagation (EuCAP)*, April 2016, pp. 1–5.
 - [12] V. Jain, S. Lapoehn, T. Frankiewicz, T. Hesse, M. Gharba, S. Gangkhedkar, K. Ganesan, H. Cao, J. Eichinger, A. R. Ali, Y. Zou, and L. Gu, “Prediction based framework for vehicle platooning using vehicular communications,” in *2017 IEEE Vehicular Networking Conference (VNC)*, Nov 2017, pp. 159–166.
 - [13] K. Karlsson, C. Bergenheim, and E. Hedin, “Field measurements of IEEE 802.11 p communication in NLOS environments for a platooning application,” in *2012 IEEE Vehicular Technology Conference (VTC Fall)*, IEEE, 2012, pp. 1–5.
 - [14] A. Bazzi, A. Zanella, B. M. Masini, and G. Pasolini, “A distributed algorithm for virtual traffic lights with ieee 802.11p,” in *2014 European Conference on Networks and Communications (EuCNC)*, 2014, pp. 1–5.
 - [15] R. Stahlmann, M. Möller, A. Brauer, R. German, and D. Eckhoff, “Exploring GLOSA systems in the field: Technical evaluation and results,” *Computer Communications*, vol. 120, pp. 112 – 124, 2018.
 - [16] S. Demmel, A. Lambert, D. Gruyer, A. Rakotonirainy, and E. Monacelli, “Empirical IEEE 802.11 p performance evaluation on test tracks,” in *2012 IEEE Intelligent Vehicles Symposium*. IEEE, 2012, pp. 837–842.
 - [17] T. Mangel, M. Michl, O. Klemp, and H. Hartenstein, “Real-world measurements of non-line-of-sight reception quality for 5.9 GHz IEEE 802.11 p at intersections,” in *International Workshop on Communication Technologies for Vehicles*. Springer, 2011, pp. 189–202.
 - [18] A. Böhm, K. Lidström, M. Jonsson, and T. Larsson, “Evaluating CALM M5-based vehicle-to-vehicle communication in various road settings through field trials,” in *IEEE Local Computer Network Conference*. IEEE, 2010, pp. 613–620.
 - [19] K. Hong, D. Xing, V. Rai, and J. Kenney, “Characterization of DSRC performance as a function of transmit power,” in *Proceedings of the sixth ACM international workshop on Vehicular InterNetworking*. ACM, 2009, pp. 63–68.
 - [20] M. Alam, B. Fernandes, L. Silva, A. Khan, and J. Ferreira, “Implementation and analysis of traffic safety protocols based on ETSI Standard,” in *2015 IEEE Vehicular Networking Conference (VNC)*. IEEE, 2015, pp. 143–150.
 - [21] J. Santa, F. Pereñíguez, A. Moragón, and A. F. Skarmeta, “Experimental evaluation of CAM and DENM messaging services in vehicular communications,” *Transportation Research Part C: Emerging Technologies*, vol. 46, pp. 98–120, 2014.
 - [22] F. A. Teixeira, V. F. e Silva, J. L. Leoni, D. F. Macedo, and J. M. Nogueira, “Vehicular networks using the IEEE 802.11 p standard: An experimental analysis,” *Vehicular Communications*, vol. 1, no. 2, pp. 91–96, 2014.
 - [23] M. Klapez, C. A. Grazia, and M. Casoni, “Application-Level Performance of IEEE 802.11p in Safety-Related V2X Field Trials,” *IEEE Internet of Things Journal*, vol. 7, no. 5, pp. 3850–3860, 2020.
 - [24] M. Klapez, C. A. Grazia, and M. Casoni, “Minimization of IEEE 802.11p Packet Collision Interference through Transmission Time Shifting,” *Journal of Sensor and Actuator Networks*, vol. 9, no. 2, p. 17, Mar 2020. [Online]. Available: <http://dx.doi.org/10.3390/jsan9020017>
 - [25] T. Paulin and S. Bessler, “A Disruption Tolerant Connectivity Service for ITS Applications Using IEEE 802.11p,” in *17th European Wireless 2011 - Sustainable Wireless Technologies*, April 2011, pp. 1–7.
 - [26] A. Paier, R. Tresch, A. Alonso, D. Smely, P. Meckel, Y. Zhou, and N. Czink, “Average Downstream Performance of Measured IEEE 802.11p Infrastructure-to-Vehicle Links,” in *2010 IEEE International Conference on Communications Workshops*, May 2010, pp. 1–5.
 - [27] J. Lin, C. Lin, C. Liang, and B. Chen, “Wireless communication performance based on IEEE 802.11p R2V field trials,” *IEEE Communications Magazine*, vol. 50, no. 5, pp. 184–191, 2012.
 - [28] P. Alexander, D. Haley, and A. Grant, “Cooperative Intelligent Transport Systems: 5.9-GHz Field Trials,” *Proceedings of the IEEE*, vol. 99, no. 7, pp. 1213–1235, 2011.
 - [29] M. Alam, B. Fernandes, L. Silva, A. Khan, and J. Ferreira, “Implementation and analysis of traffic safety protocols based on ETSI Standard,” in *2015 IEEE Vehicular Networking Conf. (VNC)*, 2015, pp. 143–150.
 - [30] “IEEE 802.11p-2010 -IEEE Standard for Information technology- Local and metropolitan area networks- Specific requirements- Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular Environments,” IEEE, 2010.
 - [31] “ASTM E2213-02, Standard Specification for Telecommunications and Information Exchange Between Roadside and Vehicle Systems — 5 GHz Band Dedicated Short Range Communications (DSRC) Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” ASTM International, 2002.
 - [32] “ASTM E2213 - 03 - Standard Specification for Telecommunications and Information Exchange Between Roadside and Vehicle Systems — 5-GHz Band Dedicated Short-Range Communications (DSRC), Medium Access Control (MAC), and Physical Layer (PHY) Specifications,” ASTM International, 2018.
 - [33] “Intelligent Transport Systems (ITS); Vehicular Communications; GeoNetworking; Part 3: Network architecture,” ETSI TS 102 636-3 V1.1.1, 2010-03.
 - [34] “Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service,” ETSI EN 302 637-2 V1.4.1, 2019-04.
 - [35] “Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 3: Specifications of Decentralized Environmental Notification Basic Service,” ETSI EN 302 637-3 V1.3.1, 2019-04.
 - [36] “Intelligent Transport Systems (ITS); Radiocommunications equipment operating in the 5 855 MHz to 5 925 MHz frequency band; Harmonised Standard covering the essential requirements of article 3.2 of Directive 2014/53/EU,” ETSI EN 302 571 V2.1.1, 2017-02.
 - [37] “FCC 03-324,” Federal Communications Commission, Report and Order, 2004-02.
 - [38] R. Lisový, M. Sojka, and Z. Hanzálek, “IEEE 802.11 p Linux Kernel Implementation,” *Tech. Rep.*, 2014.
 - [39] C. Grazia, N. Patriciello, M. Klapez, and M. Casoni, “A cross-comparison between tcp and aqm algorithms: Which is the best couple for congestion control?” 2017, pp. 75–82.
 - [40] D. Saladino, A. Paganelli, and M. Casoni, “A tool for multimedia quality assessment in ns3: Qoe monitor,” *Simulation Modelling Practice and Theory*, vol. 32, pp. 30–41, 2013.