Dynamic multipath connection for low-latency V2X communications

David Palacios, Emil J. Khatib, István Z. Kovács, Beatriz Soret, Isabel de-la-Bandera and Raquel Barco

June 7, 2018

Dynamic multipath connection for low-latency vehicle-to-everything (V2X) communications

David Palacios
Dpto. Ingeniería de Comunicaciones
Universidad de Málaga
29071, Málaga, Spain
dpc@ic.uma.es

Emil J. Khatib
Dpto. Ingeniería de Comunicaciones
Universidad de Málaga
29071, Málaga, Spain
emil@uma.es

István Kovács
Research Center Aalborg
Nokia-Bell Labs
9220, Aalborg, Denmark
istvan.kovacs@nokia-bell-labs.com

Beatriz Soret
Research Center Aalborg
Nokia-Bell Labs
9220, Aalborg, Denmark
beatriz.soret@nokia-bell-labs.com

Isabel de-la-Bandera
Dpto. Ingeniería de Comunicaciones
Universidad de Málaga
29071, Málaga, Spain
ibanderac@ic.uma.es

Raquel Barco
Dpto. Ingeniería de Comunicaciones
Universidad de Málaga
29071, Málaga, Spain
rbarco@uma.es

***Abstract*—The fifth generation (5G) of cellular communications aims at encompassing a wide variety of types of communications, among which are vehicle-to-everything (V2X) communications. These will be characterized by the interchange of both low-latency messages for control and notification purposes, and delay-tolerant messages, devoted to infotainment services. In this paper, a mechanism for dynamic traffic steering over the Uu and PC5 interfaces of V2X communications is proposed to ensure low latency for critical messages, even in high load conditions. As a proof of concept, a simulation has been carried out using long-term evolution (LTE) macrocells and IEEE 802.11n access points to provide the Uu interface and an approximation to the PC5 interface, respectively. The proposed solution has proven to effectively steer data packets regarding their criticality label, providing a delay-optimized link for critical messages and a best-effort behavior for the non-critical (delay tolerant) messages.**

***Index Terms*—Vehicle-to-everything (V2X), low-latency communications (LLC), traffic steering.**

I. Introduction

One of the main features of the fifth generation (5G) of cellular communications is to holistically cover a wide variety of communication types. In particular, big efforts are being put on integrating vehicle-to-everything (V2X) communications into 5G, leading to cellular V2X communications (C-V2X) [1]. These are characterized by holding different traffic profiles: from high-capacity and delay-tolerant traffic, corresponding to infotainment services, to low-throughput and delay-sensitive traffic, corresponding to control and mission-critical-related services.

As a result, the 3GPP (Third Generation Partnership Project) has recently developed its own specification for C-V2X, [2]. In [2], two interfaces are differentiated for the vehicles to communicate with other elements: the Uu interface, being the air interface of the long-term evolution (LTE) network, and the PC5 interface [3], operating in the band for ITS (Intelligent Transportation Systems) applications, 5.9 GHz. Whereas the Uu interface allows the vehicle to communicate with the network base stations (either macro or small cells), the PC5 interface enables sidelinks (SLs), usually deployed between vehicles or between these and roadside units (RSUs). However, as stated in [2], both interfaces are suitable for the provision of the different kinds of vehicular services; namely, vehicle-to-vehicle (V2V), vehicle-to-infrastructure or network (V2I/N) and vehicle-to-pedestrian (V2P) services, not limiting one of these to a particular interface. This fact, together with the differentiated traffic profiles of V2X, makes necessary an intelligent interface selection in each case.

Thus, for an interface to be selected at a given time, different performance metrics should be assessed, depending on the scope to be optimized. In particular, regarding the criticality of the messages, the ones labeled as critical (delay-sensitive) should be sent through the interface showing the minimum delay and maximum reliability at each time. Despite this selection could be done in a fixed manner, the time-varying nature of the network load due to the users' mobility may make one of the interfaces more suitable to carry delay-sensitive messages than the other at a given time.

In this paper, an algorithm for the dynamic selection of the most appropriate interface is proposed regarding the criticality of messages to be interchanged. That is, depending on whether a message is critical (delay-sensitive) or non-critical (delay-tolerant).

The rest of the paper is organized as follows. Section II briefly describes the different scenarios for V2X communication regarding the availability of the interfaces Uu and PC5, according to the specifications of 3GPP. In Section III, the proposed method for traffic steering between these interfaces

This work has been partially funded by Junta de Andalucía (Proyecto de Investigación de Excelencia P12-TIC-2905) and the Spanish Ministry of Economy and Competitiveness (TEC2015-69982-R). Additionally, it has been performed in the framework of the Horizon 2020 project ONE5G (ICT-760809), receiving funds from the European Union. The authors would like to acknowledge the contributions of their colleagues in the project, although the views expressed in this contribution are those of the authors and do not necessarily represent the project.

is explained. In Section IV, a proof of concept is carried out to show the benefits of the proposed method. Finally, the main conclusions are drawn in Section V.

II. V2X COMMUNICATIONS THROUGH THE UU AND PC5 INTERFACES

Depending on the availability of the Uu and PC5 interfaces, 3GPP defines three scenarios for V2X communication [2]:

- Scenario 1: Only the PC5 interface is available. Thus, both the uplink (UL) and downlink (DL) must rely on SLs.
- Scenario 2: Only the Uu interface is available. In this case, every V2X communication must rely on a direct vehicle-to-eNB link.
- Scenario 3: Both interfaces PC5 and Uu are available (see Fig. 1). This scenario is, in turn, broken down into two variants. In scenario 3A (Fig. 1a), the PC5 interface is used for UL by means of a SL to a UE RSU, which eventually interfaces an eNB via the Uu interface. In DL, the eNB directly interfaces a number of UEs using a broadcast mechanism via Uu. In scenario 3B (Fig. 1b), a UE directly transmits to an eNB via the Uu interface in UL, whereas the DL is supported over a UE RSU, communicating with different UEs through PC5 SLs.

The present work focuses on scenario 3. In particular, on the decision on whether to use a SL to access the network (scenario 3A) or use a direct link (scenario 3B) to deliver UL V2X messages.

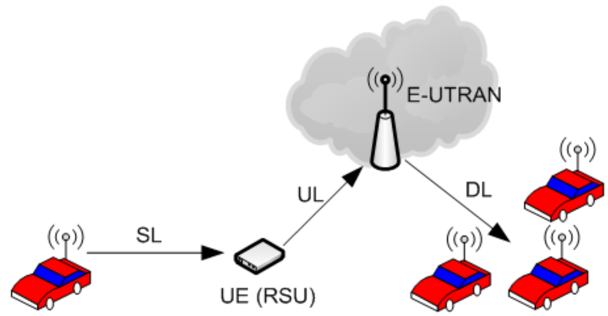

(a) Scenario 3A [2]. In the uplink direction, a UE interfaces a UE RSU via the PC5 link, which eventually interfaces an eNB.

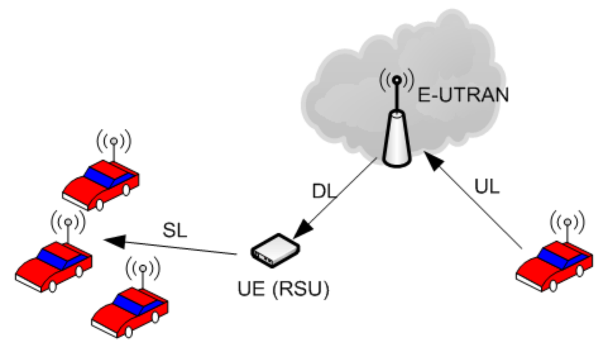

(b) Scenario 3B [2]. In the uplink direction, a UE directly interfaces an eNB via the Uu link.

Fig. 1. Scenarios supporting V2X operation using both Uu and PC5 interfaces, according to 3GPP.

III. PROPOSED METHOD FOR TRAFFIC STEERING

For the following method to be applied, it is required that UEs have both the Uu and the PC5 interfaces active. The second feature of this work, and the one which better represents its novelty, is the assumption that both the E-UTRAN and the UE can make their own routing decisions in terms of which interface to use when delivering downlink and uplink messages, respectively, based on the assessment of high-layer metrics. For example, the UE sending an UL message in Fig. 1 could send such message through either the PC5 or the Uu interface, corresponding to the scenarios 3A and 3B, respectively.

In order to choose among the different links, each of these options will have some metrics associated, like high-layer performance metrics: end-to-end delays, packet-loss rates, etc., and also low-layer metrics, like the measured signal power and quality. The different metrics that are gathered can further be grouped into different sets, which represent different scopes over which optimize the communication. For example: metrics related to a mission-critical behavior, like the delay or the reliability or metrics which inform about energy issues. Within each of these sets of metrics, every option (PC5 and Uu) can be then characterized formulating a cost function according to the corresponding set or scope. In this way, every link gets a score regarding each scope. For example, PC5 SL could be the best option in terms of energy consumption, but it could also be the worse regarding the end-to-end delay.

This grouping into sets, corresponding to high-level scopes, together with the computation of a cost function per link and scope, allows using a simple and flexible decision algorithm either in each UE or the E-UTRAN as the one shown in Algorithm 1. Whenever a new packet arrives from upper layers and needs to be sent, Algorithm 1 is executed. This is a decision process made up of nested conditions, each of which is related to one of the scopes previously defined by setting a certain criteria over them. If the condition related to the first criterion is fulfilled, then, the link with the best score in its cost function according to that scope is chosen. If it is not, the second criterion is evaluated. In case that none of the criteria are fulfilled, a default link is chosen. This type of decision structure enables the proposed solution to be used over a wide range of applications by only setting different conditions over these scopes or by switching these scopes themselves. For example, if the proposed solution is to be used on a mixed-criticality application, a criticality-related condition should be asked first and cost functions based on delay and reliability metrics should be assessed per each link in consequence. In such case, other issues like those related to the energy consumption should remain as second order scopes. Moreover, the generality of Algorithm 1 and its location under application layers makes it specially proper for the integration of information related to the context.

The cost functions have been introduced as one of the

Algorithm 1 Decision tree of the proposed solution

```
New packet from upper layers
    if condition related to criteria 1 is fulfilled then
        Send it through the best link regarding scope 1
    else
        if condition related to criteria 2 is fulfilled then
            Send it through the best link regarding scope 2
        else
            ⋮
            if condition related to criteria n is fulfilled then
                Send it through the best link regarding scope n
            else
                Send it through the default link
```

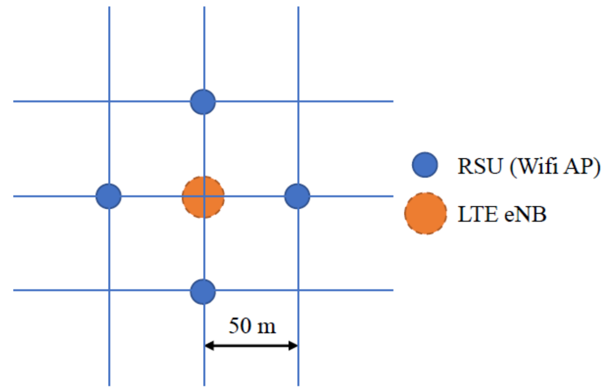

Fig. 2. Simulation scenario.

main enablers of the proposed solution. These functions are periodically computed, with a time period entailing a trade-off between energy consumption and the validity of these costs in their attempt to represent the current status of the networks. A periodical update of these costs, performed as a background process, allows the decision algorithm to take into account possible contingencies in the networks, providing the solution with a dynamic and near real-time behavior, easily assigning the best link in each case without the necessity of further assessment.

IV. Proof of Concept

A. Simulation setup

In this proof of concept, only one scope has been taken into account, leading to a decision algorithm with a single *if*. In this case, the working scope is the criticality of the messages to be sent, in order to assess the validity of the proposal in a mixed-criticality environment, in an uplink V2N/I communication. Thus, the corresponding criterion is whether the message from upper layers holds a high-critical label or not. In such case, and in order to reduce the impact and degradation of the proposed solution in the overlaying LTE network, the PC5 sidelink has been set as the default link, i.e, the link selected for non-critical messages. In this way, the Uu interface is only considered if a message is labeled as critical. In this V2N/I communication, a remote V2X server is assumed to be located beyond the PGW (packet data network gateway) of the LTE network.

Thus, it has been assumed that every UE has two network interfaces: a PC5 interface, in order to communicate with the RSU, and an LTE interface, to communicate with an eNB. In this way, two uplink transmission modes can be identified: the direct, by which the packets are sent from a UE to an eNB through the Uu interface (scenario 3B), and the indirect, by which the data packets are sent to the eNB via an intermediate access point (scenario 3A), leading to a PC5 + Uu link.

The cost functions for both direct and indirect links under the scope of criticality have been addressed by means of a delay metric, the round-trip time (RTT), measured over the UDP protocol. The RTT is included into each cost function as the moving average over the last three E2E RTT measurements in that link. In order to measure the RTT of the messages, an echo application has been enabled on both the remote V2X server and the UEs, so every incoming UDP packet is sent back throughout the same path whenever it arrives at the V2X server.

In order to assess the proposed method, a set of simulations have been carried out using the networks simulator ns-3 [4]. To that end, and as an approximation, the PC5 interface has been simulated using the IEEE 802.11n standard, in the 5 GHz band, using WiFi access points (APs) as the RSUs, and the Uu interface has been simulated using the LTE module of ns3 [5].

To show the benefits of the proposed method, the cumulative distribution function of the E2E RTT is assessed in two cases: having a low UE density (leading to a situation of low traffic) and a high UE density. Besides, two baseline situations have been also simulated, the case in which all the messages (critical or not) are sent through the PC5 interface (baseline experiment 1) and the case in which all the messages are sent through the Uu interface (baseline experiment 2). In all these cases every node sends 256-byte messages. These messages are labeled as critical with a 5 percent of probability and non-critical with a 95 percent of probability. The packet arrival is given by a Poisson process with four messages per second on average.

The simulated scenario consists in a Manhattan-like deployment in a 100 m $\times$ 100 m area. This area is made up of 3 horizontal and 3 vertical 100-m streets, in which a number of nodes have been randomly deployed. In this scenario, the UEs move at a constant speed of 15 km/h and have a 0.5, 0.25 and 0.25 probability of following straight ahead, turning left and turning right at each intersection, respectively. In order to retain the UEs in the simulated area, a wrap-around technique has been used beyond the borders of this area. In this scenario, four RSUs and an eNB have been deployed as it is shown in Fig. 2. The main simulation parameters have been summarized in Table I.

B. Results

The results of the baseline experiments 1 and 2 are shown in Fig. 3 and 4, respectively. Regarding experiment 1 (Fig. 3), it can be observed how the increase in the number of UEs, and

TABLE I
SIMULATION PARAMETERS

Parameter	**Value**
Scenario setup	
Simulation tool	*ns-3*: WiFi and LTE (LENA) modules
Number of UEs	15 (low traffic), 75 (high traffic)
UDP packet size	256 bytes
UDP packet interval: Poisson process $1/\lambda$	0.25 s
Number of WiFi APs	4
UE speed	15 km/h
Simulation time	60 s
PC5 interface	
Standard used	IEEE 802.11n
Band	5 GHz
Channel bandwidth	20 MHz
Operating mode	Infrastructure
Uu interface	
Downlink carrier frequency	945 MHz
Uplink carrier frequency	900 MHz
System bandwidth	20 MHz
Number of eNBs	1
Sectors per eNB	1
Interferences	Other cell interference was not simulated
Shadowing	log-normal, $\sigma = 8$
PGW $\leftrightarrow$ remote host delay	10 ms

therefore, the increase in the traffic load, noticeably impacts the performance of the PC5 link, increasing the 90^{th} percentile of the RTT from 37.4 ms to 529.9 ms. It is worth noting that this increase takes place for both the critical and non-critical messages, since no distinction is made at this point. This increase is mainly due to the usage of a contention-based protocol in the medium access procedure, since the probabilities of finding the medium busy increase quickly with the number of devices. Despite the current PC5 interface does not actually implement a CSMA (carrier sense multiple access) mechanism for medium access (as IEEE 802.11n does), a similar behavior to this may appear when the UEs self-allocate time and frequency resources from a relatively small resource pool. This would likely take place in a dense scenario, such as an urban area where resources may be allocated to other users.

Regarding baseline experiment 2, Fig. 4 shows that the increase in the traffic has a much lower impact in the delay, increasing the 90^{th} percentile of the RTT from 36.8 ms to 43.7 ms. Unlike WiFi, in which, every new packet to be sent needs a contention-based medium access to be performed in advance, in LTE, the contention-based medium access is only performed if frequency and time resources have not been assigned yet to the current device. Once these are allocated, the device holds them during a certain time period, not needing to further perform a contention-based random access procedure and thus, reducing drastically the E2E RTT. In these experiments, the mean time between packet arrivals is 0.25 s, whereas the time window along which the LTE resources are held in a given connection after its activity has ceased is 10 s. In this way, the time delay due to the contention-based random access is only added to the first packet transmitted by a given node, leading to the short RTTs that are shown in Fig. 4. Even though, the effect on the E2E delay of a higher number of devices contending for an access grant can be seen in this figure, as a result of some collisions during the random access procedure. Besides, a higher traffic load at the scheduler in the eNB also contributes to a higher delay. Despite these results, the direct connection through Uu should not be abused, since a congestion can be produced in the access process under certain conditions, impacting not only on the performance of the present proposal, but on the performance of traditional users employing the LTE network as an MBB (mobile broadband) communications system.

The results when the UEs follow the proposed algorithm and both interfaces are available are shown in Fig. 5 for the four combinations of traffic load and message criticality. For non-critical messages (striped green and dash-dotted blue lines) the behavior of the E2E RTT is rather similar to baseline case 1, as all these messages are sent through the indirect link. However, it should be noted that, despite having a large mean E2E RTT in the non-critical high-traffic case, this delay is noticeably lower than the one obtained in the baseline case 1. This is because in the latter, all the messages, critical or not, are forwarded throughout the same link, leading to a stronger contention in the medium access. In Fig. 5, however, only the non-critical messages and those critical messages that find the PC5 interface as a faster option are forwarded through the WiFi AP, leading to a lower number of medium access attempts and thus, to a lower delay due to this contention-based mechanism.

The case with the sidelink being faster than the direct link arises when the mean time between packet arrivals at the WiFi AP ($1/(\lambda \cdot N_{UE})$, where N_{UE} stands for the number of UEs using the sidelink towards the RSU at that time) is similar or higher than the LTE inactivity time: 10 s in this case, and N_{UE} is such, that the delay due to the contention-based access in the WiFi network is shorter that the one in the LTE random access procedure. When the messages are spread enough in time it is quite likely that every new packet to be sent through the direct link would require a random access procedure to be performed in the Uu interface, increasing the E2E delay. On the other hand, and given that the WiFi gateway holds a single connection which gathers messages from many UEs, the net mean time between packet arrivals at the gateway could be significantly lower than the defined inactivity time in LTE. This would make the allocated resources for the gateway-to-eNB connection not to be released, not needing to perform a random access procedure every time a new packet arrives at the gateway, and thus, not adding an extra term to the RTT of the packets through the indirect link.

Besides, critical messages show a more stable behavior; solid orange and dashed red lines, respectively. Fig. 5 shows that, with low traffic, both critical and non-critical messages experiment a similar and low delay profile (being 37.8 ms and

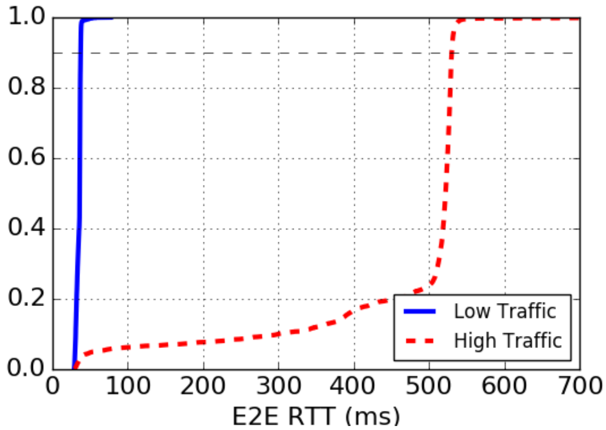

Fig. 3. Empirical CDF of the measured E2E RTT in the baseline experiment 1: always use the PC5 + Uu links.

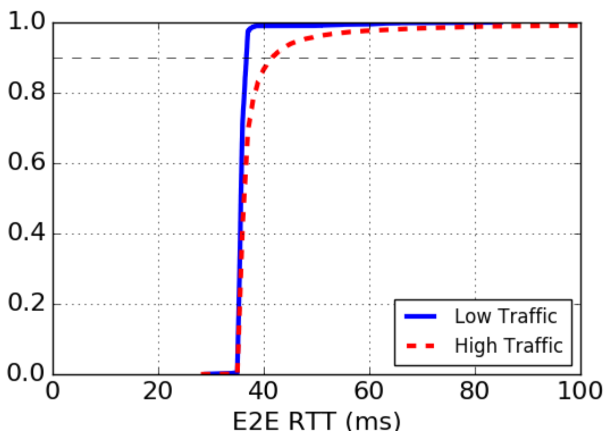

Fig. 4. Empirical CDF of the measured E2E RTT in the baseline experiment 2: always use the Uu link.

38.4 ms their RTT 90^{th} percentile, respectively) and that, in a high traffic scenario (dashed red and dash-dotted blue lines), the proposed method provides an effective differentiation for a mixed-criticality service, providing an optimal delay for critical messages (an RTT 90^{th} percentile of 37 ms) and a delay for the non-critical ones much lower than those of the only-PC5 baseline case, having an RTT 90^{th} percentile of 80.2 ms.

In light of these results, and under these assumptions, all of the use cases described in [6] for LTE-based V2N/I and some of the ones described in [7] for 5G-based V2N/I communications could be addressed in terms of expected E2E latency.

V. Conclusion

In this paper, a method has been proposed in the field of V2X communications to steer messages over the Uu and the PC5 interfaces in order to provide a minimum delay for messages labeled as critical. To that end, cost functions over delay-related metrics are computed on a per-interface basis as

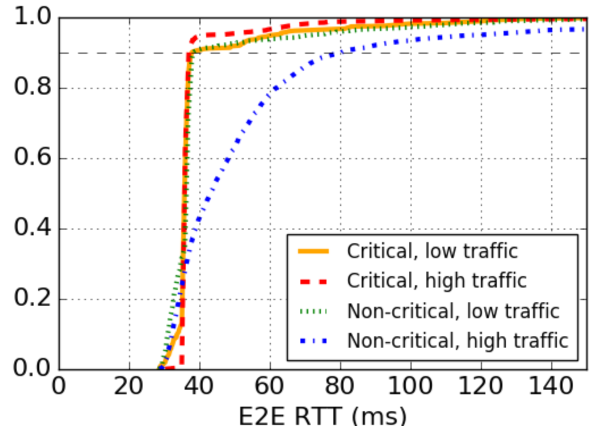

Fig. 5. Empirical CDF of the measured E2E RTT using the proposed solution to decide to use either the direct Uu link or the indirect PC5 + Uu link regarding past RTT measurements from each path.

a means to evaluate its suitability at a given time. Results from simulations show how the proposed method allows critical messages to neglect the effect of high traffic loads leading to low delays, while providing a best-effort behavior for non-critical ones.

References

[1] Qualcomm, “Accelerating C-V2X commercialization,” Qualcomm, Tech. Rep., 2017.

[2] 3GPP, “TR 36.885, Study on LTE-based V2X Services, Rel-14, V14.0.0,” Tech. Rep., 2016.

[3] 3GPP, “TS 23.303, Proximity-based services (ProSe), Rel-15, V15.0.0,” Tech. Rep., 2017.

[4] G. F. Riley and T. R. Henderson, *The ns-3 Network Simulator*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 15–34.

[5] N. Baldo, M. Miozzo, M. Requena-Esteso, and J. Nin-Guerrero, “An open source product-oriented LTE network simulator based on ns-3,” in *Proceedings of the 14th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ser. MSWiM ’11. New York, NY, USA: ACM, 2011, pp. 293–298.

[6] 3GPP, “TR 22.885, Study on LTE support for Vehicle-to-Everything (V2X) services, Rel-14, V14.0.0,” Tech. Rep., 2017.

[7] 3GPP, “TR 22.886, Study on enhancement of 3GPP support for 5G V2X services, Rel-15, V15.1.0,” Tech. Rep., 2017.