Improving the Latency of 5G V2N2V Communications in Multi-MNO Scenarios using MEC Federation

B. Coll-Perales[1], M.C. Lucas-Estañ[1], T. Shimizu[2], J. Gozalvez[1], T. Higuchi[2], S. Avedisov[2], O. Altintas[2], M. Sepulcre[1]
[1]Uwicore laboratory, Universidad Miguel Hernandez de Elche, Elche (Alicante), Spain.
[2]InfoTech Labs, Toyota Motor North America R&D, Mountain View, CA, U.S.A.
{bcoll, m.lucas, j.gozalvez, msepulcre}@umh.es ; {takayuki.shimizu, takamasa.higuchi, sergei.avedisov, onur.altintas}@toyota.com

2022 IEEE 95th Vehicular Technology Conference (VTC2022-Spring) | 978-1-6654-8243-1/22/$31.00 ©2022 IEEE | DOI: 10.1109/VTC2022-Spring54318.2022.9860639

***Abstract*—5G and multi-access edge computing (MEC) are being considered to support V2X services demanding low latency and highly reliable communications using V2N2V (Vehicle-to-Network-to-Vehicles) communications instead of direct or sidelink V2V (Vehicle-to-Vehicle). Guaranteeing V2X service continuity using V2N2V is a challenging task in multi-Mobile Network Operator (MNO) deployments where vehicles are supported by different MNOs. MEC federations have been proposed to address some of these challenges. A MEC federation is a *federated model of MEC systems enabling shared usage of MEC services and applications*. Through MEC federations, vehicles can seamlessly access V2X applications independently of whether they are hosted on their MNO's MEC, or on the MEC of a different (but federated) MNO. This paper presents the first study that analyses the impact of MEC federation on the end-to-end (E2E) latency when supporting V2X services using 5G V2N2V in multi-MNO scenarios. The paper also evaluates the feasibility to support the latency requirements of advanced V2X services in these scenarios, and the benefits introduced by MEC federation. This study considers the V2X-based cooperative lane merge service as a case study.**

Keywords*—*5G, CAVs, connected and automated vehicles, end-to-end latency, MEC federation, V2N2V, V2X.

I. INTRODUCTION

Connected and automated driving services are characterized by stringent V2X reliability and latency requirements that augment with the level of automation. 3GPP, ETSI and 5GAA recommend using 5G networks with MEC deployments for the support of these services [1][2]. This approach is fueled by the possibility of processing data and hosting V2X application servers (AS) at the edge of the network, which reduces the latency and improves the scalability compared to solely relying on cloud computing at centralized data centers [3]. The benefits of integrating 5G and MEC have raised expectations on the possibility to support advanced V2X services via V2N2V (Vehicle-to-Network-to-Vehicle) communications instead of direct or sidelink V2V (Vehicle-to-Vehicle) communications (see Fig. 1).

The 5G system architecture specifications under 3GPP TS 23.501 laid the foundation for the support of V2X services using 5G networks with MECs. Activities are now focused on more realistic (but challenging) scenarios including multi-MNO (Mobile Network Operator) deployments where vehicles are served by different MNOs. These scenarios are critical since vehicles will not be supported by a single MNO, and it is hence necessary to guarantee the V2X service requirements when communicating vehicles via V2N2V are being served by different MNOs. Multi-MNO scenarios entail numerous challenges for the support of V2X services due to the necessary communication and coordination between MNOs to guarantee seamless service provisioning. This can be challenging when MECs hosted by different MNOs are supporting the V2N2V connection due to the inter-MEC communication limitations. The multi-MNO scenarios could also lead to situations where a vehicle is out of its MNO coverage (i.e., in roaming) and its data traffic need to be routed to its home network. The longer distances that V2X data traffic travels from the visited network to the home network would result in a degradation of the V2X service performance.

To guarantee low latency and V2X service continuity in multi-MNO scenarios, ETSI proposes the establishment of MEC federations. A MEC federation is defined as a "*federated model of MEC systems enabling shared usage of MEC services and applications*" [4]. ETSI is defining and standardizing the control plane procedures, signaling and interfaces that will allow users (i.e., vehicles) to seamlessly access applications independently of whether they are hosted on the MEC of those users' MNO, or on the MEC of a different (but federated) MNO. To this aim, inter-MEC communications are considered [4] so that MEC platforms can discover and exchange information with other MEC platforms that may belong to different MEC systems even when they are hosted by different MNO. MEC federation has the potential to address some of the challenges resulting from multi-MNO deployments. However, as highlighted by 5GAA in [2], it is still necessary to analyze and quantify the end-to-end (E2E) performance that MEC federations can achieve when supporting V2X services in multi-MNO deployments, and in particular their latency budget.

In this context, this paper progresses the state-of-the-art by analyzing for the first time the impact of MEC federation on the E2E latency experienced when supporting V2X services using 5G networks with MECs via V2N2V in multi-MNO deployments. The study also analyzes the potential of MEC federations to support the requirements of connected and automated driving services using a cooperative lane merge service as a case study. The study is conducted considering the V2X scenarios identified by ETSI in [4] to illustrate the applicability of the MEC federation concept in multi-MNO deployments.

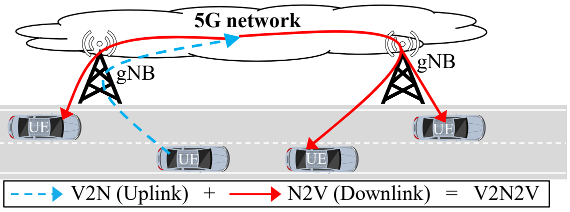

Fig. 1. Vehicle-to-Network-to-Vehicle (V2N2V) communication scenario.

II. MEC FEDERATION

A MEC system consists of the MEC hosts and the management functions necessary to run MEC applications. A MEC host contains a MEC platform and provides computing, storage, and networking resources for the purpose of running MEC applications. The MEC platform offers an environment where the MEC applications can discover, advertise, consume and offer MEC services. MEC developments were initially restricted to intra-MEC system communication within a single MNO's network. Therefore, MEC-based services could only be offered to users connected to the same MNO. Users of different MNOs could not be part of the same MEC service even if they were located nearby.

This work was supported in part by MCIN/AEI/10.13039/501100011033 (grants IJC2018-036862-I, PID2020-115576RB-I00) and by Generalitat Valenciana.

978-1-6654-8243-1/22/$31.00 ©2022 IEEE

Current efforts are focused on expanding the MEC scope and include inter-MEC communication to enable multiple MEC systems to communicate with each other. In this case, MEC platforms and applications should be able to discover other MEC platforms and applications that may belong to different MEC systems and exchange with them the necessary information. Inter-MEC communication is not restricted to MEC systems hosted in the network of a particular single MNO, but also includes MEC systems hosted in the networks of different MNOs. Inter-MEC connections can be based on a direct (private) peering point link between the associated MEC systems and/or their MNOs. Communication between different MEC systems (including those located in different MNOs) has motivated the establishment of MEC federations. MEC federations represent federated models of MEC systems that enable a shared and collaborative usage of MEC resources, services and applications. With MEC federations, MEC systems (providers) can choose to share (some of) their resources with other MEC systems and make them discoverable and accessible to other federation members.

MEC federation is envisioned as an evolution of cloud federation where users are served by the MEC system that can provide the optimal service. This is the case independent of the location of the user and the relation between the user and the MNO's network where the federated MEC system is located. Therefore, MEC federation is a key feature for guaranteeing service continuity when users move between areas covered by different MEC systems and when they move out of their MNO's network. The possibility of sharing resources among federated MEC systems also increases the capabilities of MEC for supporting demanding services at the edge of the network.

III. MEC FEDERATION SCENARIOS

ETSI motivates in [4] the establishment of MEC federations through the following scenarios aiming to improve V2X service continuity and service level in multi-MNO deployments. For illustration purposes, and without loss of generality, we present these scenarios considering that the MEC is deployed at the M1 node of the 5G transport network.

A. MEC federation for edge service availability on visited networks

This scenario considers the case illustrated in Fig. 2 in which a vehicle that is a subscriber of MNO A (home network) is roaming on MNO B (visited network). This can be either national or international roaming. Inter-MEC communications in a MEC federation involving both MNOs can be used to identify the roaming and exchange relevant information. The MECs can also evaluate whether the V2X traffic of the vehicle can be routed to the MEC located at the visited network. This is represented in Fig. 2 with a solid line that connects the roaming vehicle with the V2X AS hosted on the MEC of the visited network. In this case, the MEC federation can guarantee service continuity in multi-MNO scenarios if the MEC located at the visiting network delivers the V2X service to the roaming vehicle with the same service level as if it was delivered by the home network. If MEC federation is not supported, the roaming vehicle remains attached to the MEC located at its home network, and the V2X traffic generated by the vehicle needs to be routed to the vehicle's home network (dashed lines in Fig. 2). This longer path to reach the MEC at the home network can degrade the V2X service level compared to the scenario supporting MEC federation.

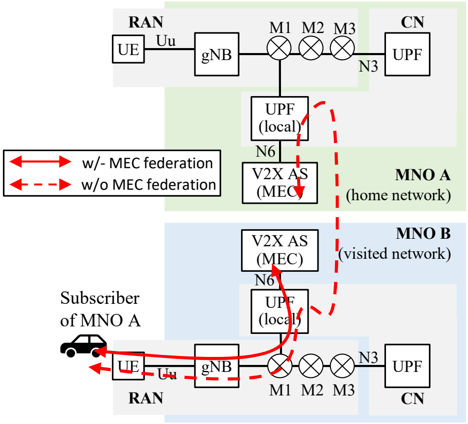

Fig. 2. Edge service availability on visited networks.

B. MEC federation for edge node sharing

The MEC federation can also be used to share the capabilities of MECs among MNOs. This scenario is illustrated in Fig. 3. The figure shows a scenario in which two vehicles use the same V2X AS. This V2X AS is only available on the MEC of the MNO A. The established MEC federation between the MNOs results in that the traffic of the vehicle connected to the MNO B is routed and processed at the MEC system of the MNO A. Note that the vehicle's V2X traffic accesses the MEC of MNO A while it remains connected to the network of MNO B. In the example illustrated in Fig. 3, the MEC federation solves possible limitations of the MEC of one of the MNOs to support a particular V2X service. There are other examples in which it could be desirable that the V2X traffic from different vehicles (also served by different MNOs) is processed at the same MEC to reduce the latency and to process the traffic at a single location to provide a unified response. This would be for example the case when different vehicles (possibly supported by different MNOs) are coordinating their maneuvers for a lane merge [4]. In this case, the V2X traffic of the vehicles would be routed to the MEC of one of the MNOs where the maneuver coordination can be more efficiently planned compared to the scenario where two MECs at different MNOs need to exchange information (with the consequent latencies) to coordinate and execute the maneuver. Without MEC federation, the maneuver coordination can only be processed at a common V2X AS if the V2X traffic from the vehicles involved in the maneuver is routed to a cloud server on the Internet (see Fig. 3); this also entails additional latencies. We should note that the solutions offered by MEC federation in this scenario would also provide MNOs the possibility to coordinate their MEC deployments, and avoid each MNO installing the same V2X services in every location.

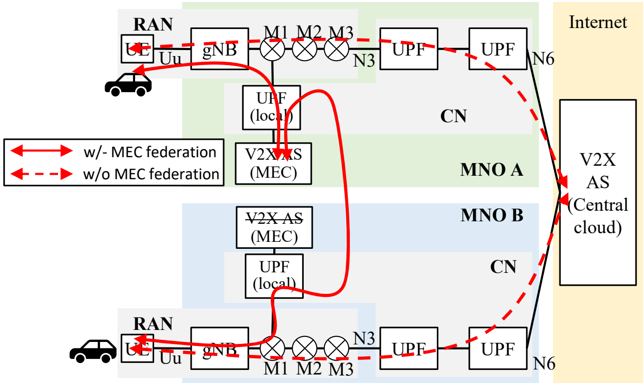

Fig. 3. Edge node sharing.

C. *MEC federation for connecting different services*

This scenario addresses the case in which different MECs need to collaborate to support a V2X service. This could happen, for example, if a MEC does not have sufficient resources to support the requested V2X service or a given set of vehicles. Using the MEC federation, a MEC could then use the resources of another MEC (at the same or different MNO network) to jointly support the V2X service (Fig. 4). If the MEC federation is not in place, the MEC with insufficient resources may have to stop supporting the V2X service or transfer it to the cloud at the cost of latency degradation.

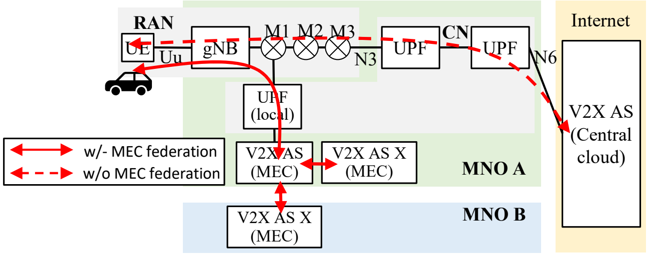

Fig. 4. Connecting different services.

IV. End-to-End Latency

A. *5G end-to-end (E2E) latency model*

This study analyzes the impact of MEC federation on the E2E latency of 5G networks supporting V2X services using V2N2V communications. To this aim, we utilize and adapt the E2E latency l_{E2E} model introduced by the authors in [3]. Fig. 5 depicts the latency components of the model for a centralized 5G network deployment. The E2E latency model accounts for the latency experienced at the radio network l_{radio}, the transport (l_{TN}) and core (l_{CN}) networks, as well as the latency generated by Internet connections, the communication link between the Core Network's (CN's) UPF node and the V2X AS ($l_{UPF\text{-}AS}$), and the processing at the V2X AS (l_{AS}). The E2E latency model also accounts for the latency introduced in the peering point between the MNOs (l_{pp}) for the scenarios when the communicating vehicles are served by different MNOs. Fig. 5 depicts the case of a centralized network deployment. However, l_{E2E} models have been derived in [3] for the most common 5G network deployments following [2]. This includes a centralized deployment in which the V2X AS is located at the cloud, and network deployments with MECs located at the CN, transport network (TN) or gNB. For all deployments, we consider that the radio access network (RAN) and the CN are interconnected using the hierarchical TN proposed by ITU-T in [5] that integrates 3 multiplexing nodes M1, M2 and M3.

The term l_{radio} accounts for the latency experienced between the vehicle (UE) and the gNB at the Uu radio interface. l_{radio} is derived by the authors in [6] considering different 5G New Radio (NR) configurations (e.g., numerology, retransmission and scheduling schemes), system parameters (e.g., bandwidth), characteristics of the data traffic (e.g., periodic or aperiodic), and density and distribution of vehicles in the scenario.

The latencies l_{TN} and l_{CN} are computed as the sum of the propagation and transit delays over the TN and CN. The propagation delay represents the time that packets need to travel through the links that interconnect the nodes of the TN or CN. It depends on the total distance that packets travel through the TN or CN, and varies with the specific 5G network deployment. In the centralized deployment, the packets travel through the entire TN and CN. The distance that packets travel in deployments with MECs depends on whether the MEC is located at the gNB, TN or CN. In deployments with MECs, the CN's UPF node is collocated with the MEC hosting the V2X AS. Therefore, the CN distance is negligible, and so is the propagation delay. The transit delay accounts for the time that packets spend at TN and CN nodes, namely, the time needed to receive, process (including dequeuing) and transmit the packets. The transit delay is computed using queueing theory and depends on the number of nodes that packets pass through, the V2X network traffic load, and the link capacities allocated to support the V2X traffic. The transit latency depends on the specific 5G network deployment, in particular the configuration of TN and CN nodes used.

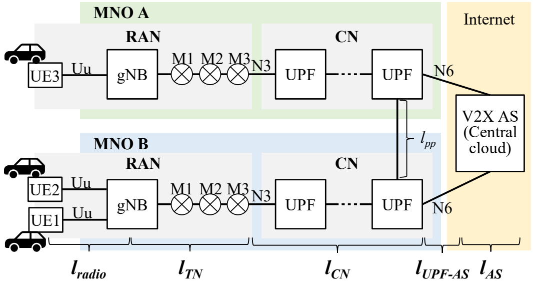

Fig. 5. Latency components of the 5G E2E latency model [3].

The Internet latency $l_{UPF\text{-}AS}$ only applies to the centralized network deployment. It is computed based on empirical measurements reported in [7] of the round-trip time observed between source-target Internet nodes located at the same country; the most likely scenario for 5G-based communications between neighboring vehicles.

The latency l_{AS} introduced by the processing of the V2X packets at the AS is computed considering that the V2X AS only forwards the received packets as in [8]. The processing power of the AS is dimensioned to avoid backlogging of packets at the AS queue.

The term l_{pp} represents the latency introduced in the peering point between MNOs. It is modeled based on the empirical study reported in [9]. This study distinguishes between local (or private) and remote (or public) peering points established between MNOs.

B. *Latency analysis of MEC federation scenarios*

The E2E latency model can be extended to account for the impact of MEC federation in multi-MNO deployments for the three scenarios identified in [4] and described in Section III.

1) *MEC federation for service availability on visited networks*

When MEC federation is used in this scenario, the V2X traffic is routed towards the MEC on the visited network that hosts the V2X AS (solid line in Fig. 2). In this case, the E2E latency can be computed as:

$$l_{E2E}=l_{radio}+l_{TN\_MEC}+l_{CN\_MEC}+l_{AS} \tag{1}$$

where, l_{TN_MEC} and l_{CN_MEC} represent the TN and CN latency for any of the possible 5G network deployments with MECs. If the MEC federation does not exist, the V2X traffic needs to be routed to the vehicle's home network through the peering point between the MNOs (dashed line in Fig. 2). In this case, the E2E latency can be computed as:

$$l_{E2E}=l_{radio}+l_{TN\_MEC}+l_{CN\_MEC}+l_{AS}+l_{pp} \tag{2}$$

2) *MEC federation for edge node sharing*

In this scenario, vehicles can use the MEC federation to access the V2X AS that is hosted in the MEC of a different MNO (solid line in Fig. 3). In this case, the E2E latency can be computed using (2). When MEC federation does not exist, if an MNO does not have a MEC or the MEC does not host a particular V2X service, the vehicles supported by this MNO might have to rely on a central cloud hosting the V2X AS (dashed line in Fig. 3). These vehicles will then experience the following E2E latency:

$$l_{E2E}=l_{radio}+l_{TN\_cent}+l_{CN\_cent}+l_{UPF\text{-}AS}+l_{AS} \quad (3)$$

where l_{TN_cent} and l_{CN_cent} refer to the TN and CN latency of the centralized network deployment.

3) *MEC federation for connecting different services*

This scenario covers the case in which a V2X service requires the collaboration between different MECs. The MECs might be hosted in the same or in different networks. For these cases, the E2E latency is equal to:

$$l_{E2E}=l_{radio}+l_{TN\_MEC}+l_{CN\_MEC}+(l_{AS}+l_{ASx})+l_{MEC\text{-}MEC} \quad (4)$$

where $l_{MEC\text{-}MEC}$ is the latency experienced in the link between the MECs, and l_{ASx} is the latency introduced by the V2X AS X (see Fig. 4). We consider the lowest-latency scenario to compute both values. In particular, we consider $l_{MEC\text{-}MEC}$ is equal to the latency experienced in a local peering link, i.e., $l_{MEC\text{-}MEC}=l_{pp}^{local}$, and that the AS X running in the second MEC just forwards the received packets, i.e., $l_{ASx}=l_{AS}$. When MEC federation does not exist, and the V2X service needs to be transferred to a central cloud because a local MEC of a MNO does not support a particular V2X service or does not have the resources to support a particular vehicle, the E2E latency is computed as in (3).

V. EVALUATION

This section evaluates the impact and benefits of establishing MEC federations to support V2X services using 5G V2N2V communications in multi-MNO deployments. We calculate the E2E latency with and without implementing MEC federations in the three scenarios previously described.

A. *Scenario*

We consider the network topology recommended by ITU in [5], i.e., a hierarchical transport network architecture that is made of 3 levels of multiplexing nodes (M1, M2 and M3). Each M1 node multiplexes traffic from 6 gNBs, each M2 node from 24 M1 nodes, and each M3 node from 12 M2 nodes. The M3 nodes serve as gateways to the 5G core network. The network is configured with distances of 3, 12, 60 and 200 km for the links gNB-M1, M1-M2, M2-M3 and M3-UPF (that connects to Internet), respectively. The network is configured with link capacities of 10, 300, 6000 and 6000 Gb/s for the links gNB-M1, M1-M2, M2-M3 and M3-UPF, respectively.

This study does not focus on a specific V2X service. Instead, we model that V2X packets generated by vehicles to arrive to each gNB at different rates λ_{gNB}^{UL} ranging from 1040 pkts/s to 41600 pkts/s. These rates correspond to different traffic densities and transmission periods [3] and allow us to analyze the impact of variable network loads. We dimension the network and determine the fraction (α) of the link capacities that should be allocated to support the V2X traffic and avoid that packets backlog at the TN and CN nodes following the methodology in [3] and network planning practices described [10] for the highest network load (i.e., λ_{gNB}^{UL} = 41600 pkts/s). The dimensioning results in an α equal to 2.12% for the network deployments where the MEC is located at the gNB or M1, and equal to 6.09% when the MEC is located at the CN as well as for the centralized network deployment. We consider the same values of α for all values of λ_{gNB}^{UL} under evaluation.

B. *Latency components*

Table I reports the round-trip latency for each one of the links intervening in the three scenarios under evaluation. Results are depicted for the Centralized network deployment and the network deployments with the MEC located at gNB (MEC@gNB), M1 (MEC@M1) or CN (MEC@CN). Average[1] (Table I.a) and 99.9th percentile latency values (Table I.b) are reported in Table I. The 99.9th percentile is chosen since it is the most common latency requirement for advanced V2X use cases analyzed by 5GAA in [11]. A range of latency values are reported when applicable for the lowest and highest network traffic loads (i.e., λ_{gNB}^{UL} equal to 1040 pkts/s and 41600 pkts/s).

TABLE I. LATENCY IN MS FOR THE DIFFERENT E2E LINK COMPONENTS

a) Average

Link	MEC@gNB	MEC@M1	MEC@CN	Centralized
l_{radio}	1.5 – 14.23			
l_{TN}	0.41 – 0.42	0.85 – 0.88	2.36 – 2.36	2.36 – 2.36
l_{CN}	< 0.001	< 0.001	< 0.01	2.005
$l_{UPF\text{-}AS}$	0	0	0	10.3
l_{pp}	0.306 (local) or 13.001 (remote)			
l_{AS}	0.5			

b) 99.9th percentile

Link	MEC@gNB	MEC@M1	MEC@CN	Centralized
l_{radio}	2.008 – 28.557			
l_{TN}	0.48 – 0.56	0.99 – 1.15	2.41 – 2.42	2.41 – 2.42
l_{CN}	< 0.001	< 0.001	< 0.01	2.005
$l_{UPF\text{-}AS}$	0	0	0	42.8
l_{pp}	1.43 (local) or 99.21 (remote)			
l_{AS}	0.7			

The radio network latency l_{radio} is calculated following the results in [6] that considers a common FDD (Frequency Division Duplex) reference configuration with SCS (Sub-Carrier Spacing) of 30 kHz and a cell bandwidth of 20 MHz. l_{TN} and l_{CN} increase with the distance and number of TN/CN nodes that V2X packets need to travel to reach the V2X AS. The Internet latency $l_{UPF\text{-}AS}$ only intervenes in the centralized network deployment, and its values are calculated using the empirical measurements reported in [7]. The peering point latency l_{pp} between MNOs differentiates between local (private) and remote (public) peering points, and is calculated based on the empirical study in [9]. Finally, the latency l_{AS} introduced by the AS is calculated in [8].

C. *Impact of MEC federation on E2E latency*

1) *MEC federation for service availability on visited networks*

The E2E latency experienced when a vehicle is roaming depends on whether there is an established MEC federation between the home and visiting MNOs (Fig. 2). If an established MEC federation exists, the V2X packets experience the same E2E latency as if the vehicle was served in its home network (Table II.a). If not, the V2X traffic of the vehicle that is roaming must be routed to the home MNO network through a peering point. This can significantly increase the latency if MNOs are interconnected through a remote peering point (Table II.c) compared to a MEC

---

[1] As it has been demonstrated in [3], focusing on the average performance could provide misleading conclusions about the capacity of certain deployments and MEC federation scenarios to support V2X services.

federation scenario (Table II.a) or if a local peering point exists (Table II.b).

TABLE II. E2E LATENCY IN MS FOR V2X SERVICE AVAILABILITY ON VISITED NETWORK

	MEC@gNB	**MEC@M1**	**MEC@CN**
	a) w/- MEC federation		
$\overline{l_{E2E}}$	2.41 – 15.15	2.85 – 15.61	4.36 – 17.09
99.9th pctl (l_{E2E})	3.19 – 29.82	3.70 – 30.41	5.13 – 31.69
	b) w/o MEC federation & local peering point		
$\overline{l_{E2E}}$	2.72 – 15.46	3.16 – 15.92	4.67 – 17.40
99.9th pctl (l_{E2E})	4.62 – 31.25	5.13 – 31.84	6.55 – 33.11
	c) w/o MEC federation & remote peering point		
$\overline{l_{E2E}}$	15.41 – 28.15	15.85 – 28.61	17.36 – 30.09
99.9th pctl (l_{E2E})	102.40 - 129.03	102.91 - 129.62	104.33 - 130.9

2) *MEC federation for edge node sharing*

When MEC federation exists in this scenario, vehicles can access the same V2X AS independently of whether it is hosted by their MNO or by a different MNO. The E2E latency experienced when a vehicle connects to the MEC in a different MNO network through a peering point is then equal to the values reported in Table II.b. When MEC federation does not exist, the V2X traffic of the vehicle is processed on the cloud. The larger distances that V2X packets travel to reach the cloud, and the latency introduced by Internet connections, result in higher E2E latency values (Table III) compared with the case where the V2X AS is running on a MEC.

TABLE III. E2E LATENCY IN MS WHEN THE APP IS AT A CENTRAL CLOUD

	$\overline{l_{E2E}}$	99.9th pctl (l_{E2E})
Centralized	16.66 – 29.39	49.92 – 76.48

3) *MEC federation for connecting different services*

The E2E latency experienced when MECs collaborate to support a V2X service through an established MEC federation is reported in Table IV. This collaboration reduces the latency compared to the case where MEC federation does not exist and the V2X service needs to be supported at a cloud for those vehicles supported by an MNO that does not deploy MECs or does not support a particular service at a MEC (Table III).

TABLE IV. E2E LATENCY IN MS FOR CONNECTING DIFFERENTE SERVICES

	MEC@gNB	**MEC@M1**	**MEC@CN**
$\overline{l_{E2E}}$	3.22 – 15.96	3.66 – 16.42	5.17 – 17.90
99.9th pctl (l_{E2E})	5.32 – 31.95	5.83 – 32.54	7.25 – 33.81

4) *Discussion*

The E2E latency of 5G V2N2V communications reduce in the considered scenarios when MEC federation is established. Higher gains are obtained under low network loads and when the MEC is located closer to the edge of the network. For example, the results reported in Tables II.a and II.c show that the average E2E latency $\overline{l_{E2E}}$ reduces by 84.3% and 46.2% under the lowest (λ_{gNB}^{UL}=1040 pkts/s) and highest (λ_{gNB}^{UL}=41600 pkts/s) network traffic loads when MEC federation is established in the MEC@gNB deployment, and by 74.9% and 43.2%, respectively, in the MEC@CN deployment. Similar benefits are also obtained in the other two MEC federation scenarios. The obtained results also show that the MEC federation gains increase when the more stringent latency requirement is evaluated. For example, Tables II.a and II.c show that the 99.9th percentile of the E2E latency reduces by 96.9% and 76.9% under the lowest and highest network loads, respectively, when MEC federation is established in the MEC@gNB scenario. This is critical to support advanced V2X services like the cooperative lane merge that requires a 20-ms E2E latency and a 99.9% reliability for the V2X messages exchanged between two vehicles coordinating a maneuver [11]. Supporting these requirements require that 99.9% of the transmitted packets are received in less than 20 ms. The obtained results show that these requirements can be fulfilled under low to moderate network loads when MEC federation is established. When a MEC federation is not established between MNOs, the latency requirements of the cooperative lane merge use case are only fulfilled if the V2X AS is hosted in a MEC and a local (or private) peering point is established between the MNOs (Table II.b). Other scenarios require hosting the V2X AS on the cloud when MEC federation is not established. In these cases, the latency introduced by the Internet connection does not allow the support of the strict V2X latency requirements of the cooperative lane merge scenario (see Table III).

VI. CONCLUSIONS

5G networks with MEC deployments have the potential to support V2X services. Multi-MNO deployments requires new solutions for guaranteeing service continuity and reduce the latency. One of these solutions is establishing MEC federations between MNOs so that vehicles can seamlessly access applications running on a MEC system, independently of whether the application is running on their MNO's MEC or on the MEC of a different (but federated) MNO. ETSI has presented in [4] several V2X scenarios envisioned for the establishment of MEC federations. Our study progressed the state-of-the art through the evaluation of the E2E latency benefits that MEC federations could bring in the envisioned scenarios for different 5G network deployments. The conducted analysis has shown that MEC federations can reduce the E2E latency under all network deployments and independently of the network load. Higher gains are observed when evaluating the more stringent latency requirements (99.9th percentile values) that are key to support safety-critical services. The latency benefits achieved through MEC federation result in a better capacity to support advanced V2X services with stringent requirements such as cooperative lane merging. Supporting such services is much more challenging when relying on central clouds, and without MEC federation would require high-cost local peering points among MNOs.

REFERENCES

[1] N. Sprecher, et al., "Harmonizing Standards for Edge Computing: A synergized architecture leveraging ETSI ISG MEC and 3GPP specifications," ETSI White Paper #36, July 2020.

[2] 5GAA WP2, "MEC for Automotive in Multi-Operator Scenarios", Tech. Report, March 2021.

[3] B. Coll-Perales, et al., "End-to-End V2X Latency Modeling and Analysis in 5G Networks", pre-print arXiv:2201.06082, Dec. 2021.

[4] ETSI, "MEC; Study on Inter-MEC systems and MEC-Cloud systems coordination", GR MEC 035 V3.1.1, June, 2021.

[5] ITU-T, "Consideration on 5G transport network reference architecture and bandwidth requirements", Study Group 15, #0462, Feb. 2018.

[6] M.C. Lucas-Estañ, et al., "An Analytical Latency Model and Evaluation of the Capacity of 5G NR to support V2X Services using V2N2V Communications", *accepted IEEE Trans. Veh. Technol.*, pre-print arXiv:2201.06083, Dec. 2021.

[7] M. Candela, et al., "Impact of the COVID-19 pandemic on the Internet latency: A large-scale study", *Computer Networks, vol. 182,* Dec. 2020.

[8] M. Emara, et al., "MEC-Assisted End-to-End Latency Evaluations for C-V2X Communications", Proc. EuCNC, Slovenia, Jun. 2018.

[9] George Nomikos, et al., "O Peer, Where Art Thou? Uncovering Remote Peering Interconnections at IXPs", *Proc. ACM IMC, pp. 265-278*, Oct. 2018.

[10] Cisco, "Best Practices in Core Network Capacity Planning", WhitePaper, Sept. 2020.

[11] 5GAA WP1, "C-V2X Use Cases Volume II: Examples and Service Level Requirements", White Paper, Oct. 2020.