Evaluation of 5G-NR V2N connectivity in a centralized cooperative lane change scenario

Federico Poli, Ngoc-Lam Dinh, Valerian Mannoni, Benoit Denis

Evaluation of 5G-NR V2N Connectivity in a Centralized Cooperative Lane Change Scenario

Federico Poli, Ngoc-Lam Dinh, Valerian Mannoni, and Benoît Denis

CEA-Leti, Université Grenoble Alpes, F-38000 Grenoble, France

e-mails: {1st name.family name}@cea.fr

Abstract—By means of system-level simulations, we analyze in this paper the performance of Vehicle-to-Network (V2N) connectivity based on the 5th Generation - New Radio (5G-NR) as a support to Cooperative, Connected and Automated Mobility (CCAM), in light of both network and Multi-access Edge Computing (MEC) deployments. Focusing on a canonical centralized Cooperative Lane Change (CLC) use case that involves three vehicles in a cross-border highway environment, we assess the link reliability and the End-to-End (E2E) latency of all the messages involved in the CLC negotiation phase (from/to inter-connected MECs hosting the centralized maneuvering application), while assuming different deployment configurations and the coexistence with a second demanding vehicular service running over the same radio resources. On this occasion, we illustrate possible benefits from Bandwidth Partitioning (BWP) on Uplink (UL) latency, as well as from an hypothetically tight cooperation between Mobile Network Operators (MNOs) on reliability and continuity, leveraging low-latency inter-MEC transactions and seamless cross-border handover capabilities.

I. Introduction

Vehicle-to-Everything (V2X) connectivity is seen as a key enabler to guarantee the seamless continuity of critical CCAM services along continental road corridors and beyond, allow higher and higher levels of driving automation (typically, up to levels 4 or 5). Related stakes concern the coexistence - and even ideally, the combination- of suitable radio access technologies (e.g., C-V2X sidelink a.k.a. PC5-Mode 4 or the cellular 5G-NR, covering respectively short-range and long-range communication needs), the definition of both network and overall system architectures and eventually, the deployment of new supportive elements of infrastructure [1].

Most of the research contributions reported in the recent literature address the previous problems from two rather decorrelated perspectives. On the one hand, some studies focus on the performance assessment of radio technologies on their own, for instance C-V2X sidelink based Vehicle-to-Vehicle (V2V) communications, most often at link level and regardless of any concrete application (e.g., [2]–[4]), or V2I communications with respect to Road Side Units (RSUs) (e.g., [5]) under various road infrastructure deployment assumptions (e.g., [6]). On the other hand, a second major research direction concerns the higher-level orchestration of networks, services and/or of MEC resources [7]–[9]. These analyses are usually intended in a global and generic 5G ecosystem, with only limited considerations with respect to the underlying radio access performance.

In contrast to the previous contributions, we herein propose a simulation-based study that aims at analyzing the performance of 5G-NR in a representative CCAM context, in light of both network and MEC deployments. More particularly, as a canonical example, we focus on a centralized CLC use case in a cross-border highway environment, assuming the support of Vehicle-to-Network (V2N) connectivity with respect to gNodeBs, as well as inter-connected MECs hosting the centralized maneuvering application. These simulations assess primarily link reliability and E2E latency performance indicators at the radio access level, while relying on simpler abstract assumptions for higher-level segments of the system architecture. Given that the CLC service is supposed to coexist with another demanding vehicular service involving a variable amount of active users, we benchmark the performances in both a baseline single-gNodeB scenario and after gradually increasing the number of gNodeBs. Finally, we illustrate the effect of 5G bandwidth partitioning, as a physical slicing means, onto the UL radio access latency, which is also shown here to be the dominating factor in the overall latency budget.

II. Use Case and Messages Sequence

We consider the CLC use case illustrated in Fig. 1 (also referred to a *CLC as planned* in [10]). Accordingly, one vehicle (A) intends to insert from the overtaking lane onto the first lane where two other vehicles (B and C) are already running. These side vehicles must hence create a gap to allow the initiator (A) to insert between them, depending on the three respective positions. In a 5G-based centralized CLC, a manoeuvring service hosted in a MEC system first interprets the intention of vehicle A and therefore advises vehicle C to increase its headway from vehicle B, before informing vehicle A that the intended lane change is possible. Vehicles typically communicate with the centralized manoeuvring service via a Uu interface relying on 5G-NR. Recommendations from the manoeuvring service are also forwarded through a Uu interface to the respective engine control units (ECUs) of all relevant vehicles so that they can act upon. The time elapsed between the initiator's request (i.e., Vehicle A's updated status) and the actual lane merge realization is hereafter referred to as the *CLC negotiation* phase. In case of automated driving, the vehicles can execute the necessary actions directly. In manual driving mode, the ECU alerts the driver and provides recommendations for the necessary actions.
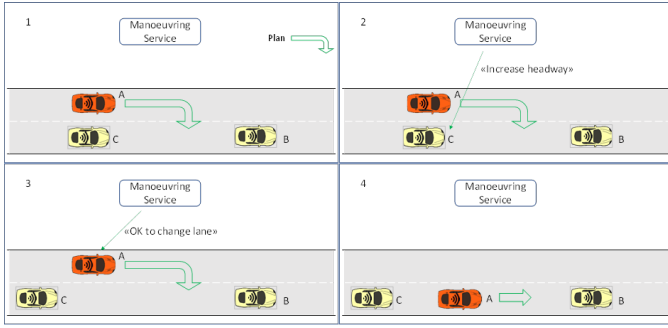
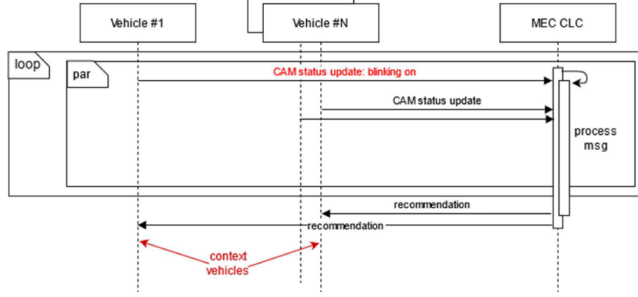Fig. 1: Illustration of a 3-phase V2X-assisted CLC use case [10].

Fig. 2: Typical message sequence in a centralized CLC [12].

In a specific cross-border context, the CLC procedure can be initiated and completed before, while or after crossing the border. In this case, vehicles geographically close to each other and participating to the same manoeuvre procedure may be connected to different MNOs, and hence, to different MEC instances. Hence, we assume herein that low-latency interfaces should be available in a reasonably near future to support information exchanges between services running on neighbouring MECs [11].

In Fig. 2, we show the corresponding messages sequence chart [12], where regular and updated status messages rely on Cooperative Awareness Messages (CAMs), whereas recommendations (incl. both *slowing down* and safe-to-merge messages, resp. to vehicles A and C) rely on Decentralized Environmental Notification Messages (DENMs). Further details of all MEC-internal information flow, which is out of scope of this paper, are available in [11].

III. SIMULATION FRAMEWORK

The system-level simulation flow originally developed to evaluate PC5-based V2V/Vehicle-to-Infrastructure (V2I) [5], [13] and LTE-based V2N [14] connectivity in cooperative manoeuvring and data fusion contexts, has been expanded for the evaluation of 5G-NR V2N connectivity. Accordingly, mobility traces (incl. time-stamped vehicles' 2D coordinates, 2D speed and absolute heading) are first generated offline using the Simulation of Urban Mobility (SUMO) tool [15], under various road densities based on real traffic data [5]. The previous mobility traces are then used for the generation of time-stamped 5G-NR connectivity traces via ns-3 simulations, accounting for the success/failure and latency of CLC transactions at both link and system levels. Mobility and connectivity simulations are deliberately decoupled here, since

we are mostly interested by the information flow prior to CLC execution.

A. Environment, Mobility and Technology Penetration

The tested environment is the same as that considered in [5] and [14]. It is composed of a 4.5 km-long portion of the A22 highway located at the border between Italy and Austria. Based on real traffic measurement data collected in 2019 at this border, macroscopic mobility parameters have been derived to configure our SUMO simulations. These measurements reveal an upper (resp. lower) limit on the traffic flow in each driving direction at approximately 2500 (resp. 750) vehicles per hour, corresponding to an average traffic density of 83 (resp. 28) vehicles per kilometre. Among the simulated flow of vehicles, we also consider a variable ratio of active connected vehicles (i.e., spanning from 1% up to 25%), which somehow reflects the technology penetration in very first approximation. Besides the vehicles involved in the CLC under test, all these other connected vehicles are supposed to transmit their own packets (e.g., regular CAMs) and hence contribute to load the network.

B. 5G-NR Radio Parameters

In our ns-3 simulations, the frequency domain granularity is the Resource Block (RB) and the time domain granularity corresponds to one Orthogonal Frequency Division Multiple (OFDM) symbol, while 14 symbols are assumed per slot. The numerology is systematically set to 2, which offers a reasonable trade-off between low latency and link robustness[1], given the imposed CLC requirements. The addressable Modulation and Coding Schemes (MCSs), which span from 5 to 25, are automatically and contextually optimized in each simulation realization on a case-by-case basis. Time Division Duplex (TDD) slots are dedicated either to Downlink (DL), UL, or both. We use a default pattern of 10 repeated slots (all flexible), each of them containing both UL and DL OFDM symbols, as well as a guard band of 2 symbols. The actual number of UL or DL OFDM symbols per slot is not pre-defined, even though higher priority is given to UL over DL. Regarding propagation modelling, a synthetic propagation channel suitable to motorway environments is generated with both large scale (path loss and shadowing) and small scale (fast fading) statistics according to the standardized models described in [16], regardless of deterministic in-site propagation effects.

The implemented Medium Access Control (MAC) supports both Orthogonal Frequency Division Multiple Access (OFDMA) and Time Division Multiple Access (TDMA) in UL and DL, where a single User Equipment (UE) is scheduled per Transmission Time Interval (TTI) in TDMA while multiple UEs are scheduled per TTI in OFDMA, using all symbols but different Physical Resource Blocks (PRBs). A mixture of TDMA and OFDMA is also feasible. In terms of

---

[1]Among other suitable configurations, a numerology of 1 was for instance considered in 5G-CARMEN field trials (e.g., for a subcarrier spacing of 30 kHz in the n78 frequency band), thus relaxing synchronization constraints for implementation practicability.

TABLE I: Main 5G-NR radio parameters used in simulations

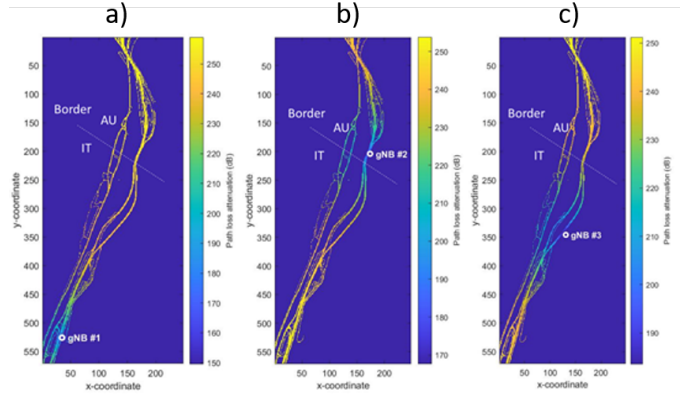| Bandwidth	100 MHz
Channel access scheme	TDMA
Scheduler policy	Round Robin
Numerology	2
Tx power at gNB	36 dBm
Tx power at UE	26 dBm
gNodeB antenna	8x16 Planar Linear Array (PLA)
UE antenna	4x8 PLA with BF
Channel model	RMa (Rural Macrocell)
MCS	from 5 to 25
Duplexing scheme	TDD with flexible pattern
BLER threshold	0.0001
PHY packet payload	100 bytes (DENM) / 500 bytes (CAM)
Regular CAM Tx frequency	20 Hz

Fig. 3: Locations of gNodeB sites considered for the evaluation of the centralized CLC scenario under gradual network deployment at the Brenner Pass, along with their respective RMa path loss maps..

scheduling, we have mainly considered a basic Round Robin scheme for simplicity to derive nominal performances as a baseline, keeping in mind that future enhancements of this scheduler would lead to substantial gains in terms of latency. Finally, Automatic Repeat Request (HARQ) is also activated with systematic prioritized data retransmission. The maximum number of retransmissions is set to 4, again as a trade-off between latency and reliability.

Table I summarizes the main radio parameters above.

C. Network Deployment and Handover Assessment

The deployment of a gradually increasing number of 5G-NR gNodeBs, all operating at 3.5 GHz, is considered along the 4.5 km of motorway around the Brenner pass. In our experiments, without loss of generality, the CLC vehicles are arbitrarily supposed to drive north from Italy to Austria, thus being primarily associated with the Italian MNO. The idea is to identify the minimum amount of infrastructure density to make sure that reliability requirements (at both radio link and overall CLC transaction levels) could be satisfied. The considered deployment configurations are summarized as follows:

- 1 gNodeB (reference/baseline for benchmark only): The location corresponds to that of the closest LTE eNodeB site currently deployed on the Italian side (See Fig. 3-a)).
- 2 gNodeBs: The first gNodeB is the same as previously, while the location of the second corresponds to that of the closest LTE eNodeB site currently deployed on the Austrian side (See Fig. 3-b)).
- 3 gNodeBs: The two first gNodeBs are the same as previously, while the 3rd one is added "manually" on the Italian side, approximately halfway between the two former to bridge possible coverage gap (See Fig. 3-c)).

Besides, handover is not explicitly simulated but treated primarily through coverage assessment, by verifying if the resulting link quality experienced by the UE (i.e., given a perfect handover) would enable the right level of link reliability, and hence ultimately, of service availability. As regards to the induced latency (i.e., beyond reliability considerations), two simplified handover hypotheses are then accounted by means of extra penalty terms in the overall latency budget:

- 5G-NR Non Standalone (NSA), as it is: Here, we consider 5G-NR NSA without any specific improvement, neither

with respect to inter-MNO handover, nor with respect to LTE core network. Since standard cross-border roaming in this situation could lead to interruptions on the order of several tens of seconds (See e.g., field measurements in [14]), CLC service continuity would simply not be ensured in this case[2]. This is also somehow reflected in our single-gNodeB evaluations, by the fact that an UEs at the cell border could experience relatively poor link quality at 3.5 GHz, while waiting for association with the next network.

- Fully-fledged 5G-NR with tight inter-MNO coordination: Just like in [14], we also assume a more futuristic - and somehow idealized- case, where inter-MNO handover would be performed transparently and almost instantaneously at the border (say, with a delay on the order of a standard intra-MNO handover), between the last serving gNodeB on the Italian and the first available gNodeB on the Austrian side. This somehow prefigures higher levels of inter-MNO cooperation (e.g., through slice-aware handover) and anticipated/non-postponed handover decisions (i.e., contrarily to current field observations [14]).

D. MEC Deployment and Latency Assessment

We also consider several MEC deployment configurations, including fully decentralized (i.e., 1 MEC per gNodeB), clustered/semi-centralized (i.e., 1 per regional concentration site), and highly centralized (i.e., one per MNO data center) schemes, as schematically represented in Fig. 4.

In ns-3, the Radio Access Network (RAN) latency between a source and a destination (e.g., typically, for an UL CAM message from a vehicle up to the gNodeB) can be simulated up to the application level, meaning that it covers contributions from lower layers (i.e., physical layer, MAC/scheduling, Radio Link Control (RLC)...) but also further delays related to higher layers of the protocol (i.e., IP, TCP/UDP encapsulation

---

[2]Recent experiments [17] suggest that the network re-selection time could be considerably accelerated (typically down to a few seconds) through advanced Fast Network Reselection (FNR) mechanisms, which do not fall in the scope of this paper.
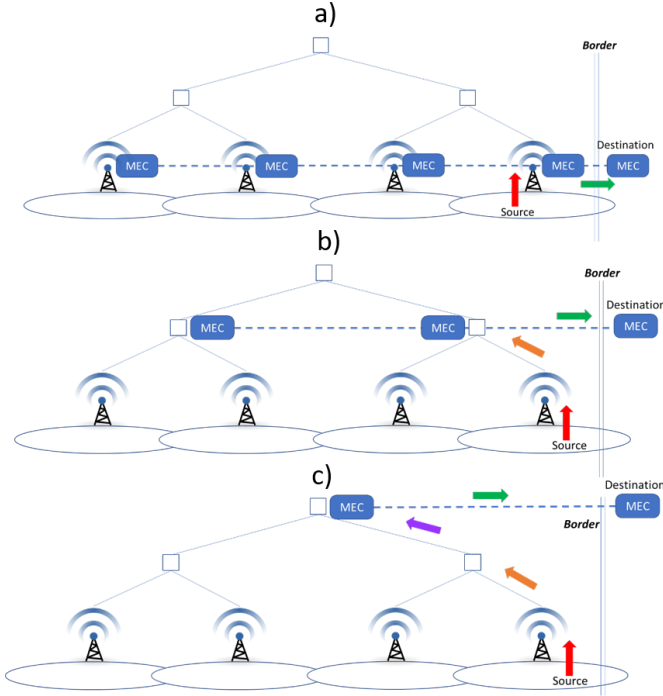
Fig. 4: Decentralized (a), clustered (b) and centralized (c) MEC deployment hypotheses, where coloured arrows represent the unilateral E2E information path from one vehicle up to the MEC (red: RAN, orange: backhaul, violet: core network, green: inter-MEC).

and transfer...), while accounting for jitter at the radio level, depending on link quality, resource scheduling and network load. Beyond the RAN aspects mentioned above, other latency components on top are not explicitly simulated either but represented in as abstract bounds. For instance, under a highly centralized MEC hypothesis, for the same UL CAM example as previously (i.e., from a vehicle up to the MEC in charge of elaborating the notifications in the centralized CLC mode), the overall E2E latency integrates extra delays related to backhaul, core network, and possibly inter-MEC connections (typically, in case of border crossing). Table II shows the mapping between the considered MEC deployment hypotheses, the situation of the vehicles involved in CLC transactions, and finally, the resulting latency terms that must be accounted in the overall E2E latency budget. For backhaul latency (e.g., between a gNobeB and a regional concentration site) and core network latency (i.e., between a radio access regional aggregation site and a data center), which are mainly distance-driven, we assume 2 ms for each component. This corresponds to lower bounds (typically, under the coarse assumption of 1ms/100km for fiber transport). Besides, although ultra-fast inter-MEC connections are mandatory (i.e., typically 1-2 ms per MEC switch for transferring service-related data/sessions, whatever the MEC deployment hypothesis), this term is left as an input parameter in the study.

E. Multi-service Assessment

Finally, we consider three different BWP schemes to illustrate possible benefits from flexible 5G slicing at radio physical
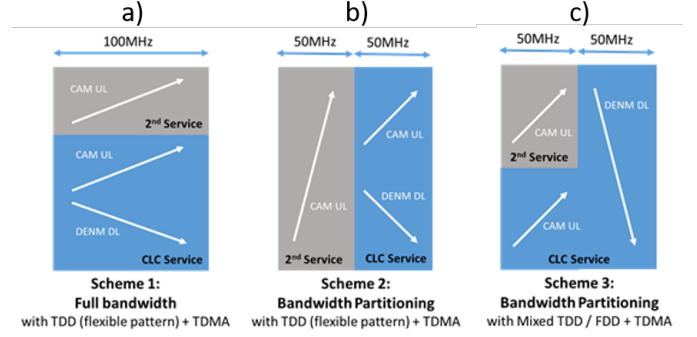

Fig. 5: Schematic representation of the different multi-service management schemes at radio physical level.

level with respect to multi-service coexistence issues. The first scheme consists in allocating the full bandwidth to both the primary CLC service (1st service) and another vehicular service (2nd service) running in the background (typically with UL transmission of regular CAMs at 20 Hz), while relying on nominal TDD with a default flexible pattern that automatically gives high priority to UL OFDM symbols. The second scheme based on BWP consists in allocating half of the available bandwidth to each service separately. The third scheme, which also uses BWP, relies on a mixed TDD/FDD scheme where each of the UL and DL flows is assigned to half of the overall bandwidth. Unless specified, the first scheme is applied by default.

IV. SIMULATION RESULTS

A. Latency in a Single-gNodeB Case

Fig. 6 shows the average UL radio access latency (averaged over simulation trials, vehicles and simulation time), as an integral component of a parametric E2E overall latency budget between the source vehicle and the MEC, depending on the different hypotheses introduced in Table II, as well as on the percentage of active users (i.e., with 1% to 25% of the simulated vehicles besides the CLC transmitting their own CAM messages at 20 Hz over V2N, while sharing the same spectrum resource), under both low (i.e., 30 veh./km, Fig. 6-a)) and high (i.e., 80 veh./km, Fig. 6-b)) traffic conditions. As expected, this overall E2E latency increases as a function of the percentage of active users in the "background" of CLC (typically from 10 ms at 1% up to 12ms at 25% in the most favourable low traffic case). In addition, the order of magnitude imputable to radio access seems to dominate by far other higher level latency components present in the budget, whatever the considered hypothesis. In line with the previous remark (i.e., since architectural aspects seem relatively more marginal), the fully decentralized MEC option could hence be only marginally advantageous at very first sight, so that the clustered MEC deployment may represent a good trade-off between low latency and reasonable deployment costs. All in all, even with a very basic scheduling policy at the radio access level, as that implemented in our simulations, a typical a priori requirement of 10 ms looks close to be achievable in most of the tested settings, thus suggesting significant margins

TABLE II: Mapping between MEC deployment hypotheses, CLC vehicles situations and latency budget components (UL message illustration)

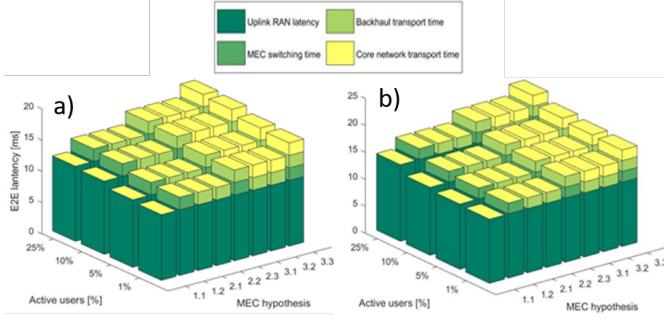| MEC Deployment	CLC Vehicles Situation	E2E Latency Components
Decentralized (1 MEC / gNodeB)	1.1) same cell	RAN (UL)
Decentralized (1 MEC / gNodeB)	1.2) different cells	RAN (UL) + Inter-MEC
Clustered (1 MEC / concentration site)	2.1) same cell	RAN (UL) + Backhaul
Clustered (1 MEC / concentration site)	2.2) different cells, same cluster	RAN (UL) + Backhaul
Clustered (1 MEC / concentration site)	2.3) different cells, different clusters	RAN (UL) + Backhaul + Inter-MEC
Centralized (1 MEC / data center)	3.1) same cell	RAN (UL) + Backhaul + Core Network
Centralized (1 MEC / data center)	3.2) different cells, same MNO	RAN (UL) + Backhaul + Core Network
Centralized (1 MEC / data center)	3.3) different cells, different MNOs	RAN (UL) + Backhaul + Core Network + Inter-MEC

Fig. 6: Average E2E latency budget for typical UL CAMs between a source vehicle and the MEC, as a function of MEC deployment hypotheses and CLC vehicles situations under low (a) and high (b) traffic conditions.

TABLE III: Link reliability as a function of the number of gNodeBs and percentage of active UEs for BWP Scheme 1 under high traffic density.

BWP1 / High Traffic		%age of active UEs			
		1%	**5%**	**10%**	**25%**
UL	**1 gNodeB**	96.67%	94.33%	90.99%	87.72%
	2 gNodeBs	97.99%	96.96%	91.12%	89.11%
	3 gNodeBs	98.98%	96.97%	93.86%	90.74%
DL	**1 gNodeB**	96.61%	95.81%	94.63%	92.01%
	2 gNodeBs	98.93%	97.72%	97.64%	94.88%
	3 gNodeBs	100%	98.83%	98.67%	96.48%

of improvement ahead with even smarter scheduling policies. Under high traffic conditions (Fig. 6-b)), UL latency seems to increase by a few milliseconds only in comparison with the low traffic density case (typically, spanning from 12 ms at 1% of active users up to 14 ms at 25% of active users). Other E2E latency evaluations regarding DL DENMs (not shown here for brevity) also confirm that the DL flow is subject to much lower latency (typically bounded down to a few ms at most), regardless of both the road traffic density and the amount of active users.

TABLE IV: Link reliability as a function of the number of gNodeBs and percentage of active UEs for BWP Scheme 1 under low traffic density.

BWP1 / Low Traffic		%age of active UEs			
		1%	**5%**	**10%**	**25%**
UL	**1 gNodeB**	98.71%	96.82%	95.78%	93.85%
	2 gNodeBs	99.96%	97.99%	96.01%	91.00%
	3 gNodeBs	100%	99.99%	97.52%	93.99%
DL	**1 gNodeB**	99.43%	97.42%	96.34%	95.79%
	2 gNodeBs	99.98%	98.72%	97.99%	96.87%
	3 gNodeBs	100%	100%	99.87%	98.02%

TABLE V: Average E2E radio access latency of UL CAM with BWP Schemes 1 and 3, as a function of the percentage of active users.

		%age of active UEs			
		1%	**5%**	**10%**	**25%**
UL	**BWP3 & Low traffic**	6.90ms	7.01ms	7.03ms	7.18ms
	BWP3 & High traffic	7.03ms	7.04ms	7.11ms	9.23ms
	BWP1 & High traffic	9.41ms	9.73ms	10.31ms	12.01ms
DL	**BWP3 & Low traffic**	0.85ms	0.87ms	0.91ms	0.95ms
	BWP3 & High traffic	0.89ms	0.95ms	1.11ms	1.68ms
	BWP1 & High traffic	0.87ms	0.92ms	0.99ms	1.67ms

B. Link Reliability under Gradual gNodeBs Deployment

Tables III and IV summarize the link reliability results, expressed as the percentage of non-corrupted Transport Blocks (TBs) over all the received TBs, accounting for the correct reception of UL CAMs and DL DENMs, as a function of both the percentage of users active in transmission and the number of deployed gNodeBs (given that seamless handover is performed), under high and low traffic conditions respectively. As expected, UL reliability decreases by far more rapidly than that of DL as a function of the percentage of users active. As expected too, in comparison with the single-gNodeB case evaluated here mostly for benchmark[3], the link reliability on both UL and DL is significantly improved while gradually adding a second gNodeB (necessitating tight inter-MNO coordination though) and beyond, a third gNodeB (thus, bridging further the "coverage gap" at 3.xGHz between the two first deployed gNodeBs). These gains are even more noticeable under high road traffic and high percentage of active users transmitting in the background of CLC, as shown in Table III. Note that the latter percentage is however likely over-pessimistic from a network load perspective in this evaluation. Much lower rates on the order of a few percents are indeed practically expected in more conservative scenarios regarding the technology penetration in the short/medium terms (typically for the 5 upcoming years or so). In case of low road traffic density and/or low penetration rate (typically less than 5%), if gNodeBs deployment is sufficient to ensure good coverage, the link reliability is even superior to 97% in all the tested configurations, which is already acceptable in comparison with standard requirements, without any further optimization.

---

[3]This scenario is just intended to reflect the lack of coverage at the cell edge of the first serving network while approaching/crossing the border, hence prefiguring service shortage in case seamless low-latency handover to another network cannot be feasible.

C. Multi-service Coexistence through Bandwidth Partitioning

Finally, we address the issue of multi-service coexistence by illustrating the impact of several bandwidth partitioning schemes on performances. In a first bunch of results (not shown here for length considerations), partitioning the overall available bandwidth is not shown to significantly improve link reliability for both DL and UL, since the main dominating parameters (i.e., propagation model, transmitted power and MCS) are invariant in comparison with the previous study. Similar trends as that of the full bandwidth case (See Tables III and Tables IV) can thus be observed. This was also foreseen in the sense the designed partitioning strategies (i.e., Schemes 2 and 3) are mostly intended to improve UL E2E latency (without degrading reliability though).

In Table V, we thus compare the impact on RAN latency of both the first and third bandwidth partitioning schemes (See Fig. 5). This table reports both DL and UL average latency, still as a function of the percentage of equipped vehicles transmitting 20Hz CAMs (with respect to all the simulated vehicles). According to the third BWP scheme, half of the overall bandwidth is thus reserved to the DL flow of the CLC service, while performing mixed TDD/Frequency Division Duplex (FDD). In comparison with Fig. 6, one can thus see that DL latency remains extremely low (slightly improved and systematically lower than 2 ms in all tested configurations), while some milliseconds can also be saved on the most critical UL latency in comparison with the first BWP scheme, where the full available bandwidth of 100 MHz is occupied by the UL and DL flows of both services. This shows that the impact onto the first CLC service of the extra network load generated by the background second service can be somehow mitigated and overcome through adequate physical slicing, even if mainly in terms of latency. Beyond, similar results have been obtained for the second bandwidth partitioning scheme, where half of the overall bandwidth is reserved to the first CLC service, including both of its UL and DL flows, whereas the second service is assigned to the other spectrum half. In this case however, we have considered only a worst case scenario of 25% of active users under the highest traffic conditions. In this situation, the second BWP strategy seems mainly relevant whenever the number of users belonging to simultaneous CLC procedures in a given environment remains relatively low, what is by the way expected in most practical real life situations (e.g., Scheme 2 achieving less than the required 10 ms for a percentage of CLC users lower than 5% among connected/active users, in comparison with Scheme 1, which was leading to 14 ms in the same operating conditions).

V. Conclusions

In this paper, we have provided and discussed a variety of simulation results that aim at illustrating the potential and limitations of 5G-NR connectivity to support CCAM services, through the example of a centralized CLC scenario.

First, in case other background vehicular services must coexist with the CLC over the same V2N 5G-NR resource, the E2E latency of the UL messages involved in a typical latency-critical service such as CLC may be problematic under high traffic density, a fortiori in case of high penetration of the 5G-NR technology on the road (typically, with up to 25% of the cars active in transmission in our evaluations). However, as a possible mitigation solution offered by 5G-NR technology (physical slicing), BWP can positively and largely contribute to relax this latency constraints at a level lower than 10 ms, in compliance with typical requirements. This kind of approach looks even more relevant when the latency-critical service to guarantee concerns a fairly low number of vehicles at a time. This eventually opens the floor to promising context-aware physical slicing in future 5G V2N connectivity to guarantee latency-critical services. Beyond, the radio access latency has been shown to dominate by far other higher-level latency components present in the budget, whatever the considered architecture hypothesis. In light of these observations, we have sorted out the clustered MEC deployment option, which may represent a good trade-off between low latency and deployment costs. All in all, even with the very basic scheduling policy that we have considered at radio access level in the previous evaluations, the typical requirement of 10 ms could be achieved in a significant number of tested configurations. Finally, regarding specifically cross-border latency issues, beyond the fine FNR capabilities already demonstrated on the field (i.e., a few seconds of interruption, as per recent field measurements carried out in a 5G NSA context), ultra-low latency inter-MNO cooperation seems almost inevitable, at least in highly localized cross-border portions of the corridor (e.g., slice-aware seamless handover), so as to keep this nominal RAN latency of 10 ms achievable, whatever the operating conditions, as already emphasized in [14].

Then, in terms of link reliability and ultimately, in terms of overall successful CLC messages exchange sequence (thus, impacting service availability), a network densification in localized cross-border areas looks also mandatory while operating at 3.x GHz (typically, up to 3 gNodeB over the tested 4.5 km of motorway), so as to bridge potential coverage gaps, given that low-latency inter-MNO handover can be performed. Overall, under such dense deployment, the demonstrated level of link reliability is also fairly close from the target KPI (i.e., typically larger than 97% even under high traffic density and for up to 5% of connected users transmitting at 20 Hz, what represents a quite aggressive operating context). Finally, these results still advocate the combination of PC5-Mode4 and V2N 5G-NR to achieve even better redundancy and hence, better robustness. Relying on the same evaluation framework, possible extensions of this simulation work will the evaluation of decentralized CLC and in-lane cooperative maneuvering scenarios, supported by the combination of both PC5 and 5G-NR technologies.

References

[1] A. Kousaridas et al., "5G Cross-Border Operation for Connected and Automated Mobility: Challenges and Solutions," *Future Internet*, vol. 12, no. 5, 2020.

[2] V. Mannoni et al., "A Comparison of the V2X Communication Systems: ITS-G5 and C-V2X," in *Proc. IEEE VTC'19-Spring*, April 2019.

[3] D. Wang et al., "System-Level Simulator of LTE Sidelink C-V2X Communication for 5G," in *Proc. ITG-Symposium on Mobile Communication - Technologies and Applications*, May 2019.

[4] M. Gonzalez-Martín et al., "Analytical Models of the Performance of C-V2X Mode 4 Vehicular Communications," *IEEE Trans. on Veh. Tech.*, vol. 68, no. 2, pp. 1155–1166, 2019.

[5] F. Poli et al., "Evaluation of C-V2X Sidelink for Cooperative Lane Merging in a Cross-Border Highway Scenario," in *Proc. IEEE VTC-Spring'21*, April 2021.

[6] A. Chiha Ep Harbi, et al., "Techno-economic and Simulation Study of a V2I-based Cooperative Manoeuvring Case in a Cross-border Scenario," in *Proc. IEEE VTC'21-Fall*, Oct. 2021.

[7] E. Coronado, G. Cebrian-Marquez, and R. Riggio, "Enabling Computation Offloading for Autonomous and Assisted Driving in 5G Networks," in *Proc. IEEE GLOBECOM'19*, Dec. 2019.

[8] N. Slamnik-Krijestorac, et al., "Network Service and Resource Orchestration: A Feature and Performance Analysis within the MEC-Enhanced Vehicular Network Context," *Sensors*, vol. 20, no. 14, 2020.

[9] N. Slamnik-Krijestorac, et al., "Collaborative Orchestration of Multidomain Edges from a Connected, Cooperative and Automated Mobility (CCAM) Perspective," *IEEE Trans. on Mob. Comput.*, 2021.

[10] M. Apruzzese et al., "5G-CARMEN Pilot Plan (Revised)," *Deliverable D5.1 of 5G-CARMEN project*, March 2021.

[11] A. Georgakopoulos et al., "5G-CARMEN Preliminary System Architecture and Interfaces Specifications (Revised)," *Deliverable D2.2 of 5G-CARMEN project*, May 2021.

[12] F. Visintainer et al., "5G CARMEN Use Cases and Requirements," *Deliverable D2.1 of 5G-CARMEN project*, May 2019.

[13] A. P. Da Silva et al., "System-level Simulation of Cooperative Sensor Data Fusion Strategies for Improved Vulnerable Road Users Safety," in *Proc. IEEE WPNC'19*, October 2019.

[14] J. Hillebrand et al., "Current 4G Networks Limitations Pleading for 5G in Cross-border CAM Services," in *Proc. IEEE 5G for CAM Summit 2021*, May 2021.

[15] D. Krajzewicz et al., "SUMO (Simulation of Urban MObility) - An Open-Source Traffic Simulation," in *Proc. MESM'02*, Sept. 2002.

[16] 3GPP, "Study on Channel Model for Frequencies from 0.5 to 100 GHz," in *3GPP Tech. Report 38.901, version 16.1.0 (Rel. 16)*, November 2020.

[17] C. Casetti et al., "Final Report on 5G Technological Enablers for CCAM," *Deliverable D3.3 of 5G-CARMEN project*, May 2021.