

A Framework for Cell Overload Mitigation in 5G Networks: A V2X Perspective

Martín Trullenque Ortiz, Oriol Sallent, Daniel Camps-Mur, Jad Nasreddine, Jordi Pérez-Romero, Jesús Alonso-Zarate

Abstract—Vehicular communications are called to be one of the most impacted sectors in the next generation of cellular networks, enabling a massive rollout of unseen vehicular use cases. However, the inherent high mobility associated to vehicles can lead to vehicular congestions, where the increase in data traffic introduced by vehicular communications can compromise network resources availability. To prevent such situations different overload mitigation approaches can be applied at both network and application layers. In this paper we present a Framework for cell Overloads Mitigation in 5G networks (FOM-5G), which is based on the O-RAN architecture and network exposure API defined in the 5G Core. Moreover, FOM-5G leverages the usage of historical information to select across network and application-level responses in each overload situation. We assess FOM-5G using realistic vehicular traces in two different traffic congestion scenarios. Simulation results demonstrate the effectiveness of customizing overload mitigation approaches to each scenario.

Index Terms—Cell overload, mobility load balancing, O-RAN, V2X, historical information.

INTRODUCTION

ONE of the verticals where 5G are expected to have the greatest impact is the automotive sector. The capability of 5G networks to provide guaranteed quality of service will become a major booster for vehicular communications and is expected to drive a massive rollout of new vehicular use cases, leveraging the need to deploy different vehicular services. A key service is the cooperative awareness basic service, standardized by the European telecommunications standard institute (ETSI) in TS 102 637 - 2, which enables vehicles to track surrounding vehicles in real-time. This service is implemented using Cooperative Awareness Messages (CAM), that are transmitted periodically to update the position and basic status of each vehicle. Excessive delays or failures to receive CAM messages may result in wrong position estimates of other vehicles, which may endanger vehicular safety.

V2X services can be provisioned either through direct short-range communications (DSRC) technology or through public cellular networks. Both support V2X communications in the unlicensed band of 5.9 GHz using either IEEE 802.11p, IEEE 802.11bd for DSRC or cellular V2X (C-V2X) for cellular networks. Using these technologies, vehicles can communicate directly between them and with roadside units (RSUs) deployed by road traffic authorities. However, the massive investment

required for a large-scale deployment of such RSUs on real roads is hindering the adoption of such technologies. Instead, a better alternative to provision V2X services is over public 5G networks, which offer a national-level coverage. In this context, the 3rd Generation Partnership Project (3GPP) defined in TS 23.286 a set of application enablers to ease the integration of V2X application functions (V2X AF) in 5G networks, where vehicles communicate with the cellular infrastructure using the Uu interface. Moreover, 3GPP TS 23.287 defined the interaction of these V2X AF with the 5G network through the network exposure function (NEF), which may request customized network analytics using the network data analytics function (NWDAF).

However, provisioning V2X services over public cellular networks poses significant radio resource management challenges. For example, the inherent mobility associated to vehicles [1] may result in vehicular congestion situations, where a large number of vehicles end up connected to the same cell causing an overload situation [2]. However, vehicular traffic congestion in certain areas is likely to exhibit periodic patterns due to activities such as daily commuting [1]. Thus, capturing these patterns in the form of historical information can be an effective approach to effectively anticipate and mitigate cell overloads. Yet, leveraging historical traffic patterns to mitigate overloads in cellular networks remains unexplored.

The classical approach to mitigate cell overloads in cellular networks is mobility load balancing (MLB), which acts on the network layer by diverting furthestmost users from an overloaded cell to the neighboring ones. In V2X scenarios, another way to relieve such congestion is to act on the application layer decreasing the frequency of CAM messages. We refer to this as V2X application congestion control (V2X ACC). An example of V2X ACC is the decentralized congestion control scheme defined by ETSI in TS 103 175 that adapts the CAM frequency depending on the radio congestion sensed by each vehicle. An adaptation of this approach for LTE C-V2X is presented in [3]. However, reducing the frequency of CAM messages increases the uncertainty about the position of neighboring vehicles, impacting their safety. To minimize this negative impact, both MLB and V2X ACC could be combined, as already identified in our prior work [4]. More recently, the growing interest in use cases in vehicular networks and Open Radio Access Networks (O-RAN) has been gathered in [5].

This work was supported by the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union – Next Generation EU, in the framework of the Recovery Plan, Transformation and Resilience (Call UNICO I+D 5G 2021, ref. number TSI-063000-2021-29).

Martín Trullenque Ortiz, Daniel Camps-Mur, Jad Nasreddine and Jesús Alonso-Zarate are with i2CAT Foundation 08034, Barcelona Spain (emails: martin.trullenque, daniel.camps, jad.nasreddine, jesus.alonso @i2cat.net.).

Oriol Sallent and Jordi Pérez-Romero are with the Department of Signal Theory and Communications, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain (emails: sallent@tsc.upc.edu, jordi.perez-romero@upc.edu).

Within this special issue, [6] proposes using an external database in an O-RAN architecture to enhance the performance of traffic steering and management algorithms in a vehicular scenario. In turn, [7] recently proposed using position information from a V2X AF to adjust beam management and optimize handovers in a synthetic vehicular scenario. However, combining both network and application layer-based algorithms is still unexplored.

With all the above, the main contributions of this paper are twofold. First, we propose FOM-5G, a comprehensive framework that combines MLB and V2X ACC to handle cell overloads caused by vehicular traffic congestions. To identify the best overload mitigation approach, FOM-5G exploits the use of historical information to match the current network situation to past patterns and selects the overload mitigation approach that proved more effective for that pattern. Beyond envisioning the framework to operate in vehicular scenarios, the design principles conceived empower FOM-5G to be applied for any application whose traffic offers a certain degree of repeatability. Furthermore, we describe a concrete implementation of FOM-5G, which leverages new cellular architectural elements like the RAN intelligent controllers (RICs) defined by O-RAN and the NWDAF function defined by 3GPP. Second, we demonstrate the gains in vehicular safety obtained with two overload mitigation approaches enabled by FOM-5G in two vehicular congestion situations obtained using realistic vehicular traces of the city of Cologne in Germany.

The remainder of this paper is structured as follows. The next section describes the design principles of FOM-5G. After that, we detail the system implementation together with the system operation. Then, we present the system evaluation in different scenarios. Finally, we provide the paper's conclusions and highlight our future work.

FOM-5G DESIGN PRINCIPLES

Fig. 1 depicts the functional architecture of FOM-5G, describing its three main building blocks. First, the cell status block maintains a set of network KPIs that characterize the conditions of the cells in the network. Second, given the current cell status, the smart cell overload reduction (SCOR) block decides for a particular cell the best overload reduction mechanism. Third, the historical information database (HID) stores past observed cell status along with the performance experienced after applying a particular overload mitigation approach. The HID is used to support the decisions taken by the SCOR block. Next, we describe the principle of operation of each module.

A Cell status

The cell status block periodically samples for each cell a set of network KPIs that are used to characterize the conditions of the cells in the network. This status is represented by a cell vector, which includes for a period of T seconds the following KPIs:

- The time of the day and day of the week.

- The geographic coordinates of the cell location.
- The number of connected users in the cell (see [7], section 5.1.1.4.1).
- The cell load (ρ_n in Fig. 1), which indicates the total demand of Physical Resource Block (PRB) in the cell (see [8], section 5.1.1.2).
- The neighboring cells load (ρ_k in Fig. 1) and their respective identifiers.
- The “transferable load”, defined as the load from the cell that could be offloaded to a neighbor cell. This can be obtained through measurement reports from the connected users.
- The incoming handovers (λ_{IN} in Fig. 1) from each neighboring cell to the interest cell (see [8], section 5.1.1.6.1.8).
- The outgoing handovers from the cell towards each neighboring cell (see [8], section 5.1.1.6.1.8).
- The users dwell time in the cell. The reader is referred to the following section for further details on how this metric can be computed.

The sampling period used by the cell status block is considered to be customized according to the vehicular dynamics of a region of interest. We estimate that a value between 30 seconds and 5 minutes should be adequate in most situations.

B Smart cell overload reduction (SCOR)

The functionality of this block in Fig. 1 can be split into three different blocks: overload detection (OD), vector identification (VI) and action selector (AS).

The OD block takes as input the current cell status and determines if it is required to apply an overload mitigation action for a given cell. A simplified implementation of this step is described in [9] and uses a threshold on the cell load, but more complex implementations could be considered.

If an overload mitigation action is required, the VI block queries the HID to select cell vectors that represent similar traffic congestion situations. To select these vectors, a distance metric across cell vectors is required. Looking at the cell status definition in the previous section, we note that Euclidean distance could be used for all entries of the cell vector, except for the time and date entry, for which a day similarity metric to capture daily events such as the one in [10] could be used. Using this metric the HID vectors that fall within a certain distance threshold from the current cell vector are pulled from the HID.

The last step in the SCOR module is to use the available cell vectors to select the overload mitigation strategy that needs to be applied in the current situation. To enable this selection, the HID stores cell vectors together with the previously selected overload mitigation strategies and the performance KPI resulting from them. The AS module uses these KPIs to rank the selected cell vectors and selects the overload mitigation strategy associated with the highest ranked vector. The used KPIs and ranking strategy are described in the next subsection.

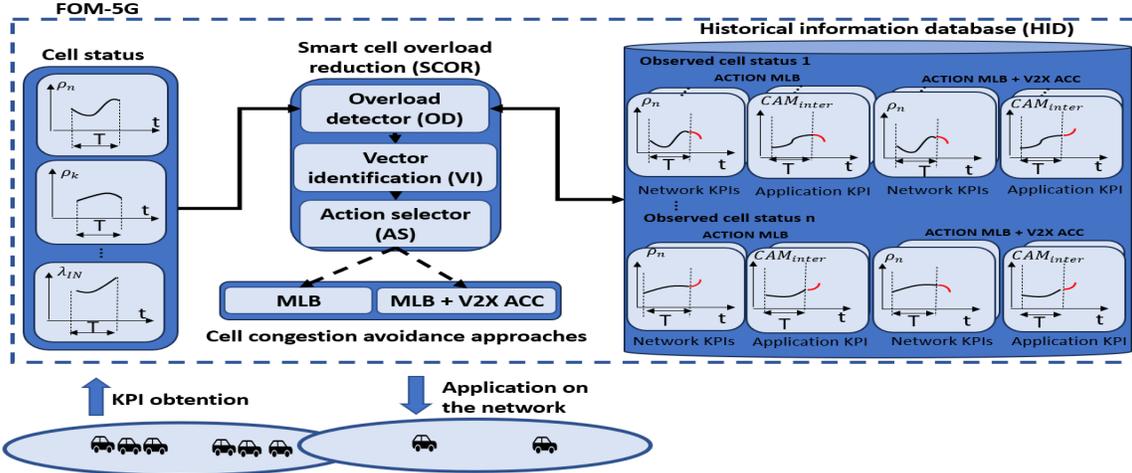


Fig. 1: FOM-5G conceptual architecture.

It is possible that the VI module does not find any cell vector in the HID that matches the current traffic conditions (e.g. due to a car crash not observed before). In this situation, a heuristic in the AS is used to select an overload mitigation strategy. For example, we prioritize maintaining the CAM frequency by applying MLB first, then adding V2X ACC if MLB alone does not suffice. Notice that, beyond taking a heuristic decision in a particular situation, this process will make SCOR to select the best overload mitigation strategy if these events happen again. This is a fundamental assumption in the design of the SCOR module.

Finally, the deactivation of the selected overload mitigation mechanism is determined by OD, again by monitoring the cell status. A simplified implementation of the deactivation decision consists of observing the cell load level below a threshold during a certain time. More complex implementations are left as future work.

Cell congestion avoidance strategies

FOM-5G considers the application of two overload mitigation approaches: MLB standalone or combined with V2X ACC. In this line, we consider V2X ACC only to be applied with MLB to avoid a potential unnecessary cooperative awareness basic service degradation.

Mobility load balancing (MLB)

In the event of performing MLB, UEs located at the edge of an overloaded cell will be reconnected to a neighboring cell. Our implementation of FOM-5G leverages the MLB algorithm proposed and evaluated in [9], which is based on adapting the cell individual offset (CIO) between two neighboring cells. The neighboring cell selection is based on their load and the transferable load obtained from the cell status metrics.

V2X application congestion-control (V2X ACC)

The V2X AF performs two basic operations: i) receive CAM packets from vehicles in the uplink, and ii) forward in the downlink each CAM packet to all vehicles that are within a certain distance of the transmitter. By default, the V2X AF forwards the received CAM packets in the downlink at the same

frequency that they are received in the uplink. However, upon network congestion, the V2X AF reduces the frequency used to transmit CAM packets in the downlink. A concrete implementation of this strategy is described in the performance evaluation section.

C Historical information database

The historical information database (HID) stores cell vectors containing three fields: i) cell status, ii) an overload mitigation action field, consisting of *MLB* or *MLB+V2X ACC*, and iii) the performance obtained after performing a given action, which we describe in this section. Storing this information enables capturing the repeatability patterns that vehicular mobility presents [1] and learning the best overload mitigation action from previous experience.

The HID cell vectors are shown in Fig. 1, where we illustrate the cell load evolution for a generic cell status after performing different actions. The performance (red plots shown in Fig. 1) reflects the impact of applying a given overload mitigation strategy to a specific cell status. To have a wider perspective on the effect an action has on the network, the performance is periodically sampled during a window of several minutes after the application of the action. Two sets of KPIs are included:

- Network-related KPIs, including cell load on the serving cell (represented as ρ_n in Fig. 1) and its neighbors, dropped packets rate, and the average delay observed by a cell when sending packets in downlink. Note that while the first captures if a cell can accommodate the demanded resources, the packet dropping and the delay reflect the safety degradation experienced at a network level.
- Application-related KPIs, including: the time between a CAM packet is received in uplink by the V2X AF and the time this CAM packet is successfully received by a target vehicle (represented as CAM_{inter} in Fig. 1). The V2X AF can use acknowledgements at the application level to confirm CAM delivery. A larger time interval between consecutive CAM receptions leads to worse position estimates and degrades safety.

The AS block from SCOR applies a ranking heuristic to the performance of different actions to sort HID cell vectors and decide on the overload mitigation strategy. A potential heuristic considers application-related KPIs first and uses network related KPIs to break ties across cell vectors. This leads to prioritizing MLB standalone if the overload can be mitigated but adding V2X ACC if avoiding severe service degradations.

FOM-5G IMPLEMENTATION FRAMEWORK

This section provides a concrete implementation of FOM-5G leveraging the standardized O-RAN architecture and the exposure APIs defined in the 5G Core, which is shown in Fig. 2. Moreover, Fig. 3 describes the detailed signaling workflow between all elements involved in the FOM-5G implementation.

Regarding the implementation, we can see that SCOR is implemented using two rApps within the O-RAN Non-RealTime (Non-RT) RIC. The first rApp, ODVI, integrates OD and VI modules explained in the previous section and the second rApp embeds AS functionalities. These two rApps implement the following interfaces: i) the R1 interface to interact with each other, ii) A1 interface to interact with the Near-RT RIC, iii) a custom interface referred as KPI enrichment interface (KEI) to push the cell status from the Near RT-RIC to the Non-RT RIC, iv) a custom interface referred as historical enrichment interface (HEI) to interact with the HID, and v) a custom interface defined as core enrichment information (CEI) to interact with the NWDAF function of the 5G Core. Notice that any commercial O-RAN system already supports the required programmability for adding custom interfaces to communicate between rApps and xApps or external databases.

A cell status composition xApp embedded in the Near-RT RIC constantly creates cell status vectors and feeds them to the ODVI rApp in the Non-RT RIC (point 1 of Fig. 3). To create these vectors, we leverage the E2 Key Performance Matrix (KPM) service model standardized in [11], which supports all the different cell status components listed in the previous section except dwell time. This can be computed by adding a counter in the E2 KPM service model that measures how long RAN IDs are allocated in a cell. The cell status composition xApp reports the measured cell status to the ODVI rApp through KEI.

Whenever the risk of having an overload is detected (overload predicted equals yes at Fig. 3), ODVI pulls matching cell vectors from the HID (point 2 of Fig. 3) using HEI. Then, ODVI forwards the information to the AS rApp through the R1 interface (point 3 of Fig. 3) and the AS rApp decides on the overload mitigation approach.

To implement MLB a dedicated xApp is deployed in the Near-RT RIC. Thus, whenever MLB needs to be conducted, the AS rApp configures the operation of the MLB xApp through the A1 interface (point 4a of Fig. 3). Then, the MLB xApp leverages the CIO control features available in the E2 RAN Control service model (point 5a of Fig. 3).

To implement V2X ACC the AS rApp notifies the network congestion to the V2X AF using the APIs exposed by the NWDAF and NEF functions of the 5G Core. In particular, a new analytics module is defined in the NWDAF, which we refer as the enriched information vehicular congestion analytics (EIVCA). The V2X AF creates a subscription to the EIVCA module, through which reports the application-related KPIs and receives from the network asynchronous notifications about the congestion status. Thus, when the AS rApp selects the V2X ACC action, it notifies the EIVCA (point 4b of Fig. 3) through the CEI interface. Then, the NWDAF resolves the geographical area covered by the overloaded cell and triggers an alert to the V2X AF through the NEF using the standardized Nnwdaf and Nnef interfaces (point 5b of Fig. 3). Finally, the V2X AF modifies its operational mode from default to congestion control mode.

To populate the performance related in the cell vectors of the HID, the ODVI rApp periodically collects network performance KPIs (point 6a of Fig. 3) from the cell status composition xApp, and application-related KPIs from CEI interface with the NWDAF (point 6.b of Fig. 3). Finally, this information is stored in the historical information database (point 7 of Fig. 3).

PERFORMANCE EVALUATION

In this section we use a simulative approach to study the effectiveness of the overload mitigation approaches considered by FOM-5G in two realistic scenarios located in the city of Cologne (Germany).

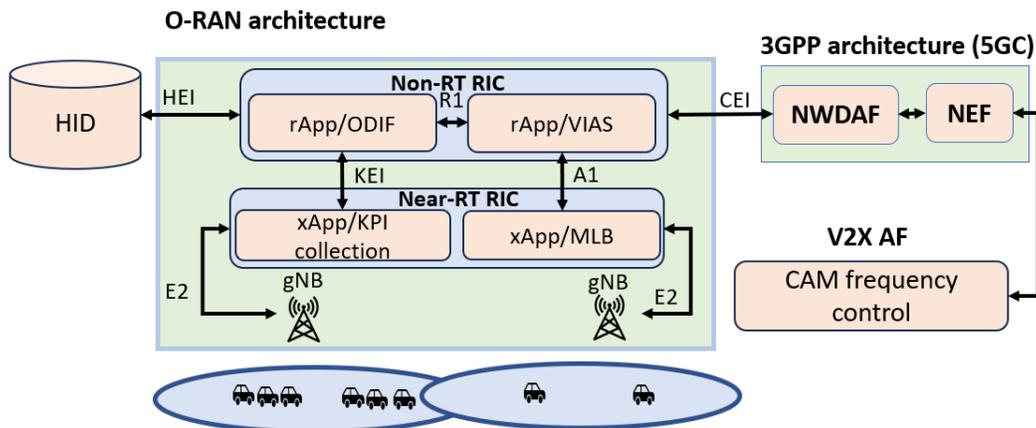


Fig. 2: Proposed implementation of FOM-5G.

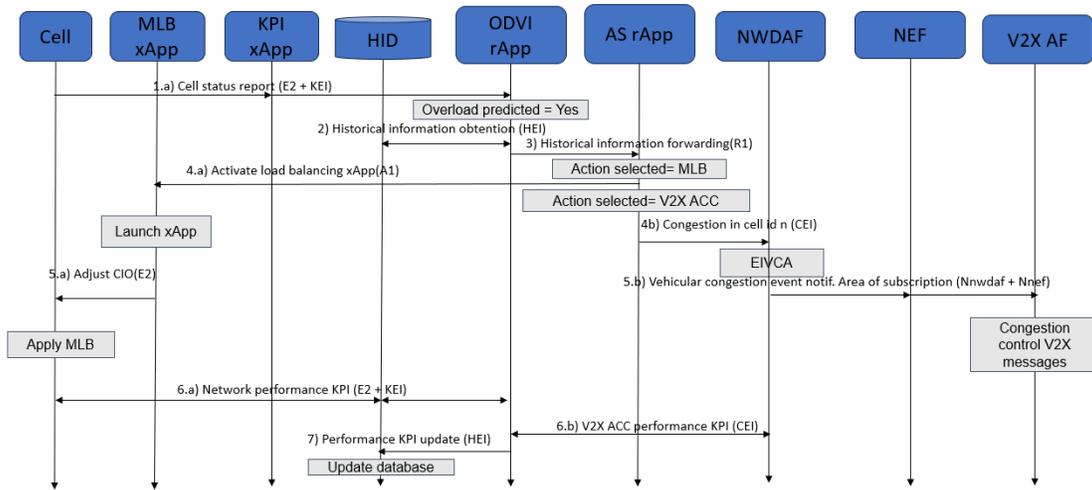


Fig. 3: Proposed workflow.

A. Scenarios considered

To evaluate FOM-5G we use a urban scenario and a suburban scenario. Due to the lack of realistic traces from other type of services, only vehicular data traffic is considered. Regarding radio deployment, in both scenarios we take the nodes position from the Deutsche Telekom network available in [12]. We assume omnidirectional cells transmitting at 2.1 GHz using frequency-division duplexing (FDD). We consider a symmetric carrier bandwidth of 10 MHz for the suburban scenario and 20 MHz for the urban scenario. The selected subcarrier spacing in both scenarios is 15 kHz, with a transmission bandwidth of 52 RBs and 106 RBs for the suburban and urban scenarios respectively. For normal cyclic prefix, 14 OFDM symbols can be transmitted in each slot.

The path loss model chosen is the urban area model, available at Section 4.5.2 of 3GPP TR 36.942. Neither slow nor fast-fading terms have been considered. Instead, the transmitted power at each gNB has been set to have a data rate around 800 kbps at the cell edge. With regards to the vehicular traces, we use the public TAPAS Cologne dataset [13], which contains 24 hours of realistic vehicular traces in a 400 km² area. Each vehicle generates a CAM packet as specified in [14]. CAM packets received in uplink are forwarded in downlink to all vehicles within an area around the vehicle of 0.16km², as in [9]. To do so we consider unicast communications, since enhanced multimedia broadcast multicast service is not ready yet in public cellular networks. To execute simulations, we use the system level simulator used in [9]. It is worth noting that the number of physical resources needed to transmit a data packet varies depending on the path loss between the vehicle and the base station.

We focus our analysis on a specific cell from each scenario. For the urban environment we focus on a cell covering a crossroad intersection regulated by traffic lights, where vehicular density is typically high and vehicular speed is low. Instead, for the suburban scenario we focus on a cell covering a highway junction, where different speeds for each lane can be envisioned.

B. V2X QoE KPI

To assess the quality of the cooperative awareness basic service we measure the errors between estimated positions and actual positions of neighboring vehicles. This error can be inferred from the age of information [15] or latency between consecutive CAM messages received from a given vehicle, combined with the speed at which the sending vehicle is moving. We consider that if a CAM packet cannot be sent due to lack of available resources, it is dropped. Thus, CAM drops during network congestion will result in higher CAM intervals, leading to larger error position estimates. Network congestion may be reduced using V2X ACC to decrease the CAM message rate, but this can also result in larger position error estimate (e.g. a one-second CAM interval will result in an error of 33 meters for a vehicle at 120 km/h). Hence, the impact of the overload mitigation strategies on the V2X QoE KPI needs to be carefully assessed for each scenario.

C. Implemented V2X ACC strategy

Annex B of ETSI TS 102 637 - 2 specifies that CAM packets should be generated every time a given UE moves 5 meters from the last CAM packet with a variable frequency, ranging from a maximum of 10 packets per second to a minimum of 1 packet per second. This is the default strategy used by the V2X AF in FOM-5G when the network is not congested.

Upon congestion, we modify the way the V2X AF forwards CAM packets in downlink for the slow vehicles. Specifically, CAM packets sent by vehicles moving at a lower speed than 0.75 m/s will be forwarded only every other second. Moreover, when a vehicle stops the CAM interval increases to 4 seconds. Two main benefits are derived from this strategy. First, the vehicular density in traffic jams is usually high, whereas the vehicles' speed is typically low. Thus, when a traffic jam occurs, this strategy is expected to apply to many vehicles, but the impact on the position estimate error of such vehicles will be irrelevant. Second, lowering CAM rate for slow users will enable fast users to observe a higher number of packets being successfully forwarded in downlink, avoiding an increase of

their position error estimate.

D. Overload reduction evaluation

We evaluate FOM-5G overload mitigation strategies in each cell of each scenario by looking both at the instantaneous downlink cell load and the cumulative distribution function (CDF) of the V2X QoE KPI. We have only focused on the downlink as the limiting link, since for each CAM sent in uplink, the V2X AF replicates it in downlink to all the vehicles in the surrounding area, leading to a very asymmetrical load. On both scenarios, we consider the HID to be already populated and FOM-5G working on inference mode.

We plot the total resources demand (requested load) normalized to 1. When the total resources demand exceeds 1, packets will be lost. For instance, a requested load of 1.5 would reflect a system throughput of one and 33% of lost packets.

Fig. 4 and Fig. 5 present the requested load evolution for a selected cell in urban and suburban scenarios respectively. For each scenario, a different busy hour has been chosen. For the urban scenario, the busy hour encompasses from 7:00 to 8:26, in which a maximum requested load of 1.25 at 7:42 can be observed (blue plot of Fig. 4). In turn, the suburban scenario considers the afternoon traffic (from 17:42 to 18:32), achieving the highest load value of 1.55 at 18:09 (blue plot of Fig. 5).

Focusing on the MLB standalone, each scenario presents a different performance. On the one hand, the urban scenario (red plot of Fig. 4) presents a limited impact with a load reduction of 0.06 (from 1.25 to 1.19). On the other hand, the suburban scenario presents promising MLB performance (red plot of Fig. 5), decreasing maximum cell load by 0.32. This effect is highlighted when observing the overload time duration, where there is a reduction from 8 minutes and 31 seconds to 2 minutes and 30 seconds (70.59%).

Applying MLB and V2X ACC also provides a different performance depending on the scenario. The urban scenario effectively mitigates the overload as many users are stopped at a traffic light. In turn, this impact is more limited on the suburban scenario (green plot of Fig. 5), where the load is only reduced by 0.16, half of the reduction achieved with MLB. Alternatively to the urban scenario, many vehicles present a high speed and only a limited number of them are stopped.

The impact of decreasing the CAM message rate has over the V2X QoE for the urban scenario is illustrated in Fig. 6. This shows the CDF of the position error estimate when the cell is overloaded (5 minutes and 31seconds). If no overload mitigation approach is in place (magenta plot), 8.88% of the vehicles contain position estimates with errors larger than 5m. This is reduced to 3.68% if MLB is in place (blue plot). Finally, using both MLB and V2X ACC ensures that the position error update is always below 5 meters (green plot).

A similar trend is observed in the suburban scenario, which we only report in writing. If no overload mitigation approach is in place 9.64% of users observe a distance estimate error larger than 5 meters. Applying only MLB decreases this number to 2.07%. If both MLB and V2X ACC are combined, the percentage of users observing an error estimate larger than 5 meters is only 0.52%.

Wrapping up the performance analysis, in the urban scenario MLB delivered an improvement of 58% in the V2X QoE KPI, whereas in the suburban case MLB delivered an improvement of 78%. This is mainly due to balanceable load difference. Although both scenarios present cell overload introduced by traffic congestions the location of each traffic jam is of key importance. While in the urban scenario users are located near the cell, in the suburban scenario the traffic congestion is located close to the cell edge, driving to a successful MLB performance (check [4] for further details). Instead, the reverse effect is observed in the performance of the V2X ACC. The higher number of vehicles stopped in the traffic light leverages a better V2X ACC performance in the urban scenario than in the suburban one. We conclude that the overload mitigation approaches considered in FOM-5G are effective in protecting V2X QoE in realistic vehicular scenarios.

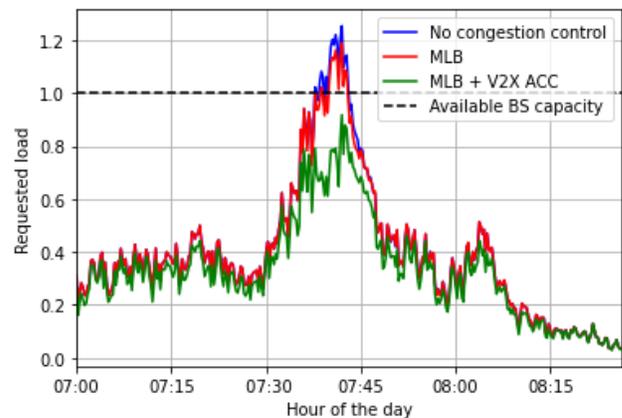


Fig. 4: Urban scenario requested load of the cell of interest over time.

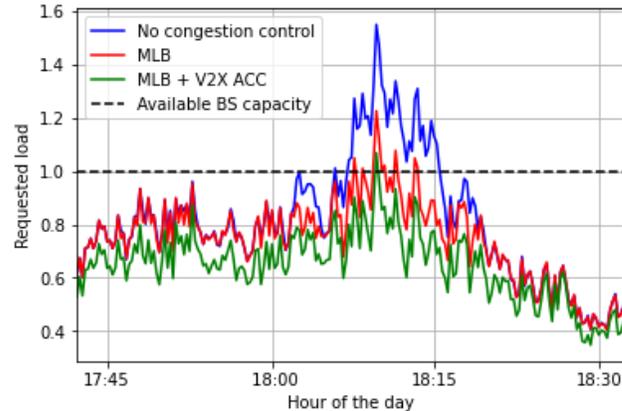


Fig. 5: Suburban scenario requested load of the cell of interest over time.

CONCLUSION AND FUTURE WORK

This work has presented FOM-5G, which is a novel framework to manage cell overloads in vehicular scenarios. FOM-5G combines the exposure capabilities of the 5G Core to interact with the application layer functions and relies on O-RAN to implement MLB algorithms that correct cell overload situations. Two realistic scenarios have been considered for evaluation, demonstrating how the nature of each scenario led

to a different effectiveness of the considered overload mitigation strategies. Through our evaluation we have showcased the need to consider both network and application KPIs for selecting the best overload mitigation action in each scenario.

Our future research will explore the design, implementation and validation of the SCOR and HID modules based not only on realistic vehicular traffic but also on real eMBB traces from a public MNO. Further research also considers investigating network slicing for different data traffic sources (e.g. V2X traffic, generic eMBB or IoT applications). This will leverage the FOM-5G framework by adding new services on the NWDAF to customize different overload reduction approaches. Another future work is the exploration of congestion in asymmetric FDD or in time division duplexing schemes, e.g. through proper adjustment of the UL/DL switching point.

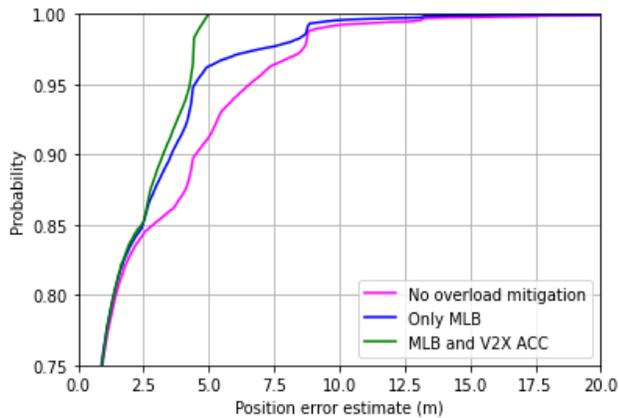


Fig. 6: CDF of the error position estimate in the urban scenario cell of interest.

REFERENCES

- [1] L. Pappalardo et. al., "Understanding the patterns of car travel," *The European physical Journal Special Topics*, vol. 215, pp. 61-73, 2013.
- [2] Z. Li et. al., "User Association for Load Balancing in Vehicular Networks: An Online Reinforcement Learning Approach," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 8, pp. 2217-2228, 2017.
- [3] M. Roshdi et. al., "Deep Reinforcement Learning based Congestion Control for V2X Communication," in *2021 IEEE 32nd Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, Helsinki, 2021.
- [4] M. Trullenque et. al., "On Alleviating Cell Overload in Vehicular Scenarios," in *IEEE Vehicular Technology Conference, 2022*, London, 2022.
- [5] K. Dev et. at., "Guest Editorial Introduction to the Special Section on Open Radio Access Networks: Architecture, Challenges, Opportunities, and Use Cases in Vehicular Networks," *IEEE Transactions on Vehicular Technology*, vol. 73, no. 7, pp. 9221-9225, 2024.
- [6] P. Sroka et. al., "Policy-Based Traffic Steering and Load Balancing in O-RAN-Based Vehicle-to-Network Communications," *IEEE Transactions in Vehicular Technology*, vol. 73, no. 7, pp. 9356-9369, 2024.
- [7] K. Suzuki et. al., "Toward B5G/6G Connected Autonomous Vehicles: O-RAN-Driven Millimeter-Wave Beam Management and Handover Management," in *INFOCOM*, Vancouver, 2024.
- [8] 3GPP, "5G, Management and orchestration; 5G performance measurements (Release 16)," 3GPP, TS 28.552 version 16.6.0 Release 16, 2020.
- [9] M. Trullenque et. al., "On the Application of Q-learning for Mobility Load Balancing in Realistic Vehicular Scenarios," in *IEEE 97th Vehicular Technology Conference*, Florence, 2023.
- [10] Z. Jankovic et. al., "Day similarity metric model for short-term load forecasting supported by PSO and artificial neural network," *Electrical Engineering*, vol. 103, p. 2973-2988, 2021.
- [11] O-RAN Alliance, O-RAN Working Group 3 Near-Real-time RAN Intelligent Controller E2 Service Model (E2SM) KPM.
- [12] CellMapper, "Signal tiles and towers," CellMapper, [Online]. Available: <https://www.cellmapper.net/>. [Accessed 4th September 2021].
- [13] S. Uppoor and M. Fiore, Large-scale urban vehicular mobility for networking research, Amsterdam: VNC, 2011.
- [14] L. Lopez et. al., "Understanding the impact of the PC5 resource grid design on the capacity and efficiency of LTE-V2X in vehicular networks," *Wireless communications and mobile computing*, 2020.
- [15] S. Kaul et. al., "Real-time status: How often should one update?," in *2012 Proceedings IEEE INFOCOM*, Orlando, FL, 2012.

Martín Trullenque Ortiz is a researcher at the Mobile and Wireless Internet (MWI) group at i2CAT Foundation. Currently pursuing a PhD at Universitat Politècnica de Catalunya (UPC) he received his master's degree at UPC in 2019. His research interests include radio resource management and vehicular communications.

Oriol Sallent is a professor at UPC. He has participated in different European projects with diverse responsibilities as work package leader and coordinator partner, and he has contributed to standardization bodies such as 3GPP, IEEE and ETSI. He has published more than 200 papers. His research interests include radio resource management, self-organizing networks and QoS provisioning in heterogeneous wireless networks.

Daniel Camps-Mur is leading the MWI group at i2CAT Foundation. Previously, he was a senior researcher at NEC Network Laboratories in Heidelberg, Germany. He holds a Ph.D. from UPC. His research interests include mobile networks and IoT.

Jad Nasreddine is a senior researcher at the MWI group at i2CAT Foundation. Previously, he was an Associated professor at Rafik Hariri University, Lebanon. He holds a Ph.D. (2005) from university of Rennes and bachelor's degree (2001) in computer science and telecommunications from the Lebanese university.

Jordi Pérez-Romero is a professor at UPC in Barcelona, Spain. He has been working on wireless communication systems, with a particular focus on radio resource management, cognitive radio networks and network optimization. He has been involved in different European projects and in projects for private companies. He has published more than 200 papers in international journals and conferences.

Jesús Alonso-Zarate received his Ph.D. (2009) and M.B.A. (2016) from UPC. He is the R&D director at i2CAT Foundation. Since 2010, he has published more than 150 peer-reviewed scientific papers. He is a recipient of various best paper awards and is very active in international R&D projects.