

Low Earth Orbit Satellite Supported Multi-Hop Dissemination of Messages in V2X Networks

Mario Franke^{*}, Roland Stroop[‡], Florian Klingler^{†‡}, and Christoph Sommer^{*}

^{*}TU Dresden, Faculty of Computer Science, Germany

[†]TU Ilmenau, and [‡]Paderborn University, Dept. of Computer Science, Germany

<https://www.cms-labs.org/people/{franke,sommer}>

<https://www.fklingler.net>

Abstract—Due to their unique mobility, Low Earth Orbit (LEO) satellites can provide network services in areas with little to no existing infrastructure and due to their low orbit they enable communication at low latencies. We thus motivate the use of LEO satellites for supporting multi-hop Vehicle-to-Everything (V2X) communication as an alternative to dense base station deployment and present CLEO, an approach that can opportunistically use LEO satellites for information dissemination among vehicles. Using extensive computer simulations, we show that not only are relaying latencies low enough for such an approach to be feasible, but we can actually improve performance of information dissemination in terms of all of Packet Delivery Ratio (PDR), Channel Busy Ratio (CBR), and end-to-end latency.

I. INTRODUCTION

Future Cooperative Autonomous Vehicles (CAVs) will rely on Intelligent Transportation Systems (ITS) in order to benefit from a variety of different vehicular applications. These applications range from safety principles (e.g., simple cooperative awareness) up to vehicular micro cloud supported Mobile Edge Computing (MEC) for large scale content dissemination in highly mobile networks. In order to address the requirements of such applications – particularly focusing on prime networking metrics such like load, reliability, coverage, and communication latency – many of those applications rely on Vehicle-to-Vehicle (V2V) communication technologies and multi-hop information dissemination [1], [2].

As a guiding example and without loss of generality, in this work we take the example of disseminating information about the presence of emergency vehicles on their way to an accident scene on a multi-lane cross-country road, which exhibits all of the aforementioned requirements, thus can benefit from multi-hop information dissemination. In almost all countries worldwide it is mandatory to clear a path for emergency vehicles – and even though some countries mandate to preemptively clear a lane as soon as a jam forms (to be used as a *rescue lane*), this is not always done for various reasons. In fact, a survey by the German Red Cross [3] (where preemptively clearing a dedicated rescue lane is mandatory) shows that in 80% of the cases, the path is not cleared in time leading to an average time loss of as much as 5 min for the emergency vehicle. A way forward might be to use ITS to disseminate information about the presence of emergency vehicles on their way to an accident scene. The timely receipt of such a message would allow a CAV or a human driver to clear

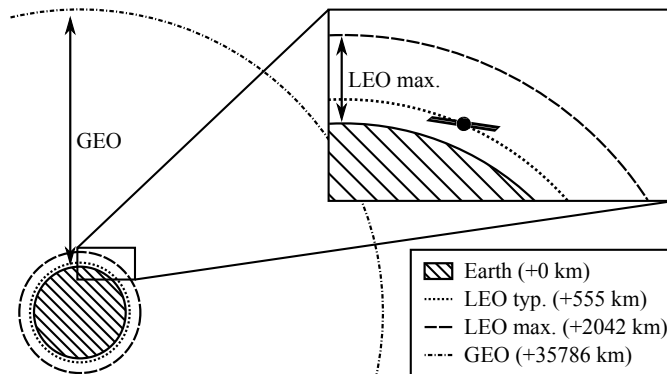


Figure 1. Illustration of orbits of different satellites

a rescue lane in time, cutting back substantially on the response times of emergency vehicles. Thus, information dissemination beyond the range of single-hop V2V communication is needed.

Initial approaches to low-latency multi-hop communications relied solely on V2V communication technologies and multi-hop information dissemination. Consequently, they suffered from high radio resource requirements and fragmented network topologies in high and low traffic density scenarios, respectively. To mitigate these issues, infrastructure is often relied on for information dissemination, though such infrastructure is often not available in the required density and coverage [4].

In this work, we motivate the use of an emerging alternative to static infrastructure that could also support vehicular networks, namely Low Earth Orbit (LEO) satellite constellations. Their most prominent advantages are wide communication coverage while still maintaining low communication latencies [5]. The main reason for achieving communication latencies in the range of very few milliseconds is the low orbit of LEO satellites, as visualized in Figure 1. However, these low orbits require highly mobile satellites. Each LEO satellite is only visible for a couple of minutes from a static point of view. On the one hand, these rapid changes in network topology mean high control overhead, on the other hand, this allows to compensate network outages within minutes or seconds since another LEO satellite is already approaching. Thus, while maintaining large LEO satellite constellations is costly, a sparse deployment (and then relying on mobility of those satellites to temporarily

fill coverage holes in time) might be cheaper than deploying static cell towers for full coverage. Moreover, such a novel networking architecture employing LEO satellites constitutes a distributed system with no single point of failure (unlike a sole cell tower in the countryside). Lastly LEO satellites can support communication even in post-disaster scenarios where ground-based infrastructure is no longer available.

We thus investigate the suitability of LEO satellites for opportunistically supporting multi-hop relaying of Vehicle-to-Everything (V2X) messages, e.g., as required for rapid rescue lane creation or any one of a multitude of applications.

In order to get performance insights of such a system, we integrate LEO satellite communication in a cluster-based multi-hop V2X communication scheme, which we refer to as *CLEO*. We then evaluate the network performance based on three complementary metrics, running detailed simulations featuring advanced LEO satellite mobility models.

II. RELATED WORK

Many different types of V2V networks that can benefit from additional support, e.g., by fixed infrastructure, have been studied in the past.

A very intuitive type of infrastructure is the use of parked cars acting as virtual network infrastructure [6]. The parked cars form clusters which provide network service accessible through gateway vehicles. In order to compensate for the mobility of driving vehicles, the cluster itself manages the handover between gateways and driving vehicles. A simulation study showed that such a virtual network infrastructure can act as robust access to backbone networks. However, parked cars are static and often concentrated in cities, thus such an approach does not benefit highways or freeways much.

A more advanced architecture, consisting of driving vehicle clusters, also called vehicular micro clouds, parked cars, Roadside Units (RSUs), and Long Term Evolution (LTE) infrastructure is presented by Pannu et al. [7]. Vehicular micro clouds are used to gather data and upload it via LTE or an RSU to a data center, which benefited upload performance.

A different approach to involve mobile broadband services for supporting multi-hop communication in vehicular networks is presented by Di Maio et al. [8]. In their approach, vehicles send routing requests to a Software-Defined Network Controller (SDNC) over a 5G network for transmitting data over several vehicular hops. Each vehicle periodically updates its neighbor list, which is stored at the SDNC. Based on this knowledge, the SDNC selects an optimal route and announces it to every involved vehicle, which can then use the route to disseminate data.

To be more flexible in deploying infrastructure and to be independent of often lacking base station coverage, researchers more recently investigated the integration of Unmanned Aerial Vehicles (UAVs) in vehicular networks. Specifically, Khabbaz et al. [9] investigate scenarios in which vehicles and drones can communicate with each other in a stable multi-hop network. However, the major drawback of drones is their limited operation time due to their battery limits.

Only little research has been conducted regarding opportunistic hybrid LEO satellite V2X networks – in fact, to the best of our knowledge the work most similar to ours presented in this paper dates back to 2012. Kloiber et al. [10] assume a low traffic density Vehicular Ad Hoc Network (VANET) in which not all vehicles can communicate with each other due to the limited communication range of existing V2V technologies. This gap is bridged by an idealized satellite connecting every car to a central *satellite hub*, thus offering a central relay to be used for flooding messages through the system.

In our work, we pursue a more general goal, namely a fully distributed system that strives to conserve channel capacity and which can benefit opportunistically from the presence of LEO satellites. We consider a broad range of network densities along with detailed models for V2V communication and, in particular, of LEO mobility. We also consider a diverse set of complementary network metrics to investigate the performance.

At a glance, this allows us to present *CLEO*, a versatile infrastructure-assisted V2V networking approach which makes use of LEO satellites to increase application performance (i.e., in our case the number of informed vehicles) while at the same time reducing communication latency as well as slightly reducing channel load – thus offering resources for other applications in the vehicular networking context.

III. THE CLEO APPROACH

We integrate LEO satellite communication into an opportunistic, cluster-based multi-hop V2X communication scheme and call this *CLEO*.

Clustering in VANETs is a well investigated research area [11], so we decided to base our approach on a proven clustering scheme. Specifically, we start with ideas of Gunter et al. [12] to develop strategies to reduce the number of vehicles using the LEO satellite channel. This is necessary to keep communication overhead on the satellite link at low values, and at the same time make specific use of the advantages of V2V communication, where necessary. To achieve this, each vehicle is equipped with a V2V communication radio module (e.g., using IEEE 802.11p) and an additional radio module for communication with LEO satellites.

In this approach, each vehicle adjusts its cluster state according to the cluster states of the other one-hop neighbors following the road in the same direction of travel. Therefore, each vehicle must periodically share its current cluster state. Without loss of generality we combine cluster state and potential payload into one message, to be (re-)transmitted on the V2V communication channel at fixed intervals; we call this message simply a Cluster Message (MSG). Choosing and transitioning between roles in the cluster takes place as follows.

Undecided Cluster Node (UCN): This is the initial and fallback role for all other roles if no neighbors are detected (via received MSGs). If a received MSG originates from a Cluster Head (CH), the UCN immediately joins the cluster of that CH as a Cluster Node (CN). Otherwise, if the UCN does not receive a MSG from another CH

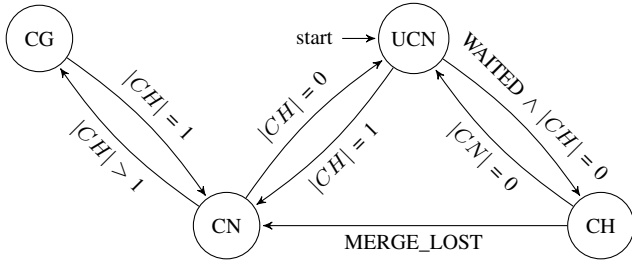


Figure 2. Cluster protocol state machine used in CLEO: State transitions are predominantly based on the number of detected neighbors belonging to a certain role.

within a predefined time, referred to as α , then the UCN becomes a CH itself.

Cluster Head (CH): Each cluster has exactly one CH. If a CH receives a MSG from another CH, then a cluster merge is triggered in which the CH with the higher vehicle ID remains CH and the other one changes to a CN.

Cluster Node (CN): A CN is a member of a cluster which receives MSGs from exactly one CH. It is not allowed to rebroadcast any MSGs. If a CN receives MSGs from multiple CHs it becomes a Cluster Gateway (CG).

Cluster Gateway (CG): A CG is a vehicle which receives MSGs from two or more CHs. By rebroadcasting MSGs, it acts a bridge between two CHs.

Figure 2 illustrates these state transitions.

Cluster operation requires the selection of two timing parameters, specifically α and a general timeout value used for declaring a connection to a CH as interrupted, referred to as β .

α must be large enough such that a UCN can receive a CH's cluster state at least once. Thus, it has to be larger than the interval at which CHs transmit their own cluster state. In order to compensate for channel access delay and transmission delay, we set α to twice the MSG generation interval.

As we use two times the MSG generation interval as α , we take the tripled MSG generation interval for β . This could result in frequent cluster state changes but reduces the number of dead connections.

Based on the dynamic assignment of roles (CN, CG, CH) to vehicles, CLEO operates as follows: Received MSGs are re-broadcast on the V2V channel by both CHs (to ensure dissemination within a cluster) and by CGs (to ensure dissemination among CHs) regardless of the origin of the MSG; to prevent duplicates, every CH and CG rebroadcasts every MSG only once. We further enable all CHs to retransmit any received MSG also via a LEO satellite in range, as illustrated in Figure 3. With this approach, we expect to gain two advantages. First, the CHs act like a satellite base station for their CNs; thus, the main advantage of satellites broadcasting messages over a large area would be fully utilized. Second, the joint use of the two types of communication should combine both advantages. If only the satellite is used, this will result in unnecessarily long delays, as the direct neighbors could be reached more quickly via V2V.




Figure 3. Communication topology of CLEO: using a Vehicle-to-Vehicle (V2V) channel, Cluster Heads (CHs) forward messages within clusters and Cluster Gateways (CGs) between clusters; in parallel, CHs exchange messages via Low Earth Orbit (LEO) satellites in range.

In summary, CLEO disseminates MSG over V2V to vehicles in proximity and provides a hierarchical structure for LEO satellite access. The LEO satellite is then used to spread MSGs over a large distance.

IV. EVALUATION

We investigate the performance of CLEO based on a simulation using OMNeT++ 5.7 as simulation engine, The INET Framework 4.2.1 for wireless networking models, Veins 5.2 and SUMO 1.8.0 for realistic vehicle mobility, as well as space_Veins¹ (pre-release version 0.3) for realistic LEO satellite mobility based on the Simplified General Perturbations 4 (SGP4) model.

The generation frequency of MSGs per vehicle is fixed at 1 Hz, with each message being 500 Byte long. The V2V link is configured to model IEEE 802.11p at a frequency of 5.89 GHz, with transceivers configured to 6 Mbit/s and a transmit power of 20 dBm.

Since our focus is on the attainable benefit when using LEO satellites, we configure a simplistic communication model of the channel between vehicles and the LEO satellite. In general, the Bit Error Rate (BER) is 0 whenever three conditions are met: a line of sight between the vehicle and the LEO satellite exists (modeling that the satellite is not hidden by the earth curvature), the corresponding elevation angle is greater than 25° (modeling that transmissions are not blocked by buildings or similar obstacles), and no transmissions that overlap in the time domain. Otherwise, the BER is 1. Regarding the downlink channel, interference or simultaneous transmissions by other LEO satellites are not assumed to happen as their transmissions are supposedly coordinated. In contrast, two uplink transmissions (by two vehicles), overlapping in the time domain, interfere each other resulting in both packets not being receivable. Uplink and downlink channel capacity are configured to 10 Mbit/s and 100 Mbit/s, respectively.

As a baseline we simulate the same system without the presence of satellites, allowing us to quantify the added benefit of LEO support.

Both the baseline and the CLEO approach is simulated for vehicles driving at 100 km/h on a 18 km long highway with

¹<https://sat.car2x.org>

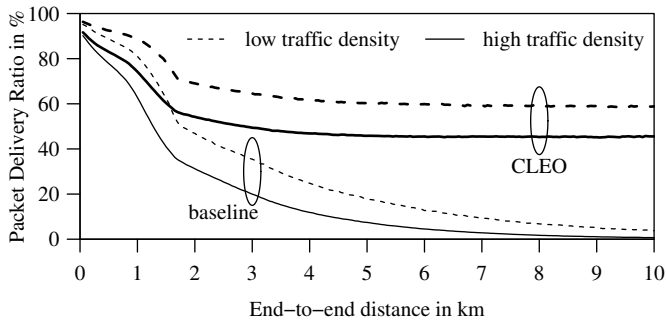


Figure 4. Packet Delivery Ratio (PDR) vs. distance

three lanes per direction using two traffic densities. The low traffic density is set to 5 veh/km for each direction, thus in total 10 veh/km. As doubling the low traffic density leads to a deterioration of the performance, 20 veh/km is considered as the total high traffic density.

To easily generalize satellite performance, we configure the highway to be centered at $0^{\circ}\text{N } 0^{\circ}\text{E}$ and the LEO orbit to the following NASA/NORAD Two-line Element Set (TLE):

```
1 51472U 22010S 22053.00000000 .00000000 00000+0 00000-0 0 9995
2 51472 70.0000 000.0000 00000000 00.0000 000.0000 15.73209361 2401.
```

For the purpose of the SGP4 model, the simulation start is configured to correspond to 22 Feb 2022 13:31:51.

A single simulation covers a time interval of 140 s. We repeat each simulation 50 times for statistical validity and take care to not record metrics during the transient period at the start of the simulation (while the highway is filling with vehicles and clustering reaches a steady state) and at the end (for not counting packets in flight as lost).

A. Packet Delivery Ratio

We calculate the PDR based on the assumption that, for every MSG generated within a 12 km long Region of Interest (ROI), all other vehicles currently in the ROI should be reached. However, we are only showing results within 10 km because for longer distances the number of communication partners does not allow to draw statistically valid conclusions.

In Figure 4 the average PDR is plotted separately for vehicles at different distances, dividing the ROI into 50 m long bins.

Three main observations can be made: First, there is a significant PDR drop after 1600 m. This drop coincides with the approximate upper limit for one-hop transmissions using V2V communication.

Second, the traffic density has an impact on both the baseline and CLEO. With higher traffic density, the number of (re-) transmissions increases as more vehicles want to transmit their MSG. But in order to benefit from the LEO satellite, at least a single CH must successfully receive a MSG. Thus, the V2V channel's PDR impacts the performance of CLEO. Secondly, in our simulation, the LEO satellite uplink channel access is unmanaged (pure ALOHA). Hence, more transmissions to the LEO satellite also increase the probability that they interfere. Likewise, the successful reception at the LEO satellite

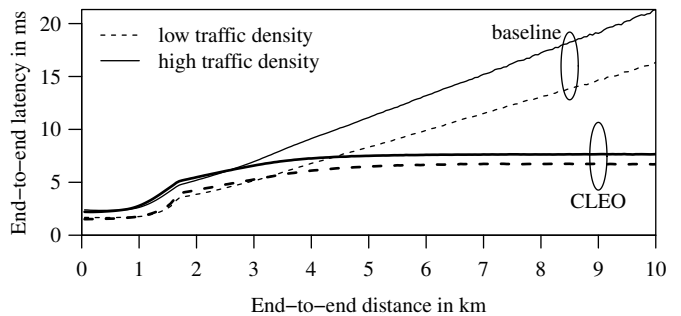


Figure 5. End-to-end latency vs. distance

decreases, too. The result is the PDR still being affected by traffic density.

The third observation is that, for long distances, the PDR of CLEO becomes constant while it keeps decreasing for the baseline approach. This is explained by the properties of the LEO satellite link. Once a MSG is received by the LEO satellite, it is relayed back to all vehicles. As discussed, the satellite downlink channel is assumed to not suffer from interference or bit errors. Consequently, all MSGs transmitted by the LEO satellite are always successfully received by all vehicles, and thus, long distances between senders and receivers do not have an impact on the PDR. In contrast to CLEO, the baseline's PDR keeps decreasing over distance due to non-resilient multi-hop communication via the V2V channel.

B. Latency

The average latency of CLEO and the baseline approach is visualized in Figure 5. We measure the end-to-end latency, from the application layer of the original sender to the application layer of every vehicle that receives that message. Thus, our end-to-end latency is including all the delay introduced by relays, no matter if the relay is a vehicle or a LEO satellite. If a message arrives at the same vehicle multiple times, we take the minimum. If a message is never received by a given vehicle, this does not impact our metric.

Again, three major observations can be made. First, the average end-to-end latency within the one-hop V2V communication distance is the minimum because the propagation delay of two vehicles directly communicating with each other cannot be further reduced.

Second, the average end-to-end latency of the baseline approach linearly increases with increasing distance, while CLEO's average end-to-end latency keeps constant after a distance of 4 km. For the baseline approach, the end-to-end latency of a MSG constantly increases with each hop the MSG has to take. Regarding CLEO, each MSG only has to be received once by a CH in order to be relayed via the LEO satellite. Since the satellite's distance to all the vehicles is roughly the same, the Round Trip Time (RTT) of a successful vehicle-to-satellite-to-vehicle transmission is equal for all vehicles. Further, after a successful vehicle-to-satellite transmission, every vehicle can soon after receive the

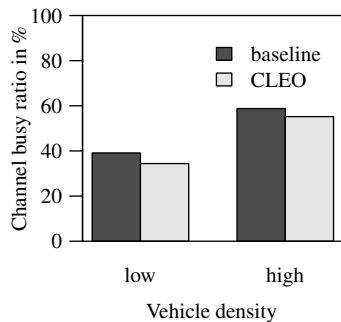


Figure 6. V2V channel busy ratio

MSG. Thus, for distances larger than 4 km, the average end-to-end latency is constant at 7 ms and 8 ms for low and high traffic density.

As a third observation, the average end-to-end latency is higher in the high traffic density scenario compared to the low traffic density scenario. This is because, with more vehicles transmitting MSGs on the V2V channel, back-off counters are decremented slower which increases the V2V channel access latency.

C. Channel Utilization

The last metric that we present is the Channel Busy Ratio (CBR) for the V2V channel, visualized in Figure 6. As a vehicle travels within the ROI, it measures the channel busy time, that is, the time that the Clear Channel Assessment (CCA) at the V2V MAC layer considers the channel to be busy (which is the case when the channel's power level is greater than -65 dBm). The CBR is the ratio of the channel busy time divided by the time a vehicle travels within in the ROI.

Using CLEO slightly reduces the CBR of the V2V channel because MSGs which are transmitted over several hops in the baseline approach, are received by all vehicles once the MSG is successfully relayed by the LEO satellite. Hence, after receiving a MSG from the LEO satellite, no vehicle rebroadcasts the MSG again.

V. CONCLUSION

In this paper, we investigated the benefits of a Low Earth Orbit (LEO) satellite acting as relay that enables large distance communication in vehicular networks. Our approach, called CLEO, integrates a LEO satellite into a cluster-based multi-hop Vehicle-to-Everything (V2X) communication scheme.

We implemented it in a simulation, modeling realistic LEO satellite mobility and accurate Vehicle-to-Vehicle (V2V) communication. In order to put more emphasis on the benefits of a LEO satellite, we used a simplistic approach like ALOHA for managing LEO satellite uplink channel access and still the simulation results show that CLEO achieves major improvements in terms of Packet Delivery Ratio (PDR) and end-to-end latency, especially for long communication distances. In addition, the Channel Busy Ratio (CBR) is slightly reduced. However, the traffic density still needs to be considered in

the design of information dissemination protocols with the utilization of the LEO satellite.

In the future, we will address two aspects. One aspect is implementing more advanced LEO satellite uplink channel access schemes in order to improve the performance of CLEO and be more independent of the traffic density. The second aspect is about improving short range (or even one-hop) communication with the help of LEO satellites, too. In our previous work [13], we presented Satellite Assisted Medium Access Control (SAMAC), which uses a LEO satellite for coordinating sidelink channel access of Device-to-Device (D2D) communication. Investigating the trade-off between relaying and coordinated D2D communication, both enabled by LEO satellites, is an open research question, which could be addressed by combining ideas from CLEO and SAMAC.

REFERENCES

- [1] L. Liu, M. Zhao, M. Yu, M. A. Jan, D. Lan, and A. Taherkordi, "Mobility-Aware Multi-Hop Task Offloading for Autonomous Driving in Vehicular Edge Computing and Networks," *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [2] L. Bedogni, M. Fiore, and C. Glacet, "Temporal Reachability in Vehicular Networks," in *IEEE Conference on Computer Communications (IEEE INFOCOM 2018)*, IEEE, Apr. 2018.
- [3] "In 80 Prozent aller Fälle funktioniert die Rettungsgasse nicht richtig," German Red Cross, Press Release 065/18, Nov. 2018. [Online]. Available: <https://www.drk.de/presse/pressemitteilungen/meldung/in-80-prozent-aller-faelle-funktioniert-die-rettungsgasse-nicht-richtig/>.
- [4] B. L. Nguyen, D. T. Ngo, N. H. Tran, M. N. Dao, and H. L. Vu, "Dynamic V2I/V2V Cooperative Scheme for Connectivity and Throughput Enhancement," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 2, pp. 1236–1246, Feb. 2022.
- [5] O. Kodheli et al., "Satellite Communications in the New Space Era: A Survey and Future Challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 1, pp. 70–109, 2021.
- [6] F. Hagenauer, C. Sommer, T. Higuchi, O. Altintas, and F. Dressler, "Parked Cars as Virtual Network Infrastructure: Enabling Stable V2I Access for Long-Lasting Data Flows," in *23rd ACM International Conference on Mobile Computing and Networking (MobiCom 2017)*, *2nd ACM International Workshop on Smart, Autonomous, and Connected Vehicular Systems and Services (CarSys 2017)*, Snowbird, UT: ACM, Oct. 2017, pp. 57–64.
- [7] G. S. Pannu, T. Higuchi, O. Altintas, and F. Dressler, "Efficient Uplink from Vehicular Micro Cloud Solutions to Data Centers," in *19th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2018)*, Chania, Greece: IEEE, Jun. 2018.
- [8] A. Di Maio, M. R. Palattella, and T. Engel, "Multi-Flow Congestion-Aware Routing in Software-Defined Vehicular Networks," in *90th IEEE Vehicular Technology Conference (VTC 2019-Fall)*, IEEE, Sep. 2019.
- [9] M. Khabbaz, C. Assi, and S. Sharafeddine, "Multihop V2U Path Availability Analysis in UAV-Assisted Vehicular Networks," *IEEE Internet of Things Journal*, vol. 8, no. 13, pp. 10745–10754, 2021.
- [10] B. Kloiber, T. Strang, H. Spijker, and G. Heijenk, "Improving Information Dissemination in Sparse Vehicular Networks by Adding Satellite Communication," in *2012 IEEE Intelligent Vehicles Symposium*, IEEE, Jun. 2012, pp. 611–617.
- [11] C. Cooper, D. Franklin, M. Ros, F. Safaei, and M. Abolhasan, "A Comparative Survey of VANET Clustering Techniques," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 1, pp. 657–681, Feb. 2017.
- [12] Y. Gunter, B. Wiegel, and H. P. Grossmann, "Cluster-Based Medium Access Scheme for VANETs," in *IEEE Intelligent Transportation Systems Conference*, IEEE, Sep. 2007, pp. 343–348.
- [13] M. Franke, F. Klingler, and C. Sommer, "Addressing the Unbounded Latency of Best-Effort Device-to-Device Communication with Low Earth Orbit Satellite Support," in *IEEE Consumer Communications and Networking Conference (CCNC 2023)*, Las Vegas, NV: IEEE, Jan. 2023.