

Topology-Aware Analysis of Context Synchronization in Opportunistic VANETs

Ergys Puka, Peter Herrmann, and Besmir Tola

Department of Information Security and Communication Technology (IHK)

Norwegian University of Science and Technology (NTNU)

Email: ergys.puka@gmail.com, herrmann@ntnu.no, besmir.tola@ntnu.no

Abstract—Vehicular opportunistic networks can extend connectivity services through cellular dead spots, but short-lived and lossy vehicle ad hoc networks leave little time to exchange the forwarding context, risking duplicate transmissions and inconsistent delivery. Using a recently proposed context dissemination protocol, this paper examines whether context synchronization is feasible under realistic ad hoc-network encounter dynamics and how it depends on network topology. In particular, we provide formal model-checker-based verification of behavior related to ad hoc networks in the dead spot dissemination protocol, yielding exact signaling-cost bounds that simulation alone cannot establish. Furthermore, we compare peer-to-peer and star networks, deriving a topology-aware design guideline: For three vehicles, the star organization achieves full context synchronization with roughly 32% fewer signals, while peer-to-peer reaches basic context convergence slightly earlier (5.35 vs. 6.40 signals on average). Finally, mobility-driven simulation of a 300 km dead spot scenario using the simulator SUMO shows that, depending on the traffic volume, 65 to 85% of vehicular ad hoc encounters are pairwise, and a further 13 to 25% involve three vehicles. This establishes that small-group encounters are the dominant regime and that our formal scope therefore covers between 90 and 98% of real contacts.

Index Terms—Opportunistic Networks, Context Dissemination, VANETs, Dead Spot Mitigation, Context-Aware Message Flooding, Opportunistic Forwarding.

I. INTRODUCTION

In environments where stable end-to-end connectivity is lacking, e.g., in sparsely inhabited areas such as the Australian Outback, Opportunistic Networks (OppNets) address communication by relying on the store-carry-forward principle and exploiting transient contacts between mobile nodes. This is beneficial in disconnected or intermittently connected ad hoc settings in which stable connections between the nodes and their environments are impossible [1], [2]. In vehicular scenarios, communications based on OppNets are particularly relevant in cellular dead spots, where vehicles can temporarily lose connectivity and must rely on ephemeral vehicle-to-vehicle (V2V) contacts to disseminate information until some node regains coverage and can deliver pending data [3].

In practical use, however, vehicular OppNets based on ad hoc networks face a key challenge: The contacts can be brief and unreliable, making it difficult to exchange enough control information to coordinate forwarding decisions. The problem can be exacerbated by the overhead and lossy behavior of ad hoc technologies. For example, establishing a connection using the widely available Wi-Fi Direct protocol [4] involves certain

discovery and group formation steps, which can consume several seconds in practice [5], [6]. When combined with mobility-limited contact durations, this can leave just a small residual time budget for exchanging the required information through higher-layer dissemination protocols.

Many opportunistic dissemination strategies assume that nodes can exchange sufficient metadata during encounters. However, the short time frames available may easily lead to packet loss and, consequently, only partial exchanges of metadata. Protocols tend to compensate for this by aggressively replicating messages, which, in turn, can create excessive overhead. In wireless broadcast environments, this can contribute to the well-known broadcast storm problem [7]. Therefore, for vehicular dead spot dissemination, there is a strong need for mechanisms that (1) accelerate delivery by exploiting vehicles that regain infrastructure connectivity earlier and (2) avoid excessive replication and catastrophic delivery failures when mobility behavior deviates from expectations.

To address this setting, we introduced ROCAMFLOOP as an efficient and robust protocol to accelerate message delivery in cellular dead spots [8]. It utilizes ad hoc networks between vehicles to forward messages to several vehicles so that their delivery may be accelerated. The protocol relies on the estimation of dead spot exit times, which is relatively easy with integrating connectivity maps [9] into modern road navigation systems, and limits message delivery to the n earliest expected leavers. Thus, duplicate deliveries are reduced compared to epidemic routing, while the network remains robust to mobility uncertainty [10]. Moreover, to avoid losses when exit-time estimates become stale, e.g., when early leavers get delayed, ROCAMFLOOP guarantees that at least one carrier remains committed to delivering messages even under incomplete ad hoc synchronization. However, the effectiveness of these mechanisms relies on timely knowledge synchronization among vehicles during opportunistic encounters, which is hindered by short contact durations.

In this paper, we study how the choice of ad hoc network organization shapes context synchronization in vehicular OppNets, using ROCAMFLOOP as a representative dead spot dissemination protocol. We contribute the following:

- *Formal verification of dead spot dissemination signaling.* Formal model checking has been applied to opportunistic and ad hoc routing (refer to Sec. V). To our knowledge, however, this is the first application of formal model

checking to ad hoc networks for dead spot dissemination and the quantification of the topology-induced signaling cost. Using the Temporal Logic of Actions (TLA), its modeling language TLA^+ [11], and the model checker TLC [12], we verify that the RoCAMFLOOP-based interaction between vehicles is free of protocol errors. Moreover, we use TLC to derive the exact minimum, maximum, and average signaling costs for full and partial message-context synchronization—bounds that simulation alone cannot establish.

- *Topology-aware design guideline.* We formalize and compare two representative ad hoc organizations, i.e., star coordination (e.g., Wi-Fi Direct) and decentralized peer-to-peer, and quantify a crossover effect: star organizations dominate for full synchronization (32% fewer signals at three peers), while peer-to-peer organizations dominate when only signal carrier enumeration is required. This crossover provides implementers with a principled basis to select a topology against the residual contact budget.
- *Empirical validation of the operational regime.* Through mobility-driven SUMO simulation [13] of a 300 km cellular dead spot, we characterize a realistic contact structure and show that more than 90% of all vehicular ad hoc encounters involve two or three vehicles. We further quantify how contact durations (under four seconds at speeds above 110 km/h) and persistent signaling loss (3% at 1 km range) constrain the feasibility of any context-synchronization scheme in this setting.

While these contributions inform the design of RoCAMFLOOP specifically, they also yield general observations about vehicular ad hoc networks. In vehicular OppNets, small-group encounters dominate, contact budgets are tight, and topology has a measurable, regime-dependent effect on signaling costs.

The rest of the paper is organized as follows. Section II summarizes RoCAMFLOOP and its mechanisms relevant to ad hoc exchange. Section III describes star-based and peer-to-peer ad hoc topologies and illustrates their signaling behavior with representative synchronization scenarios. Section IV reports the SUMO-based simulation study of contact opportunities and synchronization feasibility. It also presents the TLA^+ /TLC-based verification and signaling-cost comparison. Section V reviews related work in opportunistic routing and vehicular dissemination, followed by a conclusion in Section VI.

II. RoCAMFLOOP

As mentioned above, RoCAMFLOOP [8] utilizes ad hoc networks between vehicles in a cellular dead spot to accelerate the delivery of messages to their recipients. When two or more vehicles come close to each other, they establish an ad hoc network and utilize it to exchange stored messages. Over time, the messages are then spread among various vehicles in the dead spot. Since the first vehicle to regain cellular network coverage can deliver a stored message, RoCAMFLOOP decreases the average message delivery delay by more than 43% in large dead spots extending over 300 kilometers, compared to always delivering messages through their creator [8].

In contrast to fully epidemic protocols [10], RoCAMFLOOP markedly reduces the delivery of unnecessary message copies to the recipients. This is accomplished by restricting message delivery to the n carriers anticipated to leave the dead spot first. The value $n \in \mathbb{N} \setminus \{0\}$, which is called *robustness policy* [8], can be individually set for each message. A robustness policy with a larger value of n allows the protocol to cope better with possibly unexpected drivers' behaviors, such as making unplanned stops or deviating from the planned route, which may influence the dead spot leaving time. Conversely, using a smaller value of n will reduce the number of message copies.

To achieve the aforementioned functionality, RoCAMFLOOP uses certain protocol mechanisms that are described in detail in [8]. In the following, we discuss the two protocol mechanisms that are most relevant in the context of ad hoc networks between vehicles.

A. Ranking Table

A vehicle veh carries a separate ranking table for each message m it stores. A ranking table is realized as a set variable consisting of entries for each vehicle veh_o that also contains a copy of m according to the knowledge of veh . In an entry, an identifier of veh_o , its supposed dead spot leaving time t_l , and the freshness of this information, i.e., the point in time when t_l was last updated, are stored.

To predict dead spot leaving times, we suppose that vehicles can access so-called connectivity maps [9], in which the extensions of dead spots are listed. Then, we can utilize modern road guidance systems that are nowadays available in many cars and mobile phones. These systems can estimate the time a certain point on the route will be reached quite precisely. The road guidance system can be applied to determine when the edge of a dead spot, according to the connectivity map, will be reached. The calculated time can then be used as the dead spot leaving time t_l .

If a message m is generated in a vehicle veh , an initial ranking table is created that contains only a single entry for veh itself. When two or more vehicles form an ad hoc network, each one tries to transfer the stored messages along with their ranking tables to its peers. The receiver of a transmission aggregates the entries of its own ranking table with the received one so that the entries always contain the freshest information available.

As soon as veh leaves the dead spot, it compares the current point in time with the leaving times in the various entries of the ranking tables. A message m is sent only if there are fewer than n entries of other vehicles with earlier leaving times listed in the ranking table. In this way, the delivery of a massive number of message copies (sometimes more than a hundred, according to the simulations introduced in [8]) is prevented.

B. Guarantee Flag

The ranking tables are sensitive to the simultaneous delay of several vehicles. Given a robustness policy n for message m , let us assume that the $n + 1$ vehicles expected to exit

the dead spot first decelerate. Yet, they cannot communicate their speed reductions to one another, as they are too far apart to establish an ad hoc network. Then, each vehicle regaining cellular connectivity falsely assumes that m was already delivered by n other vehicles and therefore refrains from sending it to the recipient itself, so that m will not be delivered at all. Simulations presented in [8] show that this may happen for more than 10% of all messages using the robustness policy $n = 1$.

To alleviate this weakness, each ranking table entry is also provided with a guarantee flag, which is a Boolean variable, and a freshness value that indicates when the guarantee flag was changed the last time. If the guarantee flag is set to *true* for a vehicle veh in the ranking table of a message m , veh must deliver m independently of whether it is one of the n fastest dead spot leavers or not, according to its ranking table. To reduce the delivery of unnecessary message copies, the guarantee flag is usually set to *true* only for vehicles that are expected to leave the dead spot early anyway.

To guarantee a proper functionality, the flag has always been set to *true* for at least one bearer of the message m , which is ensured by the protocol functionality that we explain in the following: When a vehicle veh creates a new message m , its guarantee flag in the ranking table for m is set to *true* since veh is initially the only bearer of m . When veh participates in an ad hoc network and aggregates ranking table data from its peers, it sets its guarantee flag to *true* if it is supposed to leave the dead spot earlier than all other vehicles in the ranking table, and to *false* otherwise. However, if veh 's guarantee flag is currently set to *true*, it may toggle it to *false* only after receiving a ranking table from a peer veh_o that confirms that veh_o 's guarantee flag is also set to *true*. In [8], the handshake procedure is explained in detail.

The aforementioned procedure incurs the cost that more than one vehicle bearing a message m might have set their guarantee flags for m to *true*. That may happen when a handshake cannot be completed since the ad hoc network has broken down before the confirmation signal could be delivered. The result is a potential increase in the number of delivered copies of m to its recipient. However, we demonstrated in [8] that the increase in message deliveries is negligible since the average number of messages increases due to the guarantee flags by less than 0.2, e.g., from 2.41 to 2.51 for a robustness policy $n = 2$. This is considered far less problematic than the potential loss of messages due to vehicle delays.

III. AD HOC NETWORKS BETWEEN VEHICLES

The most prominent types of topologies for ad hoc networks connecting passing vehicles are star networks and peer-to-peer networks. In the following, we use the term “signal” to refer to protocol data exchanges, such as the data needed to achieve context synchronization among peers (alias information bearers), and “message” for informational data exchanged between road users and the environment; for example, data messages exchanged to inform about hazardous situations on roads.

A. Star Networks

This network type follows the so-called spoke-hub distribution paradigm, in which each network station is connected to a central hub that usually carries routing functionality.

A prominent ad hoc network example is the Wi-Fi Direct protocol [4], [14], which is available on nearly all modern Android and iOS devices; see [15], [16]. Wi-Fi Direct is based on the IEEE 802.11 WLAN protocol and offers Wi-Fi functionality through the presence of a soft access point. One of the peers, the so-called *Group Owner* (GO), is assigned the role of the soft access point (or hub). It offers its peers, the *clients*, the necessary WLAN functionality in the form of a Basic Service Set.

Wi-Fi Direct works as follows: A device that wishes to connect with others starts a discovery process. If it detects an already existing network, this network is simply joined. If such a network does not yet exist, but other devices in the vicinity are found that are also interested in establishing one, the so-called *group formation process* starts. During this establishment, the GO role is dynamically assigned to one of the devices, while the others act as clients.

To illustrate the functionality of using Wi-Fi Direct in ROCAMFLOOP, we show an example as a sequence diagram in Fig. 1. It describes an ad hoc network connecting vehicles A, B, and C, with B selected as the Group Owner (GO), while the other two vehicles are clients. Vehicle C is supposed to leave the dead spot first, followed by B, and then A. Initially, a message m is stored only in A. The diagram shows the signals exchanged between two or all three vehicles in the ad hoc network. They are provided with numbers indicating the order of their appearance. The signals are used to provide the vehicles with the correct ranking tables and guarantee flag settings. We use red lifelines to convey that the guarantee flag of a vehicle is set to *false*, while green ones indicate that it is *true*. The hexagonal markers outline the information of the ranking table for m stored in a vehicle after it has carried out an aggregation step. We also color the vehicle designators in red and green to indicate their current guarantee flag settings.

An aggregated ranking table must be forwarded to peers in the ad hoc network when the new ranking table contains information not yet known to the peer, see [8]. In star networks, information must be further forwarded by the hub to fulfill its routing obligations. In the Wi-Fi Direct solution, the GO keeps copies of the clients' ranking tables. It can then compare its own aggregated ranking table with those of the clients and forward its table to clients that lack relevant information.

A client forwards its aggregated ranking table to the GO if it contains information that the GO does not seem to possess. Using the model checker TLC, we could prove that this forwarding scheme is sufficient to ensure correct ranking tables and guarantee flag settings after exchanging all necessary signals. By using filled arrowheads in the aggregation steps depicted in Fig. 1, we indicate that the vehicle carrying out the aggregation has to forward it to other ad hoc network members, whereas non-filled arrowheads indicate that no for-




Fig. 1. Scenario with three vehicles using Wi-Fi Direct.

warding is necessary.

In the scenario illustrated in Fig. 1, eight signals are used to synchronize all three vehicles. In signal 1, client A sends message m , including its ranking table to group owner B, which aggregates the received ranking table by adding its own information. Since A's guarantee flag is *true* and B's proposed dead spot leaving time is earlier than A's, B further sets its guarantee flag to *true*. Thereafter, it broadcasts the message and the aggregated ranking table to both clients within signal 2.

In the aggregation following the reception of signal 2, vehicle A discovers that, along with vehicle B, there is another vehicle with an earlier leaving time than its own, which has its guarantee flag for m set to *true*. Therefore, A can now set its own guarantee flag to *false*. Since the peers shall be informed about this toggle, A sends the amended ranking table to the GO in signal 3. After receiving the broadcast signal 2, C stores the message m and aggregates its ranking table by adding information about itself. Since C is expected to leave the dead spot first, it sets its own guarantee flag to *true* and notifies the GO about that using signal 4.

The group owner B aggregates the ranking table from A first, which was received in signal 3. Since, according to its knowledge, A, but not C, has the newest information, the aggregated ranking table is forwarded only to C using signal 5. We assume that signals 4 and 5 pass each other, so B can aggregate signal 4 only after having sent signal 5. The aggregation of signal 4 reveals that B can toggle its own guarantee flag to *false* as well, since it now knows about vehicle C, which is supposed to leave the dead spot before B and has its guarantee flag set to *true*. This information is broadcast to both clients using signal 6. Vehicle C aggregates the ranking table received in signal 5. When B created signal 5, it was not yet aware of the information about C itself, which only arrived through signal 4. Therefore, C sends signal 7 to B, which crosses signal 6 from B. Signal 7 carries ranking table data about all three vehicles, but not about the correct guarantee flag setting of B, of which C was not yet informed.

After receiving signal 6, vehicle A aggregates its ranking table. Since it detects that it does not contain any new information that B is unaware of, it does not need to forward its ranking table anymore. The same holds for C after receiving signal 6. When B aggregates C's ranking table received in signal 7, it finds, however, that C did not know about B's guarantee flag setting. Therefore, it broadcasts the newest ranking table to the clients using signal 8. Both clients aggregate this broadcast signal but do not need to respond to it further, so the interaction for message m is complete.

The scenario in Fig. 1 illustrates the opportunistic nature of RoCAMFLOOP, as the system only replies to signal exchanges when potentially novel information is exchanged. However, the example also clarifies that RoCAMFLOOP does not guarantee that the data exchanges are always at a minimum. For instance, the fact that C sends its signal 7 before considering signal 6 causes B to resend ranking table data using signal 8 which is effectively unnecessary. However, C cannot know if or when there will be any inputs from the GO before the network collapses. Therefore, it is better to send signal 7 out as soon as possible, even if that may lead to some unnecessary signal exchanges.

B. Peer-to-Peer Networks

In a peer-to-peer ad hoc network, none of the participating peers has a particular management functionality that is offered to the others. Instead, network control is decentralized among the different participants that provide similar functionalities.

In Fig. 2, we outline the scenario already discussed in Sect. III-A, but this time based on a peer-to-peer network. Using TLC, we could prove that it is sufficient to ensure that the ranking tables and guarantee flags of all three vehicles are properly synchronized over time when each peer broadcasts its first signal and the one indicating a change of its guarantee flag to all other peers. Otherwise, a peer needs to respond only to the partner that sent a signal when the aggregation reveals that this partner has outdated information. This indicates that all partners must be informed about the mere existence of a peer, as well as any changes to its guarantee flag, while other




Fig. 2. Scenario with three vehicles using peer-to-peer.

inconsistencies with a particular partner can be solved using unicast communication with that station.

In our scenario depicted in Fig. 2, signals 1, 2, and 4 are broadcasts since they are the first signals issued by their transmitters. Signals 3 and 8 are also broadcasts, indicating that their senders recently toggled their guarantee flag. In contrast, signals 5, 6, 7, and 9 are unicast messages reacting to earlier input from the signal recipients.

A first comparison with the Wi-Fi Direct-based solution shows that, in the peer-to-peer network, nine signals, instead of eight in the star network, are needed to synchronize the information for message m . Since the number of broadcasts carried out is also larger (five instead of three), the three peers have to carry out 14 ranking table aggregations compared with only 11 in the Wi-Fi Direct-based scenario. The reason for this discrepancy is that the synchronizing effect of the GO is missing, which requires some further signal exchanges. In the

following, we discuss this issue in greater detail.

IV. EVALUATION

In this section, we discuss relevant results from simulations with SUMO, refer to findings from our formal model-checking runs, and provide more general system-level performance data.

A. Average Sizes of the Ad Hoc Networks

While cellular dead spots can occur nearly everywhere, their impact is strongest when they are large and stretch for hundreds of kilometers. Such extensive dead spots exist in (1) large, sparsely populated areas, which can be found, e.g., in Canada and Australia, or (2) in troubled areas where the cellular infrastructure may be destroyed due to, e.g., war or earthquakes. To obtain meaningful data about this, we applied the mobility simulator SUMO [13] and created various road segments of up to around 300 kilometers, without any cellular network access and with only a single place offering Wi-Fi access to stopping vehicles. By adding parking lots and side roads leading to very remote places, we also modeled unreliable driver behavior, such as pausing unexpectedly or deviating from planned routes. Thanks to published traffic data from a remote region [17], we modeled a realistic traffic flow.

First, we look at the average number of vehicles that form ad hoc networks, as well as their mean durations. The simulation results, which assume a fixed coverage range of 200 meters, as offered by Wi-Fi Direct, are depicted in Fig. 3.

On the left side, we list the distribution of the peer numbers depending on the traffic density. Here, a traffic density of 1.0 corresponds to around 340 vehicles a day on average, which is typical for a country road in a sparsely populated area [17]. The other values refer to simulations in which these traffic values are doubled, tripled, etc. The bar chart reveals that most ad hoc networks consist of only two peers. This percentage is between 65 and 85% depending on the simulated traffic density. Between 13 and 25% of the ad hoc networks consist of three members, while only between 2 and 7% of them are composed of four peers. Even larger ad hoc networks are quite rare, and depending on the traffic density, between 90% and 98% of all ad hoc networks comprise two or three vehicles.

It is not surprising that increased traffic densities lead to larger ad hoc networks on average since the likelihood of more than two vehicles coming close to each other at a certain point in time is increasing. Therefore, the share of just two peers in an ad hoc network decreases with growing traffic density, whereas the percentages of a larger number of peers increase.

On the right side of Fig. 3, we illustrate the distribution of the ad hoc network durations depending on the maximum speeds used by the vehicles. The durations get shorter the faster the vehicles move. For a speed of 50 km/h, most ad hoc network durations are between five and 13 seconds, while few networks persist longer than four seconds when the vehicles are allowed to operate at 130 km/h. This is due to the fact that, on sparsely used roads, few vehicles drive in a convoy, and most encounters occur when vehicles traveling in opposite directions meet each other. In this case, the vehicles will be shorter in the vicinity of each other when they move faster.




Fig. 3. Distributions of (1) the number of vehicles in an ad hoc network depending on the traffic density (left side) and (2) the durations of ad hoc networks depending on the vehicle speeds (right side).




Fig. 4. Signal and message losses depending on the vehicle speed and the range of the ad hoc networks.

B. Signal Loss Averages

The often very short durations of the ad hoc networks discussed above pose the question of how this circumstance affects the behavior of our ROCAMFLOOP protocol. To get answers, we analyze, by means of SUMO-based simulations, how often signal losses actually occur in a simulated system. In Fig. 4, the percentages of data losses depending on the vehicle speeds and the ranges of the underlying network technology are illustrated. A filled bar refers to lost ad hoc network signals, whereas an unfilled one describes the share of ranking table information lost due to ad hoc network collapses. As expected, shorter network ranges and higher maximum speeds of the involved vehicles lead to more signal errors, as they result in shorter ad hoc network durations.

Even with a long range of one kilometer, however, a significant signal loss rate of 3% exists. The main reason is that a new synchronization of the message data is started in ROCAMFLOOP when one of the peers in an ad hoc network detects a significant deviation from its anticipated dead spot leaving time due to unexpected breaks, etc. Of course, this process can occur shortly before the network collapses, so not all signals of this synchronization cycle can be exchanged.

An interesting result is that, with the exception of vehicles

TABLE I
MINIMUM, MAXIMUM, AND AVERAGE NUMBERS OF AD HOC NETWORK SIGNALS NECESSARY TO SYNCHRONIZE THE NETWORK PEERS (1) FULLY AND (2) THEIR RANKING TABLES WITHOUT THE GUARANTEE FLAGS.

Topology	Parameter	Full synchroniz.		All bearers in RTs	
		2 peers	3 peers	2 peers	3 peers
Star network	Minimum	2	4	2	4
	Maximum	3	11	2	8
	Average	2.50	8.54	2.00	6.40
Peer-to-peer	Minimum	2	6	2	3
	Maximum	3	15	2	7
	Average	2.50	12.58	2.00	5.35

traveling at least 110 km/h in the network with a range of 200 meters, the percentage of lost end-user messages is only a fraction of that for the lost ad hoc network signals. This demonstrates the high robustness of ROCAMFLOOP against collapsing ad hoc networks, but it also indicates that the range of 200 meters offered, e.g., by Wi-Fi Direct, is a tad too short for the higher speeds of the vehicles.

C. Comparing Star and Peer-to-Peer Networks

We complemented the SUMO-based simulations by applying the TLA⁺-based model checker TLC [12]. In particular, we modeled the ROCAMFLOOP-based signal exchange process for both star and peer-to-peer networks in TLA⁺ [11]. This allowed us to formally verify with TLC that the policies for forwarding aggregated information ensure full synchronization of the ranking tables and guarantee flags among all peers in ad hoc networks of two and three vehicles, assuming that the network persists for a sufficient duration.

Unfortunately, we could not verify ad hoc networks with four or more vehicles, as that led to a state space explosion¹ in model checking. As discussed in Sect. IV-A, however, few ad hoc networks consist of more than three vehicles (depending on the traffic density, between 2 and 10%). Therefore, verifying these scenarios is of less practical relevance.

We determined the minimum, maximum, and average number of signals required for (1) full synchronization of all ad hoc network members and (2) synchronization of only their ranking tables, but not their guarantee flag settings. The results of both measurements are listed in Tab. I. They do not contain the data exchange necessary to discover and build connections.

For ad hoc networks consisting of just two peers, there is no difference between star and peer-to-peer networks since their behavior is effectively identical in both topologies. For full synchronization, two or three signals have to be exchanged depending on whether the original holder of a message is the vehicle that is supposed to leave earlier than its peer or not. Both cases are similarly likely, so the average is 2.5. Moreover, the ranking tables contain both involved vehicles after transmitting just one signal in each direction, so the average is 2.0 in this case.

¹For example, a run using peer-to-peer technology crashed after checking 6,143,464,226 distinct states, with a runtime of more than 39 hours on a state-of-the-art PC with 14 cores.

A significant difference between the two network topologies can be seen for ad hoc networks consisting of three vehicles. The star network has a clear advantage in achieving full synchronization of the ranking table information. Here, 8.54 signals have to be exchanged on average, compared with 12.58 in peer-to-peer. Thus, the number of data exchanges for the star network is approximately 32% lower relative to peer-to-peer. In contrast, peer-to-peer networks are more beneficial when we consider only the signals until all peers in the ad hoc network have entries for all others in their ranking tables, albeit with potentially incorrect guarantee flag settings. In this category, the star network provides a disadvantage of almost 17% compared to peer-to-peer.

The likely reason for this diverging behavior is that the initial broadcast signals of the peer-to-peer networks significantly help share the original information among the peers, while in the star networks, all information between the clients has to go through the hub, causing additional signals. After the initial sharing of information between the vehicles, however, the aggregation of messages in the hub helps to reduce the number of signals needed to synchronize the guarantee flag settings. The direct interaction between the peers may lead to quite sluggish synchronization, as exemplarily shown towards the bottom of Fig. 2.

D. Handling Ad Hoc Network Collapses

As shown in Fig. 3, ad hoc networks are typically short-lived, while two-to-three-vehicle synchronization may require between 2 and 15 signals; see Tab. I. Therefore, signal loss is unavoidable, and SUMO simulations report around 13% lost packets at 110 km/h and a 200 m range [8], so any protocol on top must tolerate sudden collapses.

ROCAMFLOOP has the necessary robustness. As long as drivers act as expected, an ad hoc collapse never prevents cellular delivery. A vehicle *veh* that misses a peer signal may, at worst, remain unaware that another vehicle could leave the dead spot earlier, resulting in an extra delivery.

The guarantee flag handshake is similarly resilient. Failing to complete the handshake to switch a flag to *false* may leave several vehicles with their flag at *true*, causing redundant but never missing deliveries; see Sect. II-B. ROCAMFLOOP also limits the consequences of collapses by minimizing signaling: A vehicle retransmits its ranking table only when it carries strictly fresher or novel information for its peers.

The choice of which of the two topologies introduced in Sect. III is better suited for ROCAMFLOOP is not clear-cut. As mentioned above, the likelihood of ad hoc network collapses grows when more signals have to be transmitted between the peers. We discussed in Sect. IV-C that peer-to-peer networks need fewer signals to synchronize the ranking tables, while the star topology requires fewer signals for full synchronization, including the guarantee flags. Thus, it is not trivial to proclaim a winning topology. By considering the TLC model checker results depicted in Tab. I, however, the star network seems slightly better suited for ROCAMFLOOP, since it synchronizes the full set of data with 32% fewer signals

on average. This outpaces the approximately 17% advantage that peer-to-peer networks have in forwarding the pure ranking tables to the other peers.

E. System Level Performance

To relate our signaling results to end-to-end dissemination performance, we summarize key performance indicators of ROCAMFLOOP based on the same SUMO model for a dead spot stretching 300 km used above [8]. ROCAMFLOOP achieves a delivery probability above 99.4% for the lowest robustness policy $r_m = 1$ and nearly 100% for $r_m \geq 2$, compared to 70% for the well-known protocols Spray-and-Wait [18] and ProPHET [19]. The overhead ratio of message delivery copies varies between 1.79 for $r_m = 1$ and 4.11 for $r_m = 4$. In comparison with the 110 message deliveries for epidemic routing-based protocols [10], this results in 26 to 61 times fewer message copies. In addition, the protocol facilitates the reduction of mean delivery times by over 43%. These values are protocol-level baselines for ROCAMFLOOP based on the star topology in the form of Wi-Fi Direct. They consider the appearance of signal losses due to transient ad hoc network encounters.

The contribution of this paper is to isolate the role of ad hoc topology in the synchronization process underlying these outcomes. Prior work quantifies the resulting exposure: at $r_m = 1$, avoidable non-deliveries rise from below 1.4% when the guarantee flag handshake completes to 10–19% when it fails [8]. Since the star organization completes full synchronization, including the guarantee flag handshake, with fewer signals than the peer-to-peer organization for three peers (see Sect. IV-C), it is less exposed to the contact-window shortening discussed in Sects. IV-A and IV-B. Therefore, it is better suited to maintain the low-redundancy and low-non-delivery regime reported in [8].

V. RELATED WORK

Opportunistic networks, often framed as *Delay Tolerant Networking* (DTN), deliver data under intermittent connectivity via store-and-forward, as formalized in key architecture works [20], the DTN routing problem formulation [21], and synthesized in related routing taxonomies and surveys [22]. A baseline strategy is flooding (Epidemic Routing [10]), which, however, creates prohibitive replication in wireless vehicular settings. Consequently, many protocols either constrain replication (e.g., Spray-and-Wait [18] and its variants [23]) or use encounter-history prediction (e.g., ProPHET [19]) to balance delivery probability versus overhead.

In vehicular DTN, forwarder selection is frequently coupled with road topology, trajectories, or sparse-contact structures, e.g., VADD [24] and Geo-Opps [25]. The broader design space is surveyed in comprehensive vehicular DTN overviews [26]. Most of this work concentrates on which carriers should forward, and encounter-time metadata exchange is typically treated as a secondary assumption.

In contrast, our work targets a practical bottleneck for vehicular dead spot dissemination, where protocol-state syn-

chronization can require multiple signaling rounds. However, the contact windows may be extremely short, so that only very few signals can be transmitted before a network collapses.

Formal model checking has been applied to ad hoc and opportunistic routing protocols through statistical [27], actor-based [28], and timed-automata [29] approaches. To the best of our knowledge, however, no prior work has applied formal model checking to dead-spot mitigation protocols, nor has it been used to quantify the signaling cost induced by ad hoc topology choice.

Although Wi-Fi Direct remains the most widely deployed device-to-device standard on commodity hardware, several emerging technologies offer faster connection setups or longer ranges, including C-V2X PC5 [30], 5G NR sidelink (3GPP Rel-16/17), and IEEE 802.11bd [31]. Our topology-level findings are technology-agnostic: Star and peer-to-peer signaling costs depend on logical organization rather than the radio layer.

Building on ROCAMFLOOP's dead spot-specific ranking and guarantee flag mechanisms for robust, low-duplication delivery [8], we characterized how star/hub versus peer-to-peer technology settings affect the minimum and expected signaling needed to reach full or partial synchronization.

VI. CONCLUSION

Vehicular opportunistic networks can extend connectivity services through cellular dead spots, but brief, lossy encounters leave little time to exchange the forwarding context needed for correct dissemination. Altogether, we delivered three results in this paper: First, formal model checking established exact signaling-cost bounds for context synchronization under two ad hoc organizations, which, to our knowledge, is the first such analysis of a dissemination protocol for dead spot mitigation. Second, a topology-dependent crossover emerged: The star organization reaches full synchronization with roughly 32% fewer signals than peer-to-peer at three peers (8.54 vs. 12.58), while peer-to-peer reaches basic context convergence slightly earlier (5.35 vs. 6.40). Third, mobility simulation showed that over 90% of vehicular ad hoc encounters involve two or three vehicles, establishing that the formally tractable small-group regime is also the operationally dominant one.

In the operational regimes we evaluated, star-like encounters are the most reliable default for maintaining consistent context in infrastructure-sparse environments.

REFERENCES

- [1] L. Pelusi, A. Passarella, and M. Conti, "Opportunistic Networking: Data Forwarding in Disconnected Mobile Ad hoc Networks," *IEEE Communications Magazine*, vol. 44, no. 11, pp. 134–141, 2006.
- [2] M. Alajele, R. Doss, and A. Ahmad, "Routing Protocols in Opportunistic Networks — A Survey," *IETE Technical Review*, vol. 35, no. 4, pp. 369–387, 2018.
- [3] Y.-D. Lin and Y.-C. Hsu, "Multihop Cellular: A New Architecture for Wireless Communications," in *IEEE INFOCOM, Conference on Computer Communications*, vol. 3, 2000, pp. 1273–1282.
- [4] Wi-Fi Alliance, "Wi-Fi Peer-to-Peer (P2P) Technical Specification, Version 1.7," 2016, Wi-Fi Alliance Technical Committee P2P Task Group.
- [5] Ø. Sigholt, B. Tola, and Y. Jiang, "Keeping Connected When the Mobile Social Network Goes Offline," in *2019 International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, 2019, pp. 59–64.
- [6] E. Puka and P. Herrmann, "Data Dissemination for Vehicles in Temporary Cellular Network Dead Spots," *International Journal of Cyber-Physical Systems (IJCPS)*, vol. 1, no. 2, pp. 38–55, 2019.
- [7] Y.-C. Tseng, S.-Y. Ni, Y.-S. Chen, and J.-P. Sheu, "The Broadcast Storm Problem in a Mobile Ad Hoc Network," *Wireless Networks*, vol. 8, pp. 153–167, 2002.
- [8] E. Puka, P. Herrmann, and A. Taherkordi, "An Efficient and Robust Protocol to Accelerate Message Delivery in Cellular Dead Spots," *IEEE Trans. on Vehicular Technology*, vol. 75, no. 5, pp. 8534–8549, 2026.
- [9] F. F. Tseng, D. P. Filev, and I. H. Makki, "Vehicular Connectivity Map," U.S. Patent 9 775 128 B2, Sep. 26, 2017, U.S. Patent 9 775 128 B2.
- [10] A. Vahdat and D. Becker, "Epidemic Routing for Partially-Connected Ad Hoc Networks," Department of Computer Science, Duke University, Durham, NC, USA, Tech. Rep. CS-200006, Jun. 2000.
- [11] L. Lamport, *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley, 2003.
- [12] Y. Yu, P. Manolios, and L. Lamport, "Model Checking TLA+ Specifications," in *Advanced Research Working Conference on Correct Hardware Design and Verification Methods*. Springer, 1999, pp. 54–66.
- [13] P. A. Lopez et al., "Microscopic Traffic Simulation using SUMO," in *IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2018, pp. 2575–2582.
- [14] D. Camps-Mur, A. Garcia-Saavedra, and P. Serrano, "Device-to-device Communications with Wi-Fi Direct: Overview and Experimentation," *IEEE Wireless Communications*, vol. 20, no. 3, pp. 96–104, 2013.
- [15] Android Developers. Wi-Fi Peer-to-Peer. [Online]. Available: <https://developer.android.com/guide/topics/connectivity/wifip2p>
- [16] iOS Developer Library. Multipeer Connectivity Framework. [Online]. Available: developer.apple.com/documentation/multipeerconnectivity
- [17] New South Wales Government, "Traffic Volume Viewer," maps.transport.nsw.gov.au/egcomaps/traffic-volumes/index.html#/?z=6, 2026, accessed: 2026-02-27.
- [18] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Multiple-copy Routing in Intermittently Connected Mobile Networks," University of Southern California, Tech. Rep. CENG-2004-12, 2004.
- [19] S. D. Han and Y. W. Chung, "An Improved PROPHET Routing Protocol in Delay Tolerant Network," *The Scientific World Journal*, no. 1, p. 623090, 2015.
- [20] K. Fall and S. Farrell, "DTN: An Architectural Retrospective," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 5, pp. 828–836, Jun. 2008.
- [21] J. Sushant, "Routing in a delay tolerant network," *ACM SIGCOMM Computer Communication Review*, 2004, pp. 145–158, 2004.
- [22] Y. Cao and Z. Sun, "Routing in Delay/Disruption Tolerant Networks: A Taxonomy, Survey and Challenges," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 2, pp. 654–677, 2013.
- [23] L. Yao, X. Bai, and K. Zhou, "An Improved Spray and Wait Algorithm Based on Q-learning in Delay Tolerant Network," in *2024 International Joint Conference on Neural Networks (IJCNN)*, 2024, pp. 1–8.
- [24] J. Zhao and G. Cao, "VADD: Vehicle-Assisted Data Delivery in Vehicular Ad Hoc Networks," in *IEEE INFOCOM*, 2008, pp. 1–5.
- [25] K. Lee, U. Lee, and M. Gerla, "Geo-Opportunistic Routing for Vehicular Networks [Topics in Automotive Networking]," *IEEE Communications Magazine*, vol. 48, no. 06, pp. 164–170, 2010.
- [26] N. Benamar et al., "Routing Protocols in Vehicular Delay Tolerant Networks: A Comprehensive Survey," *Computer Communications*, vol. 48, pp. 141–158, 2014.
- [27] S. Arora, A. Rathor, and M. P. Rao, "Statistical Model Checking of Opportunistic Network Protocols," in *Proceedings of the 11th Asian Internet Engineering Conference*, 2015, pp. 62–68.
- [28] F. Yousefi et al., "VeriVANca: An Actor-based Framework for Formal Verification of Warning Message Dissemination Schemes in VANETs," in *International Symposium on Model Checking Software*. Springer, 2019, pp. 244–259.
- [29] O. Wibling, J. Parrow, and A. Pears, "Automated Verification of Ad hoc Routing Protocols," in *International Conference on Formal Techniques for Networked and Distributed Systems*. Springer, 2004, pp. 343–358.
- [30] R. Molina-Masegosa, J. Gozalvez, and M. Sepulcre, "Configuration of the C-V2X Mode 4 Sidelink PC5 Interface for Vehicular Communication," in *2018 14th International conference on mobile ad-hoc and sensor networks (MSN)*. IEEE, 2018, pp. 43–48.
- [31] G. Naik, B. Choudhury, and J.-M. Park, "IEEE 802.11 bd & 5G NR V2X: Evolution of Radio Access Technologies for V2X Communications," *IEEE access*, vol. 7, pp. 70 169–70 184, 2019.