Tobias Meuser, Björn Richerzhagen, Ioannis Stavrakakis, The An Binh Nguyen and Ralf Steinmetz. *Relevance-Aware Information Dissemination in Vehicular Networks*. In Proc. 19th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), IEEE, 2018, ISBN 978-1-5386-4725-7

Relevance-Aware Information Dissemination in Vehicular Networks

Tobias Meuser*, Björn Richerzhagen*, Ioannis Stavrakakis†, The An Binh Nguyen* and Ralf Steinmetz*
*Multimedia Communications Lab (KOM), Technische Universität Darmstadt, Darmstadt, Germany
Email: {firstName.lastName}@KOM.tu-darmstadt.de
†National and Kapodistrian University of Athens, Greece - IMDEA Networks & UC3M, Spain
Email: ioannis@di.uoa.gr

***Abstract*—As a step towards fully autonomous driving, Advanced Driver Assistance Systems provide convenience- and safety-related functions to drivers. In addition to data gathered by local sensors, these systems rely on events generated by other vehicles that need to be disseminated to a potentially large audience. Today, this geocast-functionality relies either on subscriptions covering certain areas (e.g., cities) or on individual route-based subscriptions. While the former exhibits suboptimal precision in filtering, the latter introduces significant complexity and assume that routes are known in advance. We propose a prediction-based assessment of the relevance of events without requiring prior route knowledge. Relevance is modeled based on the street network and spatio-temporal characteristics of events. We evaluate our approach in a realistic city setting, relying on the SUMO vehicular mobility simulator. Our first results show that relevance-aware information dissemination reduces the communication overhead by 68%, while at the same time achieving near perfect recall compared to route-based subscriptions.**

***Index Terms*—information dissemination, vehicular networks, information quality, relevance assessment**

I. Introduction

Advanced Driver Assistance Systems (ADAS) provide features to increase the convenience and safety of drivers [1]. These systems heavily rely on information collected by local sensors. However, as the range of local sensors is limited, vehicles exchange detected events with each other to increase their awareness beyond their local perception. Examples for such events are updates on road conditions and information about blocked lanes due to jams or accidents. Disseminating the respective events efficiently is a challenging task, involving the decision to which vehicles a specific event should be disseminated. This decision is influenced by several factors, such as the location and type of the vehicle and the expected temporal validity of the event.

In the literature, vehicles distribute events in a vehicular network using geocast-based approaches [2], relying on either area-based or route-based approaches. If the geocast is area-based, a fixed area is defined, in which the event is considered relevant. The size of this area considers event-specific parameters like the event lifetime. However, properties of the vehicles and current traffic conditions are not considered in such approaches. Consequently, events may be distributed to vehicles that will never pass the event location, leading to unnecessary data traffic. On the other hand, route-based geocast considers the future routes of vehicles to distribute information to concerned vehicles. Besides adding significant complexity to the system, the assumption that each vehicle is fully aware of its future route does not hold. Even considering autonomous driving, we still expect to have mixed traffic with both human-driven and autonomous vehicles [3]. For human-driven vehicles, the vehicle may not know its future route. This limits the applicability of the route-based approach for information dissemination.

In this work, we propose assessing the relevance of an event for an efficient event dissemination. Thereby, we limit the dissemination of events to the vehicles that are concerned the most, e.g. by using a threshold. This threshold can be determined using the information-specific properties like accuracy, which we investigated on in our previous work [4].

978-1-5386-4725-7/18/$31.00 ©2018 IEEE

II. Scenario Overview

In this paper, we focus on the dissemination of information in the form of events in the vehicular context. These events are gathered using the on-board sensors of vehicles driving on the streets. We assume that the sensed events are located on one of the road segments within our street network. These road segments are taken from map databases like OpenStreetMap and are derived by vehicles from their current location. To support the driving efficiency of other vehicles, the vehicles share their detected events with interested vehicles.

However, a detected event on a given road segment is usually not relevant to all vehicles in the network. Current approaches distribute events to all vehicles in a specific area using geocast with or without central infrastructure. This dissemination method only depends on the type of the distributed event and does not take the relevance of the event for the respective vehicle into account. This relevance considers vehicle specific properties like the future route. In this work, we focus on reducing the number of unnecessarily notified vehicles while preserving the correctly notified vehicles. For the distribution of events, we extend the existing Publish/Subscribe (Pub/Sub) system Bypass.KOM [5] to distribute events based on a relevance value.

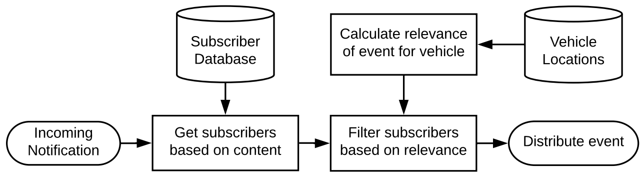

Fig. 1: Processing of an incoming event at the broker.

A. The Publish/Subscribe System

We use an attribute-based Publish/Subscribe (Pub/Sub) system with a central broker entity to distribute incoming events. This broker entity is a pure forward-based broker, i. e., the broker has no storage capabilities for a delayed distribution of events. Once a vehicle senses an event, it sends the event to the broker for further distribution. The broker then decides which vehicles need to be notified. As our scalable information dissemination system distributes events based on the location of the subscribers, the broker requires location information of all vehicles in the system. To this end, the broker monitors the location of those vehicles via the cellular network.

B. Relevance-Aware Filtering

Compared to distribution approaches from the literature, we do not distribute events in a fixed area or over a certain amount of road segments, but use the temporal and geographical relevance of an event to distribute it. This relevance value is not binary 0 or 1 like in existing Pub/Sub systems, but may be any value between 0 and 1. We will explain the relevance assessment in section III.

Once a message arrives at the broker, the broker selects the concerned vehicles as shown in Figure 1. First, the broker requests the current subscribers of that message based on the message content. The returned list of subscribers contains all subscribers regardless of their distance to the event. In the second step, the broker performs additional filtering of the subscribers based on their relevance. This relevance considers location of the vehicle and event-specific properties. As the relevance value is not binary, the broker needs to decide on the concerned vehicles based on fuzzy logic. Simple approaches to that issue are probabilistic and threshold-based approaches. The usage of the relevance limits the dissemination area and protects the network from getting overloaded. In the last step, the broker distributes the event to the selected subscribers.

III. Information Relevance Assessment

In our system, a central broker calculates the relevance of events for the vehicles in the network. The relevance of the event for a close vehicle is generally high, while decreasing with increasing distance. The event is most important for the next vehicle arriving at the event's location, as this vehicle will not receive any updates of the event state from other vehicles anymore. However, due to the lack of route knowledge, the next vehicle cannot be determined with certainty in most situations. Thus, we utilize knowledge about traffic flows to predict the future movement of the vehicles.

Consequently, our goal is to ensure that every vehicle receives an event only once and has received it before entering the route segment the event is located. To achieve this goal of receiving every event only once per vehicle, we calculate the probability that the vehicle will be the next vehicle at the event location. There are two factors influencing this probability. The first factor is the path probability of the vehicle encountering the event. The second factor is the temporal validity of an event.

a) Path Probability: If the route of a vehicle is known, the probability of the vehicle encountering the event is either 0 or 1. However, most commonly the broker will have no information about the vehicles' routes. In this case, the broker predicts the routes considering the general traffic flow. The knowledge about the traffic flow can be obtained using methods proposed in the literature, e. g. the method of Xue et al. [6]. In our work, we predict the path using a Markov chain of order 1, using the road segments of the scenario as states of the Markov chain.

b) Temporal Probability: Events typically have limited lifetime. In order to prevent unnecessary event distribution, the broker should not distribute events to distant vehicles, which will possibly reach the location of the event after the expiration of the lifetime of the event. In a real-world scenario, the exact lifetime of an event cannot be ascertained and is subject to further research. However, the broker can estimate the lifetime of some event by observing past events of the same type. Mostly, the lifetime of an event is not static, but varies even for events of the same type. We model the different lifetimes of events using a lifetime function. This lifetime function $f(t)$ has a value range between 1 and 0 and states the probability that an event is still valid after the time t. Using this function and the expected travel duration, we can estimate the probability that the event is still valid once a vehicle at a certain position arrives at the event location.

c) Information Relevance: In this section, the route probability and the temporal probability are combined to calculate the relevance of an event for each vehicle. As we only inform the next vehicle arriving at the event location, we need to determine the probability that a specific vehicle reaches the event location next. Moreover, the event needs to be valid at this time. For our relevance assessment method, we require all possible routes of all vehicles and the lifetime function of the event as input and determines the relevance value between 0 and 1 for each vehicle. We prune all routes with a probability $\leq 0.1\%$ due to computational complexity. Thus, we limit the computational complexity of the problem to be solved.

We calculate the relevance of an event for a vehicle driving along the investigated route by combining the temporal and geographical properties of the route. The geographical relevance is the probability that a vehicle arrives first at the location of the event based on the current traffic environment. The temporal relevance is the probability that the event is still valid once the vehicle arrives at the event location. As the exact event lifetime and the vehicle route are not known, the temporal relevance is

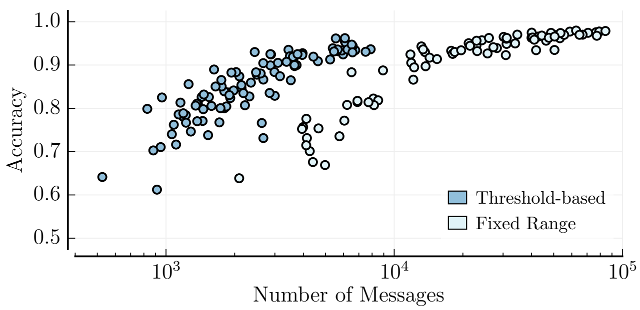

Fig. 2: Produced cellular traffic and achieved distribution quality without local communication. Our approaches use significantly less network resources.

calculated for each route individually and combined using the respective route probabilities.

IV. Preliminary Results and Discussion

We implement our scenario in the event-based Simonstrator framework [7], as it enables evaluation of communication scenarios using cellular and ad-hoc communication technologies and supports the Publish/Subscribe paradigm. We added a connection to SUMO [8] using the Traffic Control Interface (TraCI) to study our system under realistic vehicular movement. To this end, we use the TAPAS cologne scenario [9] being one of the largest freely available datasets. We investigate an area of $2km \times 2km$ for one hour, executing every setup with 20 different random seeds. To allow the system to gather historical data for the prediction, we do not produce any events during the first 30 minutes. After that time, an event is spawned randomly on one of the road segments in every time interval. The size of this interval is chosen dependent on the event lifetime.

To evaluate the performance of our relevance assessment method, we developed a threshold-based approach, which distributes events to all vehicles exceeding a certain relevance threshold. To assess the benefits of our approach, we compared it to a simple geocast mechanism as a baseline.

Figure 2 displays the performance of our algorithm. We added multiple simulation runs with different parameter to this plot, classifying them by their produced traffic. On the x-axis, the produced traffic is displayed, while the y-axis displays the achieved distribution quality. For both the baseline approach and our approach, the distribution quality decreases with decreasing traffic. However, our approach outperforms the baseline significantly. For runs with similar mobile traffic, our approach achieves on average a 35% higher quality. Similarly, for similar quality, our approach uses on average 68% less mobile traffic to achieve this quality level. However, the slope of the derivate function decreases for both approaches with increasing traffic, i. e. the cost for additional quality is much less if the current quality is low. This is an issue for both approaches, resulting from the forward-only Pub/Sub. In order to resolve this issue, we will analyze the impact of storing some events at the server in future work.

V. Conclusion

In this paper, we present a context-aware information dissemination approach for vehicular networks. Compared to the state-of-the-art dissemination approaches, our approach does not require a fixed area in which events need to be distributed. It uses the relevance of an event to distribute the event efficiently. Our relevance assessment approach is based on several event-specific and vehicle-specific properties. The considered event-specific properties are the temporal validity, and the vehicle-specific properties are the location of the vehicles. Unlike many interest-based approaches from the literature, our relevance assessment approach does not require the availability of route knowledge. For vehicles without route knowledge, we predict the possible future routes and consider these routes with their respective probability. If a vehicle is considered the first at the event location, the event is distributed to this vehicle. However, the continuity of the relevance compounds the transmission decision. We developed a treshold-based approach to incorporate the relevance in the distribution process. The evaluation shows that our developed approach achieves similar performance as geocast in terms of distribution quality, while simultaneously producing significantly less traffic compared to geocast-based approaches.

In future work, we will investigate on other methods to incorporate the relevance in the distribution process and possible improvements when storing events at the server. Moreover, we will further improve our results using direct Vehicle-to-Vehicle (V2V) communication.

Acknowledgments

This work has been funded by the DFG within the CRC 1053 - MAKI (B1, C2); the LOEWE initiative (Hesse, Germany) within the NICER project; and by the Chair of Excellence UC3M - Santander Program.

References

[1] A. Tigadi, R. Gujanatti, A. Gonchi, and B. Klemsscet, "Advanced driver assistance systems," *IJERGS*, vol. 4, no. 3, pp. 151–158, 2016.

[2] F. Li and Y. Wang, "Routing in vehicular ad hoc networks: A survey," *IEEE Vehicular Technology Magazine*, vol. 2, no. 2, 2007.

[3] P. Bhavsar, P. Das, M. Paugh, K. Dey, and M. Chowdhury, "Risk analysis of autonomous vehicles in mixed traffic streams," *Journal of the Transportation Research Board*, no. 2625, pp. 51–61, 2017.

[4] T. Meuser, P. Lieser, T. A. B. Nguyen, D. Böhnstedt, and R. Steinmetz, "Adaptive information aggregation for application-specific demands," in *Proc. of BalkanCom*, May 2017. [Online]. Available: ftp://ftp.kom.tu-darmstadt.de/papers/MLA+17.pdf

[5] B. Richerzhagen, "Mechanism transitions in publish/subscribe systems - adaptive event brokering for location-based mobile social applications," Ph.D. dissertation, TU Darmstadt, Aug 2017.

[6] G. Xue, Z. Li, H. Zhu, and Y. Liu, "Traffic-known urban vehicular route prediction based on partial mobility patterns," in *Proc. of IEEE ICPADS*. IEEE, Dec 2009, pp. 369–375.

[7] B. Richerzhagen, D. Stingl, J. Rückert, and R. Steinmetz, "Simonstrator: Simulation and prototyping platform for distributed mobile applications," in *Proc. of SIMUtools*. ICST, 2015, pp. 99–108.

[8] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz, "Sumo–simulation of urban mobility: an overview," in *Proc. of SIMUL*. ThinkMind, 2011.

[9] S. Uppoor, O. Trullols-Cruces, M. Fiore, and J. M. Barcelo-Ordinas, "Generation and analysis of a large-scale urban vehicular mobility dataset," *IEEE TMC*, vol. 13, no. 5, pp. 1061–1075, 2014.