Performance Evaluation of ETSI GeoNetworking for Vehicular Ad hoc Networks

Sebastian Kühlmorgen, Ignacio Llatser, Andreas Festag, Gerhard Fettweis
Vodafone Chair Mobile Communications Systems, Technische Universität Dresden, Germany
{sebastian.kuehlmorgen,ignacio.llatser,andreas.festag,gerhard.fettweis}@tu-dresden.de

***Abstract*—The GeoNetworking protocol provides single-hop and multi-hop communication in vehicular ad hoc networks based on IEEE 802.11p/ITS-G5. It has been standardized by the ETSI Technical Committee ITS as part of its Release 1 set of specifications and is expected to be deployed in the next years. This paper presents a performance evaluation of the GeoNetworking protocol in its recently published version. Our study assesses the performance of the broadcast forwarding algorithms for multi-hop packet transport that are used to disseminate information in geographical areas for road safety and traffic efficiency applications. From the algorithms specification in the standard, we derive six variants with different combinations of protocol mechanisms with increasing complexity and assess their performance in terms of reliability, latency, and overhead. The algorithms are evaluated in a reference freeway scenario with bidirectional road traffic and a realistic trace-based mobility model with varying vehicle density. The obtained results indicate that the combination of contention-based and greedy forwarding shows the best overall performance; further functional improvements have a limited performance gain in the studied scenario.**

I. Introduction

Communication among vehicles and roadside stations is expected to significantly enhance safety and traffic efficiency. This gain is achieved by improving the driver awareness, in-vehicle signage of traffic signs, road warnings and speed advisory, amongst other use cases. In recent years, major efforts have focused on designing a vehicular communication system based on wireless LAN operating at 5.9 GHz frequency range allocated for road safety and traffic efficiency. Several standards have been defined covering vehicular ad hoc networks (VANETs), including messaging protocols, management procedures as well as security and privacy support.

Recently, a Release 1 of vehicular communication standards has been completed in Europe [1]. There, the GeoNetworking protocol resides the core of the European Telecommunications Standards Institute (ETSI) protocol stack and provides the transport of messages over ITS-G5, the European variant of the vehicular communication support in IEEE 802.11-2012 (formerly known as the p-amendment). GeoNetworking is highly tailored to the characteristics of VANETs with sparse and dense distribution of vehicles and frequent topology changes, and meets the requirements of safety, traffic efficiency and driving comfort applications. Being part of the Release 1 of cooperative ITS standards in Europe, it was also adopted in the standards' profile of the Car-2-Car Communication Consortium (C2C-CC) – a European industry consortium of car makers, suppliers and research organizations – for an initial deployment of vehicular communication planned for the upcoming years. Implementations have been tested under real-world conditions in field operational trials (FOTs) and validated by tests for conformance and interoperability.

Geographical routing schemes in wireless networks have been widely studied for several years [4] – [8]. Initially, they addressed general mobile ad hoc networks, later in particular sensor networks and vehicular networks were focused. Their main principle is based on utilizing geographical information for packet transport in an ad hoc network without coordination by a central entity. A destination can either be an individual node or a geographical region. The latter mode allows for efficient distribution of packets inside a geographical area (in short, geo-area), which requires mechanisms to alleviate the well-known broadcast storm problem [9]. Compared to topology-based routing schemes, such as AODV and OLSR, with geographical routing the source node does not need to compute end-to-end routes. Instead, packets can be sent 'on the fly', which makes the routing robust against frequent topology changes, as typically found in VANETs. ETSI GeoNetworking [2] is a standardized variant of geographical routing, derived from existing research approaches and embedded into the ETSI protocol stack.

The GeoNetworking standard has evolved over the past years and finally published as a European norm [2]. The present paper evaluates and compares the broadcast forwarding algorithms of the GeoNetworking Release 1 standard version. Only few publications assess the performance of the standardized solutions [10] – [12], though none of them addresses the latest version of the standard. Our main contributions are as the follows: First, we classify the protocol mechanisms used in the standardized forwarding algorithms and define relevant performance metrics. Second, we evaluate the six forwarding variants, assess the contribution of individual mechanisms to the overall performance, and finally, propose improvements to the standardized algorithms based on the obtained results.

The remaining sections of this paper are organized as follows. Sec. II provides relevant details of the GeoNetworking protocol where Sec. III defines the performance metrics that will be considered for the evaluation and Sec. IV presents the studied scenario, the simulation environment and performance results, followed by the conclusions in Sec. V.

II. ETSI GeoNetworking Protocol

The ETSI GeoNetworking protocol, standardized in the EN 302 636 standard series[1], controls the transport of data packets from a source node to the destination, which can be either an individual node (GeoUnicast), all nodes/any node inside a geo-area (GeoBroadcast/GeoAnycast), or all nodes

[1]ETSI standards are available at http://www.etsi.org/standards.

in a one-hop/n-hop neighborhood (single-hop/topologically-scoped broadcast). Every ad hoc router has a location table that maintains the position of its known neighbors and is used to make forwarding decisions; it also has packet buffers for location service, store-carry-and-forward and forwarding algorithms (see below).

For the safety and traffic efficiency use cases described in Release 1 and the C2C-CC profile, two packet transport types are relevant: first, single-hop broadcast for the transmission of periodic Cooperative Awareness Messages (CAM), and second, GeoBroadcast for the multi-hop distribution of event-driven messages within a geo-area, known as Distributed Environment Notification Messages (DENM) [3]. The ETSI GeoNetworking protocol [2] specifies three main forwarding algorithms to distribute messages in a geo-area: Simple GeoBroadcast and contention-based forwarding as base schemes, and advanced forwarding, which combines both base schemes and comprises a set of protocol mechanisms to improve their performance. In order to allow a systematic evaluation of the standardized forwarding algorithms, we separate them by their technical features:

Simple GeoBroadcast (*S*) is a packet flooding scheme that restricts the packet re-broadcasting to the geo-area, i.e. every node inside the area retransmits the packet immediately after reception. The packet forwarding is controlled by a duplicate packet detection based on its source identifier and sequence number, maximum hop count and lifetime.

Contention-based forwarding (*CBF*) is an algorithm where the sender broadcasts a packet, the receiving neighbors buffer the packet and start a timer whose duration depends on their distance from the sender (the larger the distance, the shorter the timeout). When the timer expires, the packet is re-broadcasted. If the packet is received a second time before the timer expires, the timer is canceled and the packet is regarded as duplicate and discarded.

***CBF* with greedy forwarding** (*GFC*) combines *CBF* with a sender-based selection of the next hop for immediate forwarding (greedy forwarding) in *CBF* in order to avoid the forwarding delay caused by the *CBF* timer. The source defines a target area behind itself. Then it selects a 'virtual destination' position at the opposite border of the area, determines the best neighbor, i.e., that with the largest forward progress, and transmits the packet to the unicast MAC address of the best neighbor. Upon reception, the selected neighbor immediately forwards the packet. Other neighbors which overhear the packet transmission switch into *CBF* mode and start a timer. In case the greedy forwarding transmission gets lost, the nodes in *CBF* mode are able to trigger a retransmission when their *CBF* timer expires.

GFC with retransmission threshold (*GFC-RT*) extends *GFC* by a functionality for nodes in *CBF* mode: When such node receives a duplicate packet in the contention period, it keeps the packet and increments a Retransmit Counter (RC) instead of discarding the packet. If the RC exceeds a pre-defined threshold $\mathrm{RC_{th}}$, the packet is discarded and removed from the *CBF* buffer; otherwise, the packet is re-broadcasted when the *CBF* timer expires. *CBF* can be regarded as a special case of *GFC-RT* with a $\mathrm{RC_{th}} = 1$. This extension allows for redundant retransmissions to increase the reliability of the packet distribution at the expense of a higher data overhead.

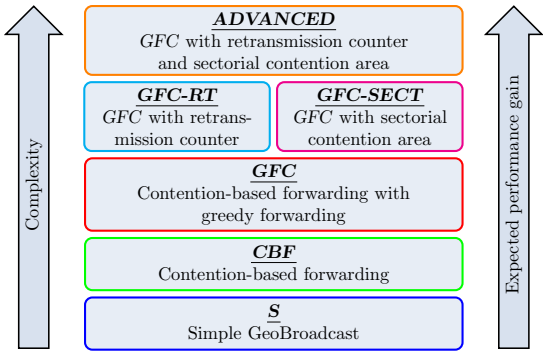

Fig. 1. Overview of the GeoBroadcast Forwarding Algorithms in the ETSI GeoNetworking protocol.

GFC with sectorial contention area (*GFC-SECT*) extends the *GFC* scheme and restricts the potential number of forwarders in *CBF* mode to those located outside a sector. The sectorial area is determined by the angle φ conformed by the forwarder (F) selected by the greedy forwarding algorithm, the sender (Se) and the overhearing relay (R): $\varphi = \angle \mathrm{FSeR}$. When the sender transmits the packet to the selected forwarder, the packet is overheard by all neighbors. The nodes outside the sector ($\varphi > 30°$) switch into *CBF* mode, whereas the nodes inside the sector do not contend ($\varphi \leq 30°$) and drop the packet. Thereby, the sectorial extension limits the number of potential forwarders to reduce the data traffic overhead.

Finally, the latter two extensions, *GFC-RT* and *GFC-SECT*, can be further combined, which corresponds to the 'advanced forwarding' algorithm in EN 302 636-4-1 [2] and will also be referred to as *ADVANCED* in the remaining sections of the paper. The six protocol variations as illustrated in Fig. 1 can be sorted by increasing complexity from *S* to *ADVANCED*[2]. An important question that this paper aims to answer is whether the increasing algorithm complexity results in a higher communications performance.

III. Performance Metrics

In order to evaluate the performance of the six variants of GeoBroadcast forwarding algorithms in ETSI GeoNetworking, we define the following metrics, which cover the protocol requirements in protocol overhead, reliability and latency.

Data Traffic Overhead (DTO): Naive multi-hop forwarding algorithms may cause many retransmissions [9], which result in higher contention on the wireless medium, increased number of collisions in the medium access and potentially a data load beyond the network capacity. To alleviate these issues, more advanced forwarding schemes attempt to reduce the number of redundant retransmissions to a minimum. The DTO allows measuring the data load in the network, as the total number of all nodes' data frames entering the MAC or the PHY layer, divided by the number of DENM messages

[2]For better readability, we defined a color code for the algorithms, which is reused in Sec. IV.

sent by the facility layer:

$$DTO_{\text{MAC}} = \frac{\text{No. of frames at MAC layer}}{\text{No. of facility DENM Tx}}, \tag{1a}$$

$$DTO_{\text{PHY}} = \frac{\text{No. of frames at PHY layer}}{\text{No. of facility DENM Tx}}. \tag{1b}$$

In order to understand the difference between DTO_{MAC} and DTO_{PHY}, it is important to note that a data packet that enters the MAC layer does not necessarily result in a PHY layer transmission. For example, during an ongoing MAC transmission or when the channel is sensed to be busy, new transmission requests are stored in the limited-size MAC queue. If the channel is congested, the MAC frames remain in the MAC queue until the CSMA/CA algorithm finds a free contention window or stops the transmission attempts for the current frame. In extreme cases, it may happen that the lifetime of the packet expires or the MAC queue overflows and further MAC frames are discarded.

Node Coverage Ratio (*NCR*): For road safety and traffic efficiency applications, such as hazardous location notification, it is important that every vehicle (we use the terms 'node' and 'vehicle' interchangeably) in the area reliably receives this message to inform the driver about potential risks and dangerous situations. In order to quantify the reliability of the message dissemination in a geo-area, the NCR of a message transmitted by the facility layer is defined as the ratio between the number of nodes in the geo-area receiving that message at least once and all nodes located in the geo-area:

$$NCR = \frac{\text{No. of Rx nodes in the geo-area}}{\text{No. of all nodes in the geo-area}}. \tag{2}$$

Nodes outside the area are not taken into account. Since the geo-area information is carried in the GeoNetworking header and does not change during the multi-hop transport of the packet.

End-to-End Delay ($E2ED$): In addition to the NCR, also the latency, i.e. the delay between the message transmission and its reception by the destination, is a key metric. A small $E2ED$ is critical for safety applications, such as a 'traffic jam ahead' warning, where a fast driving vehicle may quickly approach the rear end of a traffic jam with very low speed vehicles. A very short communication delay is required to raise the awareness of the driver or even to trigger the emergency brake automatically. The $E2ED$ is defined as the difference of the timestamp when the message is received by the destination and the timestamp when it is generated by the source including all forwarders:

$$E2ED = t_{\text{Rx}} - t_{\text{Tx}}. \tag{3}$$

We have identified DTO, NCR and $E2ED$ as the most relevant networking-related performance metrics for the GeoNetworking protocol. Next, we will use them to evaluate and compare the forwarding algorithms.

IV. Performance Evaluation

This section describes the considered scenario and the simulation environment, presents the results of the performance evaluation and proposes protocol improvements.

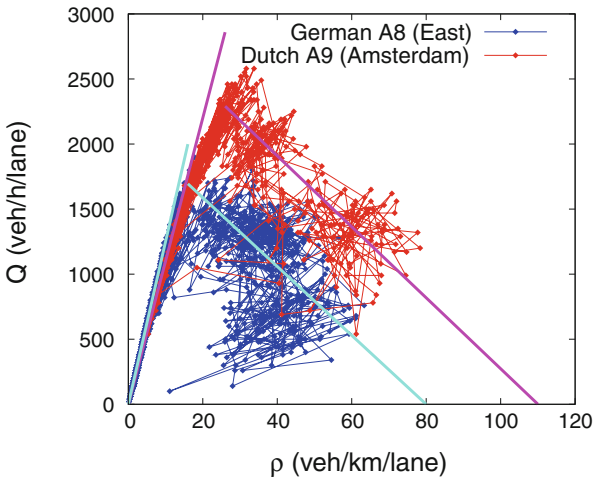

Fig. 2. Flow-density diagram for sections of two typical freeways. The solid lines are extrapolations of the data points [15]

A. Simulation Scenario and Environment

We assume a typical communication scenario with vehicles driving on a circular bidirectional freeway with three lanes per direction, a radius of 3 km and a speed limit of 120 km/h (33.3 m/s). However, this is only a limit, the vehicles have different speeds. The circular structure is needed to avoid edge effects which would occur if we just had a straight freeway. The vehicle density varies from 5 to 60 veh/km/lane (which corresponds to 500 and 6,500 vehicles in total, respectively), and is kept constant. These traffic densities cover free and congested road traffic conditions, which indicates a transition from free flow to congested road traffic, in line with those found on typical freeways shown in Fig. 2. The figure illustrates the flow-density diagram, where the vehicle flow Q steeply increases up to a density of $\rho = 20\,\text{veh}/\text{km}/\text{lane}$ and later declines until 60 veh/km/lane, which indicates the upper bound for our simulations. The distribution of vehicles is uniform along the whole freeway.

The simulated vehicles implement the ETSI protocol stack covering ITS-G5 PHY and MAC, GeoNetworking as specified in EN 302 636, and the facility layer protocols CAM and DENM. Every vehicle sends CAMs every 500 ms and 500 vehicles uniformly distributed over the circular freeway periodically send DENMs with an interval of 1 s. We use a simplified DENM protocol, where only new DENMs are sent and other DENM message types (update, negate, cancel) are not used. A vehicle is equipped with a single ITS-G5 transceiver and operates in a single channel. A simulation run is executed for 100 s; the DENM transmissions start at $t = 5\,\text{s}$ to ensure a steady state simulation where the GeoNetworking location tables are filled up with position information of neighboring vehicles and finish 5 s before the end to allow the last DENMs to be forwarded within the geo-area. The protocol stack of the vehicle models is simplified and does not include functionality for decentralized congestion control (DCC), i.e., gatekeeper functionality, QoS/EDCA, security and privacy, nor multi-channel/multi-transceiver operation. This simplification ensures that we are able to study the performance of the forwarding algorithms isolated from the effects of other protocols

TABLE I. SIMULATION PARAMETERS

Parameter	Values
Vehicle density	5-60 veh/km/lane (18.3 km, 6 lanes)
PHY / MAC protocol	ITS-G5 (802.11p)
Data rate/ Tx Power	6 $\frac{\text{Mbit}}{\text{s}}$ / 20 dBm
Channel bandwidth	10 MHz at 5.9 GHz
Simulation duration	100 s
CAM interval / size	0.5 s / 512 bytes
DENM interval / size	1 s / 512 bytes
Total DENMs sent	45,000 (500 vehicles)
Geo-area	$1,500 \times 100$ m behind the vehicle
Sectorial angle / RC_{th}	30° / 2

or functions.

The simulations are carried out with the network simulator ns-3 [16], version 3.19, using a log-distance propagation model with Nakagami fading [13]. The off-the-shelf version of ns-3 has been extended with the GeoNetworking and facility-layer protocols, i.e. CAM and DENM. The vehicle mobility is obtained from traces generated by the traffic simulator SUMO [17]. Due to the huge amount of nodes and since there is no difference in the results by changing the seed it is sufficient to run every simulation once. A summary of the main simulation parameters is given in Tab. I.

B. Data Traffic Overhead (DTO)

Fig. 3 presents the data traffic overhead at the MAC layer DTO_{MAC} and at the physical layer DTO_{PHY} (defined in Sec. III) as a function of the vehicle density for the three main GeoNetworking algorithms: Simple GeoBroadcast (*S*), contention-based forwarding (*CBF*) and contention-based forwarding with greedy forwarding (*GFC*). For better readability, the *GFC* extensions are elided; they have a similar behavior to *GFC*, shifted downwards for *GFC-SECT* and *ADVANCED*, and upwards for *GFC-RT*.

We can detect when the channel becomes congested by observing at which vehicle density DTO_{MAC} becomes greater than DTO_{PHY}, which means that some MAC frames could not be sent due to a busy channel. Indeed, we observe that *S* and *CBF* experience channel congestion from 10 and 40 veh/km/lane, respectively. In the case of *GFC*, it is more difficult to see when the channel becomes congested since DTO_{PHY} also includes the MAC retransmissions of the unicast greedy forwarding algorithm; in this case, $DTO_{\text{PHY}} > DTO_{\text{MAC}}$. We can find when *GFC* starts experiencing channel congestion by comparing its DTO_{MAC} to that of *CBF*. The particular feature of *GFC* is that a forwarder vehicle is selected to avoid additional delays, but the number of transmitted packets from the GeoNetworking layer is the same. Therefore, the channel congestion in *GFC* starts at 25 veh/km/lane (indicated as low density in Fig. 3), when its DTO_{MAC} becomes higher than that of *CBF*. We confirm this result with the increase in the end-to-end delay of *GFC* at the same vehicle density in Fig. 4b.

We observe that *S* has a much higher DTO_{MAC} than *CBF* and *GFC* and it increases with the vehicle density, due to the fact that all nodes re-broadcast the packets immediately after reception, which results in a huge amount of redundant transmissions. Furthermore, for low to medium vehicle densities, *CBF* has a slightly lower DTO_{MAC} than *GFC*, since the channel is already influenced by a higher data traffic overhead.

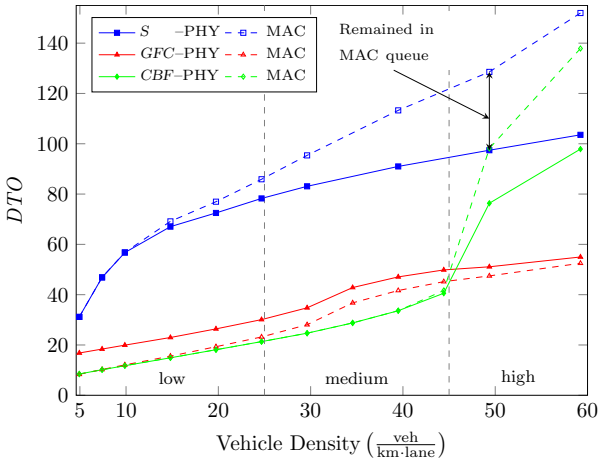

Fig. 3. Data traffic overhead DTO_{PHY} (solid lines) and DTO_{MAC} (dashed lines), both normalized by the number of sent DENMs.

For higher densities, however, the DTO of *CBF* increases abruptly and tends to that of *S*. The reason is that when the channel is congested, the packets must wait in the MAC queue before being transmitted; during this time the others nodes' CBF timers expire as well. In this case it is no longer possible to discard the packet after overhearing a duplicated packet, which leads to a high DTO.

Finally, we can see that, on the one hand, the DTO_{PHY} of *GFC* is greater than that of *CBF* for low vehicle densities, since the unicast greedy forwarding path causes MAC retransmissions that contribute to the DTO_{PHY} of *GFC*. On the other hand, these MAC retransmissions allow *GFC* to maintain a lower overhead than *CBF* for high vehicle densities, since when other nodes in the neighborhood overhear a packet transmission for the second time, they cancel their own retransmission timers and drop the packet. This mechanism prevents many redundant retransmissions and gives *GFC* the lowest DTO for high vehicle densities.

C. Node Coverage Ratio (NCR)

Fig. 4a presents the NCR of the GeoNetworking forwarding algorithms, which reflects the reliability of the packet transport in the network. From a high level perspective, we observe that the NCR is roughly inversely proportional to the data traffic overhead: a high DTO results in channel congestion and a large number of collisions, leading to a low NCR. Furthermore, when the channel is congested, some packets stay indefinitely in the MAC queue instead of being transmitted, as shown in Sec. IV-B, which further reduces the NCR. Vice versa, scenarios with a low overhead typically correspond to a high NCR.

For instance, we observe on the one hand that the NCR of *S* is less than 90 % even at low vehicle densities, and it further decreases than the vehicle density grows. On the other hand, for *CBF* and *GFC*, the NCR is almost 100 % up to a vehicle density of 30 veh/km/lane. For a vehicle density of 40 veh/km/lane, the NCR of *CBF* starts to decrease. However, for higher vehicle densities, the NCR of *CBF* decreases at a faster pace than in *GFC*. These trends are consistent with the data traffic overhead behavior observed in Fig. 3.

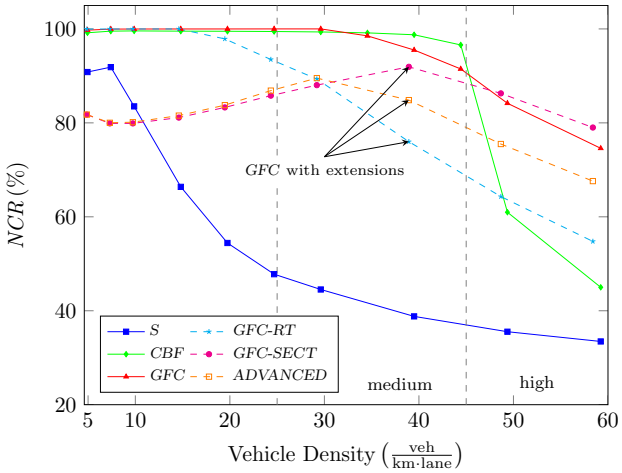

(a) Average node coverage ratio NCR

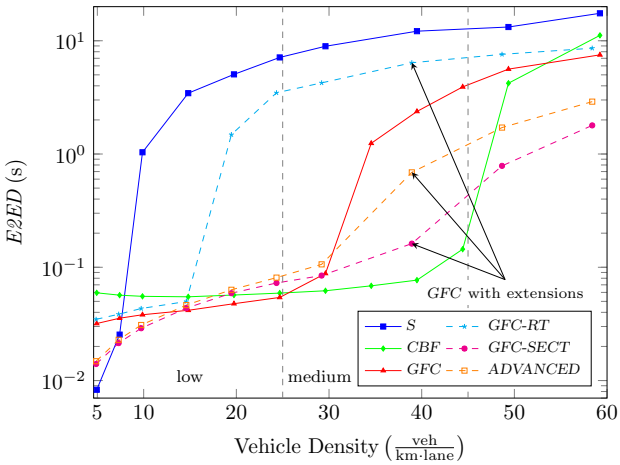

(b) Average end-to-end delay $E2ED$ of sent DENM

Fig. 4. Reliability and latency for Simple GeoBroadcast (*S*), Contention-based forwarding (*CBF*), and Contention-based forwarding with greedy forwarding (*GFC*) with its extensions Sectorial contention area (*GFC-SECT*), Retransmission threshold (*GFC-RT*), and altogether (*ADVANCED*).

We further observe in Fig. 4a that the *GFC* extensions, namely, *GFC-RT* and *GFC-SECT*, in general do not improve the NCR of *GFC*. In particular, *GFC-RT* shows a lower NCR than *GFC* for medium to high vehicle densities, whereas *GFC-SECT* outperforms *GFC* only for vehicle densities higher than 45 veh/km/lane. The reason for the low NCR of *GFC-RT* is that the additional packet retransmissions cause a higher data traffic overhead, which leads to a more congested channel. *GFC-SECT*, on the other hand, prevents most nodes overhearing a packet transmission from setting a retransmission timer; this mechanism reduces the data traffic overhead, but it may interrupt the packet dissemination in case that the greedy forwarding chain fails and all potential retransmitters are located inside the sector refrain from retransmission. Only at very high densities the sectorial contention area pays off and *GFC-SECT* achieves a higher NCR than *GFC*. The combination of both extensions, *ADVANCED*, achieves an intermediate NCR between the values obtained by the extensions implemented separately.

D. End-to-End Delay (E2ED)

Fig. 4b presents the $E2ED$ for all routing algorithms. *S* shows the shortest $E2ED$ at the lowest vehicle density (5 veh/km/lane), since there is no channel congestion (see Fig. 3) and nodes forward the packets immediately. However, as the vehicle density rises, the $E2ED$ increases very quickly to values longer than 1 s, which is unacceptable for safety-critical vehicular applications. The $E2ED$ of *CBF* is almost constant at 60 ms until 40 veh/km/lane, which coincides with the point at which the channel becomes congested (see Fig. 3). For higher vehicle densities, many packets cannot be transmitted due to a busy channel, the MAC queues fill up and the $E2ED$ grows rapidly. *GFC* shows a similar behavior to *CBF*, but the sharp increase in the $E2ED$ occurs at 30 veh/km/lane. Therefore, with *GFC* the channel becomes congested earlier than with *CBF* due to the additional load created by the unicast MAC retransmissions of the greedy forwarding algorithm. Finally, the *GFC* extensions show a shorter $E2ED$ for both low and high vehicle densities than standard *GFC*, whereas *GFC* has a lower $E2ED$ for medium vehicle densities.

E. Performance Summary

To summarize, we observed that for low vehicle densities (below 25 veh/km/lane) *GFC* shows the best overall behavior, with the highest reliability (NCR of 100 %), and a bounded latency with $E2ED$ below 50 ms. Beyond 25 veh/km/lane the performance of the algorithm deteriorates, i.e. its NCR decreases and the $E2ED$ rises. For medium vehicle densities (between 25 to 45 veh/km/lane), *CBF* outperforms the *GFC* algorithm with a higher reliability and a shorter latency, mainly because it creates less data overhead, which results in a less congested channel. For high vehicle densities (beyond 45 veh/km/lane), *GFC* has a clear advantage over *CBF* in terms of reliability at a comparable (high) latency.

We also observe that the extensions to *GFC* do not necessarily lead to a performance gain in most scenarios. *GFC-RT* allows for redundant retransmissions for higher reliability, but as *CBF* and *GFC* already achieve a NCR of 100 % at low to medium vehicle density in our simulation setting, *GFC-RT* cannot go beyond. On the contrary, *GFC-RT* adds additional load to the network such that the channel saturates at a lower vehicle density than *GFC*. However, we still expect that *GFC-RT* may be useful under worse channel conditions where *GFC* does not achieve a NCR of 100 %, such as in urban settings with buildings or obstacles blocking the propagation. The extension *GFC-SECT*, which restricts the number of forwarders to those outside the sectorial contention area, shows a good performance in high traffic densities, with the lowest end-to-end delay and a comparably high reliability. The combination of both *GFC* extensions (*ADVANCED*) achieves an intermediate performance between *GFC-RT* and *GFC-SECT*, as it could be expected.

F. Suggested Algorithm Improvements

Based on our performance study, we derive the following possible improvements for the standardized GeoBroadcast forwarding algorithms:

Forwarding range. The greedy forwarding algorithm used in *GFC* and its extensions selects the neighbor with the maximum progress towards the destination as the next hop.

This node is usually located at the edge, i.e. the "gray zone", of the communication range [14], which has a poor link quality. This implies that unicast transmissions to these nodes will likely fail and require MAC retransmissions. If the greedy forwarding algorithm would select only those neighbors with a minimum link quality, the number of MAC retransmissions required would be much lower, leading to a lower channel load and effectively contributing to a shorter $E2ED$ and higher NCR. This can be most simply achieved by defining a maximum forwarding range as an upper limit for the distance at which the selected next hop may be located.

CBF timer. In scenarios with a high channel load, the $E2ED$ is dominated by the MAC queuing delay instead of the *CBF* timer delays. In this case, the *CBF* overhearing mechanism does not work anymore and the performance of *CBF* approaches that of Simple GeoBroadcast. If the *CBF* timeout value took into account the lower layer processing, in particular the MAC queuing delay, we expect that the DTO of *CBF* at high vehicle density would decrease as more vehicles suppress their own retransmissions by correctly overhearing other vehicles' transmissions.

MAC layer retransmissions. When a source sends a unicast packet with *GFC* to the selected best hop and the transmission fails, the source attempts to retransmit the MAC frame up to a predefined number of times. These MAC layer retransmissions compete with the network-layer retransmissions triggered by *CBF*; when nodes in *CBF* mode overhear the MAC layer retransmissions from the source, they interpret them as re-broadcasts from other nodes and incorrectly refrain from retransmitting. As opposed to this, it might be more suitable to enable multi-path transmissions by letting other vehicles to retransmit instead to enforce a transmission over a lossy link. This could be achieved by reducing the maximum number or even disabling the MAC layer retransmissions.

While the first proposal to define a forwarding range can be directly applied to the 'advanced forwarding' algorithm in EN 302 636-4-1, the *CBF* timer modification requires knowledge of the expected layer-2 processing delay. A possible implementation could use cross-layer information of the estimated delay from the layer-2 queuing functionality. The last proposal implies a slight modification of the IEEE 802.11p MAC protocol to allow a more flexible control of the maximum number of retransmission attempts. Finally, we note that all three proposals suggest improvements for the *CBF* principle and therefore also apply to *GFC* and its variations.

V. Conclusion

This paper evaluated the GeoNetworking protocol as standardized by ETSI in its Release 1 for cooperative ITS and compared the GeoBroadcast forwarding algorithms, namely, Simple GeoBroadcast (*S*), contention-based forwarding (*CBF*), contention-based forwarding with greedy forwarding (*GFC*) with its extensions sectorial contention area (*GFC-SECT*), retransmission threshold (*GFC-RT*), and finally the combination of both *GFC* extensions (*ADVANCED*). Performance metrics and a highway simulation scenario were defined to quantify the reliability, latency and traffic overhead of the forwarding schemes. While *S* disqualifies due to its very high overhead and resulting high latency and low reliability, *CBF* outperforms *GFC* over a wide range of vehicle densities. The *GFC* extensions provide little benefit in the considered scenario, but they may be beneficial in urban scenarios with worse wireless links. Based on the performed evaluation, improvements to the GeoNetworking standard algorithms are proposed.

Acknowledgment

This work has been performed in the framework of the FP7 project ICT-619555 RESCUE (Links-on-the-fly Technology for Robust, Efficient and Smart Communication in Unpredictable Environments), which is partly funded by the European Union. We also thank Dr. Martin Treiber for valuable discussions.

References

[1] A. Festag, "Cooperative Intelligent Transport Systems (C-ITS) Standards in Europe," *Communication Magazin*, vol. 12, no. 52, Dec. 2014.

[2] ETSI, *Intelligent Transport Systems (ITS); Vehicular Communications; GeoNetworking; Part 4: Geographical Addressing and Forwarding for Point-to-point and Point-to-multipoint Communications; Sub-part 1: Media-Independent Functionality (ETSI EN 302 636-4-1 V1.2.1)*, Jul. 2014.

[3] ——, *Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 3: Specifications of Decentralized Environmental Notification Basic Service (ETSI EN 302 637-3 V1.2.2)*, Nov. 2014.

[4] B. Karp and H. T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks," in *ACM MobiCom*, New York, NY, USA, Oct. 2000, pp. 243–254.

[5] I. Stojmenovic, "Position-based Routing in Ad Hoc Networks," *IEEE Communications Magazine*, vol. 40, no. 7, pp. 128–134, Jul. 2002.

[6] H. Füßler, J. Widmer, M. Käsemann, M. Mauve, and H. Hartenstein, "Contention-based Forwarding for Mobile Ad Hoc Networks," *Ad Hoc Networks*, vol. 1, no. 4, pp. 351–369, Nov. 2003.

[7] T. Osafune, L. Lin, and M. Lenardi, "Multi-Hop Vehicular Broadcast (MHVB)," in *2006 6th International Conference on ITS Telecommunications Proceedings*, Jun. 2006, pp. 757–760.

[8] M. Torrent-Moreno, J. Mittag, P. Santi, and H. Hartenstein, "Vehicle-to-Vehicle Communication: Fair Transmit Power Control for Safety-Critical Information," *IEEE Trans. Veh. Technol.*, vol. 58, no. 7, pp. 3684–3703, Sept 2009.

[9] O. Tonguz, N. Wisitpongphan, J. Parikh, F. Bai, P. Mudalige, and V. Sadekar, "On the Broadcast Storm Problem in Ad hoc Wireless Networks," in *BROADNETS*, New York, NY, USA, Oct. 2006, pp. 1–11.

[10] A. Festag, P. Papadimitratos, and T. Tielert, "Design and Performance of Secure Geocast for Vehicular Communication," *IEEE Trans. Veh. Technol.*, vol. 59, no. 5, pp. 2456–2471, 2010.

[11] V. Sandonis, I. Soto, M. Calderon, and M. Urueña, "Vehicle to Internet Communications Using the ETSI ITS GeoNetworking Protocol," *Transactions on Emerging Telecommunications Technologies*, pp. 2161–3915, Oct. 2014. [Online]. Available: DOI:10.1002/ett.2895

[12] P. d'Orey and M. Boban, "Empirical Evaluation of Cooperative Awareness," in *IEEE VTC Spring'14*, Seoul, Korea, May 2014.

[13] T. Abbas, K. Sjöberg, J. Karedal, and F. Tufvesson, "A Measurement Based Shadow Fading Model for Vehicle-to-Vehicle Network Simulations," *arXiv:1203.3370 [cs]*, Mar. 2012. [Online]. Available: http://arxiv.org/abs/1203.3370

[14] H. Lundgren, E. Nordström, and E. Tschudin, "Coping with Communication Gray Zones in IEEE 802.11b Based Ad Hoc Networks," in *WOWMOM*, New York, NY, USA, Oct. 2002, pp. 49–55.

[15] M. Treiber, A. Kesting, and C. Thiemann, *Traffic Flow Dynamics : Data, Models and Simulation*. Springer, 2012.

[16] "NS-3 Network Simulator 3," Dec. 2014. [Online]. Available: http://www.nsnam.org

[17] D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker, "Recent Development and Applications of SUMO - Simulation of Urban MObility," *International Journal On Advances in Systems and Measurements*, vol. 5, no. 3&4, pp. 128–138, December 2012.