Journal of Mobile Multimedia, Vol. 5, No. 2 (2009) 158–180

DISTRIBUTED INFORMATION MANAGEMENT AND PUBLISH/SUBSCRIBE IN VANETS: REQUIREMENTS, STATE OF THE ART AND A NOVEL P2P-BASED APPROACH

VIVIAN PRINZ
Department of Informatics, Technische Universität München, Boltzmannstr. 3
85748 Garching, Germany
prinzv@in.tum.de

MICHELE BROCCO
Department of Informatics, Technische Universität München, Boltzmannstr. 3
85748 Garching, Germany
brocco@in.tum.de

WOLFGANG WOERNDL
Department of Informatics, Technische Universität München, Boltzmannstr. 3
85748 Garching, Germany
woerndl@in.tum.de

Received July 21, 2008
Revised March 25, 2009

Wireless communication is particularly powerful in vehicular ad hoc networks (VANETs) as it implies important possibilities to enhance traffic safety. For this purpose, a car-to-car (C2C) communication system should provide distributed information management. For example, it should guarantee the availability of a black ice warning for the duration of its validity. In addition, the Publish/Subscribe (Pub/Sub) paradigm allows for information filtering and tracking of changes in the information environment. Thus, vehicles could monitor traffic changes, for example.

This article introduces the application of structured peer-to-peer (P2P) algorithms and P2P/Distributed Hash Table (DHT)-based Pub/Sub for the named functionalities. To allow for their application given high mobilities and to provide location-awareness, the vehicular network is segmented into separate, interacting P2P networks. Cars manage information of their own segment and exchange them between adjacent segments if necessary. The Pub/Sub functionality is built on top of the P2P segments' DHTs. This way, Pub/Sub can be applied as there is no need to maintain a vehicular broker tree. The proposed solution enables C2C applications to publish information referring to certain areas or validity durations. In addition, they can request, modify and delete this information and solely interested vehicles can be notified about these events.

Keywords: vehicular ad hoc networks, peer-to-peer algorithms, distributed hash tables, publish/subscribe, wireless communication systems, car-to-car communication, car-to-car applications

Communicated by: D. Taniar

1 Introduction

Despite the multitude of vehicle security components, the number of damages to property and persons caused by traffic accidents is still worrying. For instance, in 2006, the European Commission reported more than 40.000 fatalities only within the European Union [1]. *Vehicular ad hoc networks (VANETs)* are distributed, self-organizing and highly mobile ad hoc networks. Wireless communication between these vehicles is referred to as *car-to-car (C2C)* communication. It enables vehicles to recognize situations that are not detectable by local sensors, the driver himself or single cars. Therefore, C2C communication is considered as one future key enabler for increasing traffic safety.

Consequently, this research area has gained much attention from both academia and industry in recent years. There are outstanding solutions for many features concerning C2C communication systems. Yet, distributed information management is still challenging. For example, a C2C communication system should facilitate the location-specific publication of black ice warnings or parking lot information. Thereafter, these information items have to be available for the duration of their validity. Currently, each piece of information is therefore repeatedly broadcasted. Particularly with regard to information not of general interest (e. g. available parking lots) and given resource constraints, improvements are of relevance. In addition, subsequent data handling of already distributed information is a matter in this field. A further topic is *Publish/Subscribe (Pub/Sub)* communication. With this technique, a C2C communication system could actively deliver warnings or traffic changes to concerned vehicles only, for example. However, Pub/Sub realization in large scale and highly mobile networks is demanding as well. Hence, Pub/Sub has not been applied in VANETs yet.

In this article, we focus on the application of structured *peer-to-peer (P2P)* algorithms for the named functionalities. Amongst other things, P2P algorithms are concerned with reliable distributed information management. Furthermore, we show that VANETs and P2P networks bear strong resemblances. To give an example, both network types lack central control and are self-managed. Finally, *Distributed Hash Tables (DHTs)* of structured P2P algorithms can be used for a broker tree-independent implementation of the Pub/Sub communication primitive. This is desirable as vehicular broker trees are hard to maintain due to frequent topology changes. However, there are factors like the high mobility of vehicles or the required street topology awareness, which have not been considered when designing existing P2P algorithms. Furthermore, also P2P/DHT-based Pub/Sub relies on an event distribution tree for scalability reasons. Thus, we introduce mechanisms to adapt these differences in order to apply structured P2P algorithms and P2P/DHT-based Pub/Sub for C2C communication systems.

The remainder of this paper is structured as follows: Section 2 describes C2C applications and accompanied requirements on vehicular communication systems. Section 3 provides an overview on related research areas and technologies. Afterwards, our solution is introduced and discussed in section 4. Section 5 presents our evaluation environment. It comprises our simulation environment as well as an evaluation of built-on state-of-the-art information management as groundwork for our current work. Finally, the article gives a summary and an outlook on our further work in section 6.

2 C2C Applications and Accompanied Requirements

This section describes different classes of C2C applications. Afterwards, it derives demands on the underlying C2C communication system with regard to these applications.

2.1 C2C Applications

C2C communication provides the basis for future applications in the field of *active safety*. For example, vehicles can inform each other about preceding accidents or black ice. Another example is cooperative traffic jam detection performed by multiple cars. This bears the possibility to warn vehicles about traffic jams or even to avoid them by propagating appropriate speed limitations.

Currently, only road signs, present police officers or radio news can inform vehicles. A radio message implicates the disadvantage that it is not limited to the affected area. For some drivers the message is uninteresting. Other drivers forget the message until they reach that dedicated region. Furthermore, plenty of time passes until a report reaches a radio station, road signs are placed or the police is on the spot. Particularly with regard to local relevance and fast propagation of information, C2C communication offers important potential for reducing the number of accidents.

Besides active safety applications, the so-called *deployment applications* provide additional benefits for drivers. Deployment applications are valuable add-on applications like assistants helping drivers to find dedicated points of interest in their vicinity. It is assumed, that these applications considerably contribute to the dissemination of C2C communication systems. This is important as active safety applications usually rely on a certain degree of dissemination.

To give an example, deployment applications can provide drivers with information about nearby parking lots by accessing the C2C network: Car parks are able to provide information about available parking lots and their costs using wireless communication endpoints. Vehicles leaving a public parking lot can announce this as well. By aggregating this information, C2C applications can visualize nearby parking lots and, as may be required, their costs, through in-car navigation systems. Like active safety applications, deployment applications benefit from regional relevance and up-to-dateness of information accessible via the C2C network.

2.2 Requirements on C2C Communication Systems

Obviously, a multitude of possible applications exists within the named two classes. C2C communication systems should support this wide range of applications. Therefore, it is necessary to provide a communication platform with a well-defined interface arbitrary C2C applications can be built on.

In this paper, we are investigating functionalities of decentralized C2C communication systems. This is because it is not possible to seamlessly provide and maintain infrastructure nodes in the whole road network. Therefore, decentralized organization offers some inherent advantages over centralized systems. First, there is no need for every vehicle to keep an internet connection. Hence, there are no barriers concerning costs or accessibility. Second, there are no bills to pay for central authorities hosting and maintaining corresponding servers. Third, central servers can become a bottleneck or even fail – decentralized systems avoid scalability problems, which is particularly important in large-scale vehicular networks. The following requirements on decentralized C2C communication systems can be identified:

- **Resource Efficiency:** Vehicles communicate over the air interface. Hence, many messages lead to higher collision probabilities and longer access times. Both consequences have to be avoided particularly with regard to their impact on active safety applications.
- **Persistent Availability:** It is not sufficient to send a black ice warning to all successive vehicles once. Vehicles reaching the hazard zone later have to be warned as well. Consequently, information has to be available over a variable period.
- **Retrievability:** Due to given resource constraints, information should not be provided by continuous retransmissions. Data like parking lot information is interesting for only a few, dedicated vehicles. Thus, it should be possible to store information in a distributed manner and to allow arbitrary vehicles to retrieve this information. In the following, we will refer to this kind of distributed storage as *publishing*.
- **Modifiability:** Car parks have to be able to modify published information about their capacity when parking lots are taken or released, for example.
- **Deletion/Expiration:** Drivers must not be warned about black ice after road conditions have normalized again.
- **Location Dependency:** It should be possible to publish information referring to certain geographical regions. A warning about bad road conditions is irrelevant in far-off regions. Hence, it should not be managed there for efficiency reasons and for fast propagation.
- **Reliability:** As we presume no centralized infrastructure, cars manage information themselves. Thereby, information must not get lost if a vehicle gets out of reach unexpectedly. This happens if a vehicle passes a tunnel or is being parked, for example.
- **Pub/Sub:** Information is rather directed at cars with certain properties than at vehicles with a known address. For example, some warnings are only relevant for cars driving on a particular lane. Other warnings may be of concern just for certain vehicle types. Heavy trucks in case of bridge damage, for instance. Thereby, the vehicle which detects the hazard and produces the warning does not know the cars interested in that message. Apart from that, it would be of advantage to allow vehicles to be informed if certain published information – the capacity of a car park or traffic density in a dedicated area – change. The Pub/Sub asynchronous messaging paradigm allows for loosely coupled asynchronous messaging covering information filtering and the tracking of changes (see Section 3.3).

To summarize, a C2C communication system should realize reliable distributed information storage. Thereby, it should allow for publishing information with respect to certain geographical regions or specified validity durations if necessary. Furthermore, it should be possible to modify, delete and to search for this information. In the following, we will refer to these mechanisms as distributed information management. Apart from that, loosely coupled asynchronous messaging over Pub/Sub would be of advantage. From the application point of view, these requirements imply the provision of publish, delete, search, modify and subscribe primitives besides send and receive.

3 Related Work and Technologies

This section gives an overview on distributed information management and Pub/Sub in C2C communication systems. Thereby, we describe related approaches focusing on C2C communication systems. Furthermore, we introduce P2P algorithms as they address distributed information management. In addition, we discuss current related work concerning the application of P2P algorithms in C2C networks. Finally, we give an overview on Pub/Sub as well as on P2P/DHT-based Pub/Sub and describe approaches aiming at applying this communication primitive in C2C networks.

3.1 C2C Communication Systems

Regarding C2C communication systems, the state of the art in direct communications as well as in information management are of relevance.

3.1.1 Direct Communication

A lot of the current work in C2C research is on media access, multi-hop routing and group communication primitives. Media access mechanisms control air interface access among several vehicles (e. g. [2]). Multi-hop routing protocols use intermediate vehicles to transmit data to otherwise unreachable receiver vehicles (e. g. [3], [4], [5]). Multicast approaches apply group communication primitives migrated from wired networks to mobile networks (e. g. [6]). Finally, Geocast protocols enable the transmission of a message to multiple vehicles within a dedicated geographical area (e. g. [7]).

Moreover, the so-called *Dedicated Short Range Communication Standard (DSRC)* [8] is currently being developed in America, Europe and Japan. It is designed to support wireless communication for C2C and *car-to-infrastructure (C2I)* applications. As the groundwork for DSRC, the standard IEEE 802.11p is used, which is derived from the IEEE extension 802.11a. The standard envisions vehicle speeds of up to 200 km/h and direct communication within a distance of up to one kilometer or 3000 $feet$ respectively (see Fig. 1).

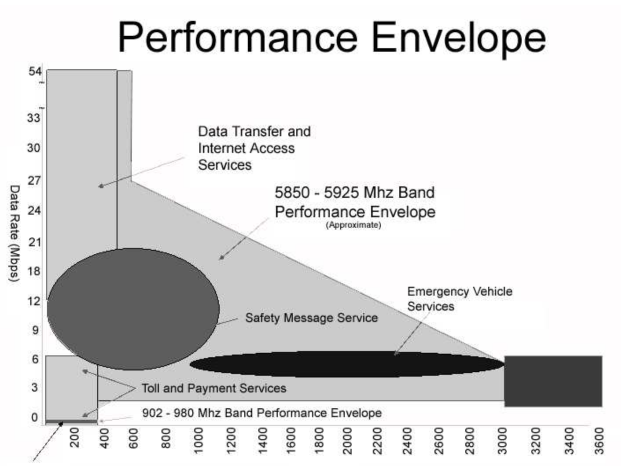

Fig. 1. DSRC performance [9]

3.1.2 Information Management

The named mechanisms are used to transmit messages over the air interface to single vehicles or vehicle groups. They are not concerned with distributed information management. In this context, several approaches discuss intelligent retransmissions of broadcast messages to realize location-aware persistent availability. That is, vehicles within an area retransmit information according to dedicated rules. This way, vehicles reaching the area later receive the messages as well.

For example, the *self-organizing traffic information system (SOTIS)* [10] implements a decentralized system for traffic density analysis. Each car propagates its position, its speed and its driving direction repeatedly. With this information, every vehicle is able to determine the current traffic situation by calculating the average speed of the traffic participants. In [11], the authors introduce message priorities influencing the repetition interval of SOTIS-messages. This way, traffic jam warnings can be propagated faster than regular SOTIS messages, for instance. Xu et al. discuss in [12] a location based C2C broadcast protocol. Thereby, vehicles generate warning messages in fixed time intervals as well. Their protocol design focuses on handling repeated warning messages triggered by several cars. The considerations made are based on the fact that usually several vehicles are involved in a traffic accident. Therefore, the authors propose mechanisms to differentiate warnings and to eliminate redundant ones. Thereby, the repetition interval is optimized with the objective of minimizing the channel delay for each vehicle. In contrast, [13] presents an analysis of random repetition intervals for safety critical messages in a highly mobile environment. A part from that, a lot of work has been done on different research questions arising in conjunction with *mobile ad hoc networks (MANETs)* (e. g. [14]). In this context, Williams and Camp [15] give a general overview on algorithms that can be used to minimize the number of redundant rebroadcasts in MANETs.

Yet, none of these approaches is able to avoid the high number of redundant messages. This is because every piece of information is repeatedly sent by multiple cars until its validity expires. An alternative approach is the so-called *Stored Geocast*. Geocast messages are sent to one node. Afterwards, that node delivers them within the desired region. For example, Imielinski and Navas [16] divide the ad-hoc network into equal cells. One *GeoNode* is assigned to each cell. The GeoNode stores the messages persistently and propagates them. In [17], an infrastructure independent solution is proposed. An initial Geocast is followed by a distributed election algorithm. This algorithm is applied to determine the node that will manage and deliver the Geocast message. As soon as the node leaves the region, it initiates a new election to guarantee further message delivery.

Also in Geocast approaches, every piece of information is continuously retransmitted. Alternatively, every vehicle can broadcast its location periodically. This way, Geocast servers can recognize a vehicle entering their area of interest. Hence, they can deliver their Geocast messages on demand. However, continuous rebroadcasts and inherited redundancies are required in either case.

3.2 Peer-to-Peer Algorithms

Among other things, P2P algorithms address reliable, distributed information management. Therefore, this section focuses on P2P algorithms and their application in C2C networks.

3.2.1 P2P/DHT-based Information Management

P2P algorithms are applied to realize P2P networks. Pure P2P networks are fully decentralized, distributed systems. That is, they do not provide central controlling components, but are self-organized [18]. So far, P2P algorithms were mainly used to implement internet applications. In this case, peers are the users' desktop PCs connected over the internet. P2P algorithms provide features like redundant, persistent and distributed storage of data including data expiration. Further example features are efficient data lookup and data handling routines, support for trust and anonymity as well as reliability. Therefore, P2P algorithms can be used for a wide range of different application domains. Popular example applications are distributed file systems, file sharing services and communication clients such as Skype.

Different applications are based on different P2P algorithms to manage the underlying P2P network. These networks are usually classified into *structured* and *unstructured* P2P networks. An unstructured P2P network (e. g. [19]) consists of peers joining and leaving the network according to only a few rules. There is no prior knowledge about the network topology. Likewise, distributed information management is topology independent: Objects are randomly located at peers and searched by flooding the network with requests. Various optimizations exist in this context.

In contrast, structured P2P networks base on a strictly controlled topology. Both data as well as information about other peers are stored in a distributed manner by the participating peers. Thereby, structured P2P networks utilize Distributed Hash Tables as carriers.

A DHT is a data structure used to determine where to store a piece of information. Likewise, it allows finding it again: Every piece of information has a unique key. Information keys and peer addresses are mapped onto identifiers using a hash function. Information is then placed on the peer whose identifier is most similar to the information's identifier. Every peer maintains a small routing table containing the identifiers of neighboring peers and their addresses. They are used to route requests to the peer with the required identifier. This way, information can be published, modified, deleted or simply retrieved. Expiration is realized by passing a validity period when publishing. The peer storing the information deletes it as soon as the validity expires. Corresponding structured P2P algorithms basically differ in their data object schemata, their address spaces, their routing strategies and their replication strategies to achieve fault tolerance. Usually, neighboring peers in the address space take over a peer's identifier interval and corresponding information if it fails.

A popular example is the structured P2P algorithm *Chord* [20] and its ring topology. Chord virtually places peers onto a ring ordered by ascending identifiers. As described, the identifiers are calculated using a hash function. Figure 2 illustrates a Chord ring. Peers (nodes, N) manage information items (key, K) with certain identifiers. For example, the peer with identifier 56 is currently responsible for the information with identifier 54. There is no other peer between peer 56 and peer 51. That is, there is no peer whose identifier is greater than 51 and also smaller than 56. Because of that, peer 56 is responsible for all information items whose identifiers lie in between 56 and 51. Generally speaking, peers are always hosting information, whose identifiers lie in the interval bounded by the peer's identifier and his predecessor's identifier. Hence, the aforementioned similarity of peer and information identifiers equates to unidirectional proximity on a ring as far as the Chord algorithm is concerned. Likewise, the named neighborship relation refers to peers that are next to each other on the ring. For more detailed information on P2P algorithms see, for example, [21].

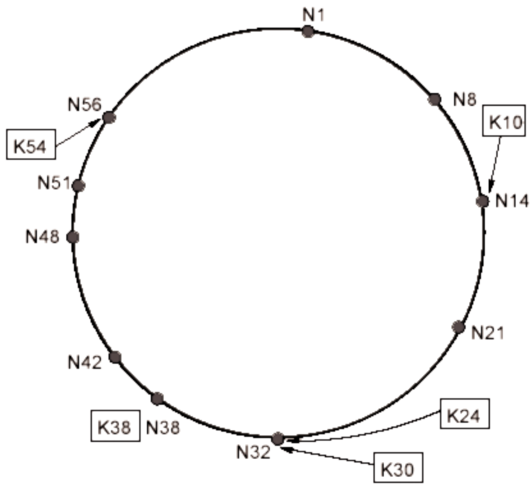

Fig. 2. The Chord Ring [20]

3.2.2 P2P Algorithms for C2C Communication Systems

C2C and P2P networks bear strong resemblance. Both are distributed, self-organizing networks. Furthermore, nodes act autonomously and anonymously. Internet users as well as vehicles joining or leaving the network are not known in advance. In addition, P2P and C2C networks tend to be very large – internet and road network. Finally, both suffer from frequent topology changes. PCs are switched on and off arbitrarily. Vehicles are started and shut off or come to pass street sections without radio connectivity. Hence, P2P algorithms take properties of C2C networks into account.

There are solutions concerned with integrating mobile nodes into fixed network infrastructure-based P2P service provisioning (e. g. [22]). Regarding our field of infrastructure independent information management, several solutions propose to combine MANETs and P2P (e. g. [23], [24]). However, they do not regard C2C networks. Consequently, they do not have to consider criteria like high vehicular speeds, the size of vehicular networks or references to and influences of the street topology.

In the field of C2C research, the SmartWeb project [25] and Huang et al. [26] address the issue of both P2P and C2C communication. These approaches mean a direct endpoint-to-endpoint connection between two vehicles when talking about P2P communication. As a result of the SmartWeb project, a direct WLAN-based connection between two real vehicles can be established. This connection is used to propagate sudden braking between these two vehicles. Huang et al. use a direct connection between two vehicles driving towards each other to exchange data about their driving speed and direction. Afterwards, the probability for a collision is calculated based on the driver's reaction time. This way, collision warnings can be displayed to the driver at the right time.

In [27], Lochert et al. compare the central server approach and the possibility to regard the vehicular network as one big structured P2P network. For the latter, application specific information, requests and topological information have to be exchanged over vehicular multi-hop connections. These connections may exceed many kilometers. Thereby, information exchange has to take place in a fast and reliable way, which poses a problem in VANETs (see Section 4.1). In addition, the authors point out the huge network load. All named information would have to be routed frequently and most often over long distances. Thus, they conclude that, despite the substantial benefits of P2P, infrastructure-based solutions will be the ones preferred – at least in the near future. In [28] Lochert et al. suggest a solution for an infrastructure-based P2P network. The network is made up of vehicles each of which maintains an internet connection to be able to access fixed network infrastructure. This way, they take advantage of P2P functionalities while avoiding the stated problems.

As motivated in section 2, we are focusing on infrastructure-independent solutions. One can summarize, that there are no approaches that allow for fully decentralized P2P/DHT-based information management in VANETs.

3.3 Publish/Subscribe

Besides information management, we are concerned with the provision of messaging according to the Pub/Sub paradigm. This section first introduces this paradigm. Afterwards, it presents solutions that realize Pub/Sub in an entire different way by utilizing the DHT of a structured P2P network. Finally, we describe the current state of Pub/Sub deployment in C2C networks.

3.3.1 The Publish/Subscribe Messaging Paradigm

Pub/Sub is an asynchronous messaging paradigm where message producers and consumers are distinguished. Producers do not send messages to specific consumers. Rather, published information is characterized by topics, without knowledge of what consumers there may be. Consumers express interest in one or more topics and only receive messages of interest, without knowledge of what publishers there are. Hence, producers and consumers are loosely coupled.

Consumers declare their interest on a certain topic by issuing a *subscription* to that topic. Thereby, they can give particulars about themselves or the intended message contents. As soon as an information item is published, it is forwarded to all corresponding subscribers. To this end, brokers manage subscriptions and forward information items. Multiple brokers are applied to achieve scalability. A subscription is always located at one of these brokers. When publishing an information item, the publisher connects to one broker – probably the nearest in the network. Afterwards, this broker forwards the information along a tree of other brokers (*Pub/Sub tree (PST)*). Thereby, the brokers filter information according to the topics specified in their subscriptions (*topic-based filtering*). In addition, they may evaluate if the attributes or content of those information match the constraints defined by their subscribers (*content-based filtering*). Finally, they may not only inform subscribers if information of interest is published, but also if it is altered or deleted. This presumes information to be stored at some node of the PST. The modification or deletion request then causes the node to edit or delete the information. Afterwards, a corresponding update is sent along the PST and subscribers are informed. In Pub/Sub, communication is therefore anonymous and inherently asynchronous. Hence, a system is able to quickly adapt in a dynamic environment. For more detailed information on Pub/Sub systems see, for example, [29].

3.3.2 P2P/DHT-based Publish/Subscribe

Several approaches apply Pub/Sub on top of a DHT of a wired, structured P2P network (e. g. [30], [31]). The basic idea is not to distribute subscriptions over a PST which information items have to traverse in order to be checked against every existing subscription. Rather, it is utilized that information items of the same topic can always be stored on the same peer: Instead of unique keys, topics are used to calculate information identifiers. Hence, each topic determines one additional peer to store an information item on. Subscriptions refer to a particular topic and are stored in a distributed manner themselves. By hashing a subscription's topic, the subscription is stored on the peer that also stores all information of the referenced topic. This peer knows when topic-related information is published, modified or deleted because only the peer itself carries out the operations on request. After execution of such an operation, the peer checks if there are any related subscriptions. If so, it compares their properties with those of the information and informs suitable subscribers.

To ensure scalability, also P2P/DHT-based Pub/Sub solutions rely on distribution trees. This is because a peer solely processing information filtering and message delivery for certain topics may soon become a bottleneck in large-scale networks. Thus, *rendezvous-based routing algorithms* are applied on top of the structured overlays [32]. The basic idea is to build a multicast tree for each topic. These trees are single-rooted diffusion trees. They are rooted in their topic's responsible peer and are spanning all corresponding subscribers. When a peer publishes a new subscription, the subscription is routed to the responsible peer. This peer then updates the tree structure in order to include the new subscriber. When an information item of a certain topic is published, it is routed to the responsible peer as well. But then, matching is reduced to identifying the correct multicast tree and notification routing is performed by diffusing the notification through the tree. Accordingly, subscriptions and information related to a certain topic still remain on one peer, but information delivery is enhanced.

3.3.3 Publish/Subscribe for C2C Communication Systems

Regarding existing Pub/Sub systems, there are many implementations for wired networks (e. g. [33], [34], [35]). In addition, there are approaches that apply Pub/Sub communication in MANETs. Examples are [36] and the middleware described in [37]. Generally speaking, each node thereby acts as a broker [32]: Information is flooded and nodes locally maintain their subscriptions. Subscriptions are used to filter incoming messages. This mechanism is basically chosen because there is no fixed infrastructure to place brokers at and maintaining a PST on mobile nodes is difficult. Regarding C2C communication, every information item would have to be reflooded periodically if vehicles locally kept their subscriptions. Otherwise, vehicular subscribers newly arriving, newly being started or newly issuing a subscription would not be informed. This is because there are no brokers for checking a new subscription against existing information.

[38] describes a protocol to build and reconfigure PSTs in MANETs. The solution is implemented in the REDS framework [39]. Yet, the protocol is not applied for high mobilities. One can summarize that running algorithms for keeping a tree-topology over a set of mobile nodes is expensive. Thus, corresponding approaches are not suited to highly mobile environments [32]. Consequently, realizing classical Pub/Sub in VANETs is a subject for further exploration.

4 Approach: Interacting P2P Segments for C2C Communication Systems

This section specifies challenges that arise in conjunction with the application of structured P2P algorithms and P2P/DHT-based Pub/Sub in C2C communication systems. Furthermore, our solution for these challenges is introduced and discussed.

4.1 Challenges

As explained in section 3, structured P2P algorithms can be used to realize distributed information management. In addition, they consider network characteristics P2P and C2C networks have in common. Finally, the DHT of a structured P2P algorithm can be utilized to realize a PST-independent implementation of the Pub/Sub messaging paradigm. However, with C2C there are network properties that have not been considered when designing existing structured P2P algorithms and P2P/DHT-based Pub/Sub.

One important difference concerns communication connections. P2P algorithms and Pub/Sub come from the internet world. On the internet, the quality of connections varies only very little. VANETs mainly pose the problem that vehicles are not able to communicate just as fast and reliable [40].

On the one hand, reliability is restricted because obstacles like trees may jam wireless connections. On the other hand, communication is slower because multi-hop routing protocols for VANETs require a lot of message exchanges. This is due to the high vehicle speeds, which necessitate many routing table updates. Multiple messages again result in higher collision probabilities and longer access times. Alternatively, connections are established on demand, which results in delays. Finally, established connections over multiple vehicles are not very stable again due to the high vehicular velocities.

The problem is that structured P2P algorithms and Pub/Sub require fast and reliable exchange of data and information. First, a slower transmission of topology information may result in inconsistencies in the structured P2P network. Also, rapid (re-)allocation of application data in case of failure recovery might be a problem resulting in lack of fault tolerance. Hence, P2P information management might fail if messages are transmitted too slowly or not reliably.

Second, regarding Pub/Sub, the issue arises that building and maintaining a PST is difficult in VANETs. This is because there are frequent topology changes due to mobility in addition to the ones due to churn and subscription changes. Assuming a structured P2P overlay, it would be possible to avoid the maintenance of a PST by applying P2P/DHT-based Pub/Sub. However, P2P/DHT-based Pub/Sub solutions again use distribution trees to ensure scalability (see Section 3.3.2). The problem persists. In addition, slower and less reliable exchange of application data and subscriptions poses a challenge in this context as well.

Apart from that, information like a black ice warning is only relevant in a particular region. Consequently, information items like these should not be managed in far-off areas for efficiency reasons and for fast propagation (see Section 2.2). P2P algorithms are not concerned with information management referring to certain regions of the road network.

4.2 P2P/C2C Segments

A solution for adapting these different conditions is not to regard the vehicular network as one big structured P2P network. Instead, it is divided into segments each forming a structured

P2P network. The segment sizes are chosen according to the available communication range. For example, in urban regions street sections of up to 1000 *meters* would be chosen when assuming the maximum DSRC communication range. The address space of a P2P network is limited to these segments. Thereby, neighboring segments overlap logically at their borders. Figure 3 illustrates these overlapping segments.

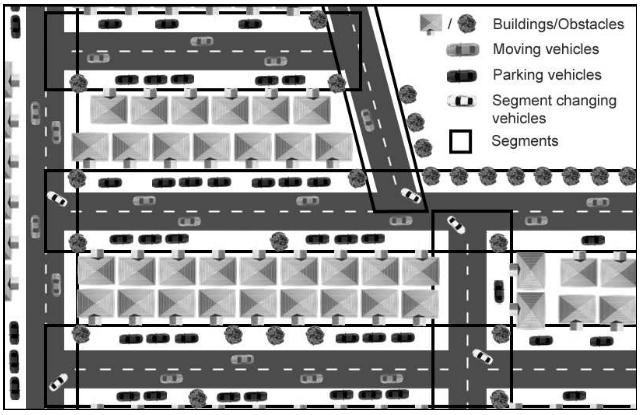

Fig. 3. Interacting P2P/C2C segments

Within one segment, all vehicles are able to communicate directly via DSRC. Consequently, they can exchange information needed to keep up the P2P topology in an adequate manner. As described in section 3.2, the addresses of vehicles are mapped onto identifiers using hash functions. An information item is always placed on the vehicle whose identifier is most similar to the information ones. Hence, it is known at any time on which car to store an information item. Thereby, a validity period can be specified. Furthermore, it is known where to search for a particular information item in order to retrieve, modify or delete it.

Vehicles know their position, their direction and their velocity. This knowledge is provided by in-car sensors and navigation systems. Moreover, vehicles know whether they cross segment borders by using digital maps.

As soon as a vehicle approaches a new segment, it contacts the first vehicle of the corresponding P2P/C2C network. This is performed with the aid of beacon messages. In doing so, a segment-changing vehicle is able to obtain all required information to join the new topology. Additionally, it announces its departure from the old network to avoid overhead from failure recovery. Figure 3 highlights vehicles crossing segment borders using a lighter color.

4.3 P2P/C2C Segment Interactions

By default, information is published, and then managed, in one segment. Yet, a piece of information might be of relevance beyond the borders of that segment. In the following, we will refer to information, whose scope is specified for more than one segment, as *segment overlapping information*.

Segment overlapping information has to be forwarded between neighboring P2P/C2C networks. To this end, a vehicle crossing segment borders first searches for that kind of information in its old segment. Afterwards, it publishes found items in the new segment. Finally, it marks these items as being forwarded by modifying them. Thus, there is no redundant forwarding. To be able to specify segment overlapping information, a corresponding information type is defined. Hence, segment overlapping information becomes searchable and markable for all vehicles. If segment overlapping information is modified or deleted, the forwarded mark is reset again and, in case, set deleted. This way, the information item is passed to adjacent segments again. There, the responsible peer has to overwrite or to delete the existing information.

Segment overlapping warning messages are an exception in this context. If they are of relevance in neighboring segments as well, these messages are broadcasted. It is not possible to wait until the next vehicle leaves the segment. If warning messages are still valid in the next but one-segment warnings have to be flooded. However, as this would correspond to at least one kilometer in DSRC, it can be assumed that the scenario is going to emerge rather seldom.

The introduced approach corresponds to the C2C communication system. The segmentation of the vehicular network and all other mechanisms are applied and realized at the communication system layer. C2C applications access the functionalities using publish, delete, search and modify primitives provided by its interface. When an information item is published, its scope can be specified in meters, for example. This parameter is passed to the interface when calling the publish method. The communication system selects relevant segments and publishes the information accordingly. Thus, C2C applications can be developed without knowledge about the segmentation of the vehicular network.

4.4 Publish/Subscribe based on P2P/C2C Segments

To avoid the need for building and maintaining a tree structure, our approach suggests the application of P2P/DHT-based Pub/Sub. Thereby, Pub/Sub is set up on top of the described C2C/P2P segments. In doing so, it is not necessary to maintain a PST as the segments' DHTs can be used to realize Pub/Sub. In addition, there is no need for a vehicular distribution tree for P2P/DHT-based Pub/Sub: Scalability poses no problem, as each segment is relatively small. Finally, subscriptions and notifications can be exchanged in an adequate manner, as exchange is limited to the segments.

However, application level subscriptions need to be valid not only in one segment. Therefore, a vehicle's subscriptions are published in its new segment on crossing segment borders. Furthermore, they are deleted in the old segment. As information items are forwarded between segments according to their relevance range, it is sufficient to manage subscriptions in the originators current segment. They will be checked against all relevant information items.

From the application level point of view, a subscription is issued only once indicating a validity period. The Pub/Sub functionality refers to the communication system as well. The system updates subscriptions for the duration of their validity when moving into a new segment. If no validity period is specified, subscriptions will be updated until a default time period expires. C2C applications must renew their subscriptions after expiration of the respective validity period. Hence, the Pub/Sub mechanism can be used without knowledge of the VANET segmentation, but expiration has to be kept in mind.

We will now review the sample application mentioned in section 2.2 to illustrate the described functionalities: In urban environments, parking lots are sometimes hard to find. A

(a) Example subscription

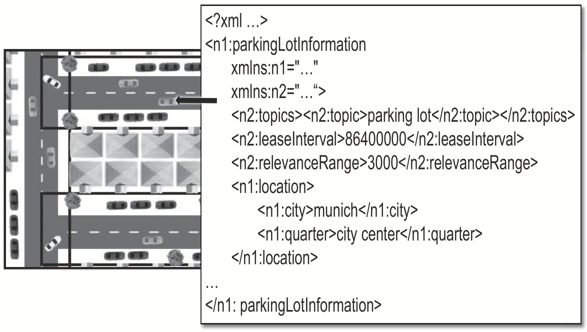

(b) Example Topic-hashed Information

Fig. 4. Publish/Subscribe based on P2P/C2C Segments

C2C application could issue a subscription on parking lots as soon as the driver claims he wants to park his car. Thereby, the driver might specify the region he wants to park his car in by selecting it on the navigation system's screen, for example. This region is set as additional attribute within the subscription. The subscription is then published using the topic "parking lot" (see Fig. 4(a)).

If information producers use the same topic, the subscription is placed at the vehicle also storing parking lot information. Therefore, topics must be known in advance. For example, the application helping drivers to find parking lots might also be able to detect when a car leaves a parking lot. It might publish this each time using the topic "parking lot" (see Fig. 4(b)). Car parks would need to be informed about this consensus to allow for the integration of their capacities. Also, the C2C application will only work if parking lot information is deleted as soon as a parking lot is taken. Of course, a lot of hurdles arise in this context, but they are not in the scope of this paper.

After publishing the subscription, the responsible vehicle checks it against existing parking lot information. If there are parking lots available in the desired area, the searching vehicle is notified immediately. If there is no free space, the vehicle is informed as soon as a new parking lot is published or a car park changes his capacities. The vehicle responsible for the information topic "parking lot" stores the information or alters it. Afterwards, it checks existing subscriptions and performs necessary notifications.

The given example shows basic aims of the proposed approach. There may be a lot of free parking lots in a city. Yet, there is no need to repeatedly transmit each single parking lot information item within its entire relevance range to ensure that a searching vehicle receives them. Furthermore, a searching vehicle does not need to listen to all messages sent over the air interface in order to retrieve the right parking lot messages. Finally, parking lot information can not only be provided but also modified and deleted and interested vehicles can be notified about these events.

Likewise, a warning might be of relevance only for trucks. Thereby, its relevance range might be a region trucks pass rather seldom. The warning would not need to be repeated at high frequencies to make sure trucks are warned early enough. In addition, it would not have to be retransmitted during its entire validity duration although trucks do not turn up often.

4.5 Discussion

Within one segment, vehicles are able to communicate directly as segment sizes are chosen according to the available communication range. Consequently, P2P stabilize routines can be performed fast. Furthermore, P2P/DHT-based Pub/Sub can be utilized, as there is no need to keep up a tree structure for scalability reasons. Finally, application specific data can be exchanged efficiently. Apart from that, single pieces of information are managed in a distributed manner only within relevant segments. However, the segmentation leads to the problem that P2P fault tolerance mechanisms are bounded to the segments. Within one segment, a car's neighbors assume the responsibility for its range of identifiers and corresponding information. This is necessary as soon as a vehicle leaves the network or fails, for example, from being parked. Therefore, information is not lost as long as one vehicle equipped with a C2C communication unit is left in the P2P/C2C segment. This presumes the car is able to store all information items. As segments are comparatively small and the approach is concerned with vehicles only, storage capacity is assumed to be not as critical.

Yet, all information is lost as soon as the last equipped vehicle leaves a segment. Subscriptions of vehicles parking within that segment are lost as well. To handle this, we suggest that the relevance range of every information item covers adjacent segments by default. In addition, subscriptions have to be passed to neighboring segments using the described forwarding mechanisms. The default relevance range might even be selected wider depending on the preferred level of fault tolerance.

Consequently, information items and "backup subscriptions" would always be managed in adjacent segments. If a car recognizes that a P2P/C2C network it wants to join does not exist anymore, it is able to request corresponding information items and subscriptions. Thereby, information items and subscriptions can be recovered until one segment and all its adjacent segments fail. Hence, a failure would not take place as long as one equipped vehicle is left within an area much bigger than the available communication range. This might even be of advantage regarding given deployment issues of vehicular communication technologies (e. g. [41]).

Yet, a clear disadvantage lies in the fact that structured P2P networks imply overhead. This overhead is caused by messages needed to keep up the network topology when peers join and leave the network. The smaller the geographical perimeter of a P2P/C2C network is, the fewer vehicles will join and leave it at its borders during a certain time period. At least, when assuming an equal distribution of vehicles. Nevertheless, joins and leaves will occur at high frequencies. Even in urban regions, a vehicle driving 60 km/h would spend only about one minute within a segment, assuming there are no traffic lights to stop at.

5 Evaluation Environment

Our next step is to realize the proposed approach in a framework for C2C communication in order to evaluate it. In particular, this will enable us to compare the possibility to repeatedly broadcast location- or application specific data (see Section 3) with the overhead caused by P2P topology information.

This section first presents the VISSIM/VCom environment. It is necessary to realistically simulate factors like car movements and wireless vehicular communications as they strongly influence evaluation results. The VISSIM/VCom environment allows us for advanced traffic simulation as well as for C2C communication simulation. Furthermore, we describe our implementation of the black ice warning scenario, which is based on state-of-the-art information provisioning. In addition, we present results of its evaluation. These results built the groundwork for comparing our approach to state-of-the-art information provisioning when evaluating it.

5.1 VISSIM/VCom

VISSIM by PTV Vision [42] is a time step-oriented and behavior-based simulator for urban and highway traffic simulation including pedestrians, cyclists and motorized vehicles. A simulation's result is online a traffic flow animation – as may be required two- or three-dimensional. VISSIM visualizes the animation using a graphical user interface (see Fig. 5). Furthermore, different traffic characteristics can be logged and analyzed offline.

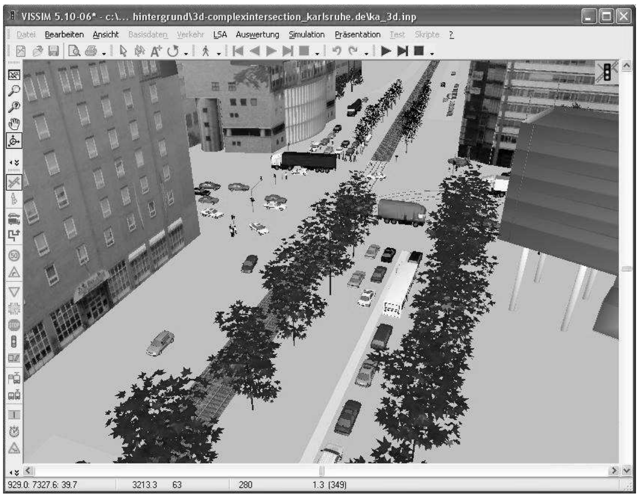

Fig. 5. VISSIM graphical user interface and a 3D traffic model

Traffic-dependent control logic is modeled using external modules for traffic signaling and control. Concerning the traffic flow model, so-called driver-vehicle entities move through the transport network. Each driver and his behavioral patterns are associated with a dedicated car. This way, driving behavior and a vehicle's technical possibilities correlate. A driver-vehicle entity is characterized by different attributes. Besides the vehicle's technical specification, they span psychophysical perception barriers like estimation abilities or the readiness to assume risk, for example. Moreover, dependencies between different driver-vehicle entities influence an entity's behavior.

VISSIM is controllable using the graphical user interface shown in Figure 5. In addition, VISSIM provides a *Component Object Model (COM)* interface. Using this interface, external software components are able to control VISSIM and to request its data.

To simulate C2C communication, we are working with *VCom* [43]. VCom is a VISSIM module, which has recently been developed by the VISSIM producer PTV in collaboration with the Technical Universities of Karlsruhe and Munich. The motivation behind VCom is that a network simulator becomes a bottleneck when coupling it with VISSIM in case there is a multitude of nodes. Hence, VCom is a module for simulating communication, which does not use a network simulator to simulate the communication process. Rather, the network simulator *ns-2* [44] was used to generate communication statistics. VCom accesses these statistics to simulate C2C communication.

Accordingly, the VISSIM/VCom system architecture consists of three components. It comprises the traffic simulator VISSIM, the VCom module for simulating the wireless communication processes and interfaces for integrating C2C/C2I applications (see Fig. 6).

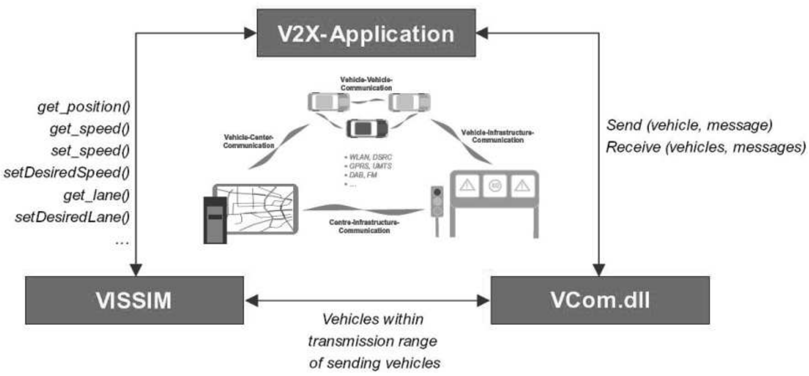

Fig. 6. The VISSIM/VCom system architecture [43]

5.2 Implementation of a Simple Application using State-of-the-art Information Provisioning

Using VISSIM/VCom, we implemented the introduced black ice warning scenario. Thereby, we applied the state-of-the-art repeated broadcasting mechanisms discussed in section 3 to realize a warning's persistent availability.

Regarding the traffic network, black ice is modeled as a dedicated part of a lane on a 4-lane motorway. There are two lanes for each driving direction, whereas a passenger car's desired speed is 140 km/h or 87 mph respectively. Trucks are assigned a desired speed of 80 km/h equating to 50 mph. The black ice area is a slow down area covering $10 \times 3,5$ $meters$ ($33 \times 11,5$ $ft.$). The area enforces a vehicle to slow down sharply when passing. This is realized by setting the vehicles desired speed to 30 km/h – 18,6 mph – as soon as the black ice area is 50 $meters$ (164 $ft.$) ahead. 50 $meters$ is an estimated value derived from the distance the glassy surface is assumed to be visible factoring the driver's reaction time. In addition, the vehicle's braking deceleration is assigned 8 m/s^2 resulting in hard braking. On passing the black ice area, the desired speed stays the same but braking deceleration decreases to 1 m/s^2 as a result of the bad road condition.

As soon as the black ice application detects a vehicle braking this way, it forces the vehicle to broadcast a black ice warning. Thereby, the vehicle passes its current location's coordinates, the affected lane and a default warn distance of 1 km or 3300 ft. respectively.

Every vehicle receiving a warning checks whether it is affected. That is, it calculates whether it is on the referred lane driving towards the hazard. If so, it decreases its speed.

Apart from that, every vehicle within warning distance radius becomes a repeater. In this context, also vehicles driving on the opposite lane or veering away from the hazard are concerned, as they help to carry the warning and to keep it alive. Each repeater retransmits a received warning according to a variable repetition frequency starting from the warning's reception time. Thereby, a repeater calculates his current distance to the hazard each time to be able to stop broadcasting after leaving the warning distance radius.

Given the described parameters, vehicles passing the black ice area uninformed traverse it driving about 70 km/h (43,5 mph). They are unable to decelerate to their desired speed. If an affected vehicle's speed is set to 75 km/h after it was warned, it is able to moderate its speed to 30 km/h as desired.

5.3 Evaluation of the Simple Application

We are aiming at comparing the overhead caused by P2P topology information with that of state-of-the-art information provisioning. Hence, evaluation results of the presented application are a major groundwork for our own evaluation.

First, it is of relevance to determine the retransmission interval needed to guarantee an acceptable rate of warned vehicles within warning distance. The percentage of warned vehicles within warning distance depends on the number of vehicles equipped with a C2C unit as well as on the warnings' retransmission interval. In our traffic network, we modeled a traffic load of 2000 vehicles per hour for each driving direction. That is, there are 2000 vehicles per hour distributed over two lanes. In addition, we defined an average rate of 25 percent C2C unit-equipped vehicles. Consequently, there are about 500 vehicles able to communicate wirelessly per hour and direction.

Figure 7 charts the percentage of warned vehicles within warning distance per time using varying retransmission intervals. Obviously, it is not possible to keep the black ice warning alive when repeating vehicles retransmit the warning every 60 seconds. Also when applying a retransmission interval of 10 seconds, we tend to have about 10 percent of the vehicles within hazard zone uninformed (see Fig. 7(a)). Yet, we are able to maintain a predominant information rate of 100 percent using a retransmission interval of 5 seconds (see Fig. 7(b)).

Figure 8 illustrates the number of transmitted warnings required to achieve the named rates. The visible steps result from the fact that different vehicles receive a broadcasted warning only fractions of milliseconds delayed. Thus, the repeating vehicles' iteration schedules reappear quasi-simultaneous for the given time axis resolution.

Regarding the measured number of warnings transmitted, one can state that there are about 150 messages transmitted per minute when applying a retransmission interval of 5 seconds. An average density of 500 equipped vehicles per hour and direction corresponds

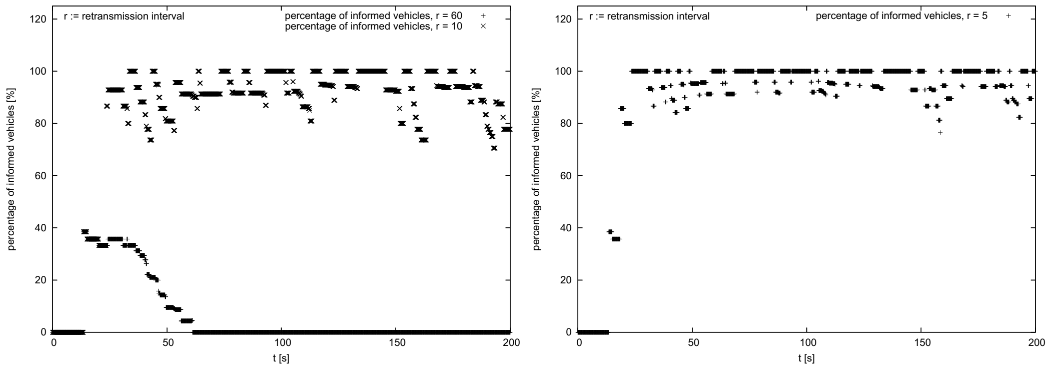

(a) Retransmission interval: 10 and 60 seconds

(b) Retransmission interval: 5 seconds

Fig. 7. Percentage of informed vehicles per time

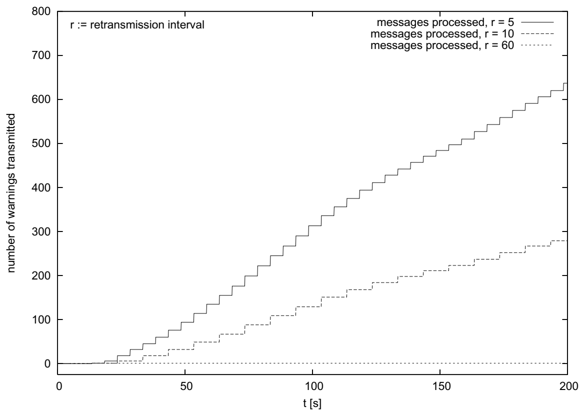

Fig. 8. Number of messages transmitted per time and retransmission interval

to a density of 16, 66 vehicles per minute. This results in an average transmission load of 9 messages per minute per vehicle (150 ÷ 16, 66) or 0, 15 messages per second per vehicle respectively. The load refers to each vehicle traversing the warning distance radius. Thereby, it averages the influences of message loss, message delay or varying regional distributions of C2C equipped vehicles, which affect a vehicle's starting time of becoming a repeater.

Besides sending, cars have to check whether they already know a warning for each incoming one as broadcasted warnings do not specify receivers. Figure 9 shows the number of processed messages for each retransmission interval. We will again consider the retransmission interval of 5 seconds as it leads to an adequate warning ratio: There are about 1250 messages processed per minute. This results in 75 processed messages per minute per vehicle corresponding to 1, 25 messages per second. The higher number of processed messages is due to the possibility that vehicles are in the communication range of multiple repeaters. This results in multiple redundant receptions.

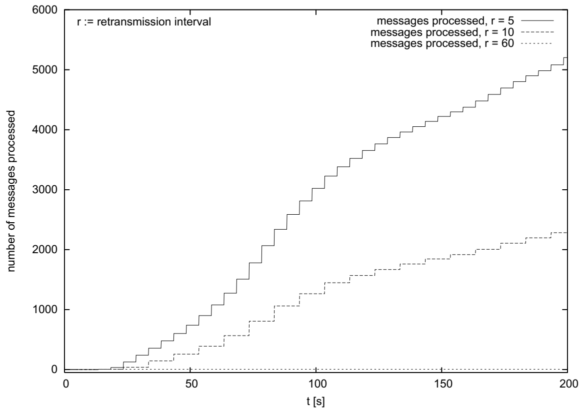

Fig. 9. Number of messages processed per time and retransmission interval

To conclude, we will compare the loads of 0, 15 sent messages and 6, 7 processed messages per vehicle per second with our approach's load. Thereby, we will have to consider that the growth of the depicted message load will show linear behavior for additional information items (multiple available parking lots, gas stations and their prices et cetera). Apart from that, overhead may be sustainable to a certain degree regarding our approach's further functionalities of distributed information handling and Pub/Sub communication.

6 Conclusions and Future Work

In this paper, we focused on the application of structured P2P algorithms and P2P/DHT-based Pub/Sub for C2C communication systems. The overall goal is to utilize C2C communication to support active safety and deployment applications. In this context, wireless communication is particularly powerful because it allows for the fast provision of regionally relevant information.

To support different kinds of C2C applications, a communication system should manage information like black ice warnings in a distributed manner during their validity. Furthermore, the system should allow for publishing information with respect to certain geographical regions and for handling them later on. In addition, it should provide Pub/Sub messaging to facilitate information filtering and the tracking of changes in the information environment. We gave an overview on research areas and technologies related to these functionalities. Thereby, we arrived at the conclusion that distributed information management and Pub/Sub provision poses an open challenge in C2C networks.

Against this background, we presented our approach. The approach utilizes mechanisms of structured P2P algorithms for distributed information management in VANETs. These mechanisms are basically hard to apply because vehicular networks cannot provide communication properties similar to those in wired networks. To meet resulting problems and to allow for location aware information management, our approach divides the vehicular network into segments. These segments form separate, interacting P2P networks. Cars manage information of their own segment and exchange them between adjacent segments if necessary. This is the case if information is of further relevance or if the level of fault tolerance shall be increased.

In addition, the Pub/Sub communication primitive is applied. The primary challenge in this context is that building and maintaining a PST proves to be difficult in VANETs. Also P2P/DHT-based Pub/Sub solutions rely on a distribution tree for scalability reasons. To handle this, Pub/Sub is built on top of the P2P/C2C segments' DHTs. In doing so, P2P/DHT-based Pub/Sub can be used whereas scalability problems are avoided. To provide seamless Pub/Sub functionalities, subscriptions follow their originator when passing segment boarders. Furthermore, they are exchanged between adjacent segments in order to enhance fault tolerance.

Subsequently, we presented our evaluation environment. Thereby, we described the recently developed module VCom. VCom allows for simulating C2C communication within the traffic simulator VISSIM. Furthermore, we introduced our implementation of the black ice warning scenario, which is based on state-of-the-art information provisioning. Finally, we presented results of the implementation's evaluation.

Our next step is to realize the proposed approach in a framework for C2C communication in order to evaluate it. The framework is supposed to provide an interface offering publish, delete, search, modify and subscribe primitives besides send and receive. This way, different C2C applications can be implemented on top and can access our solution's functionalities. In combination with our simulation environment, the framework will allow us to analyze our solution. In particular, it will enable us to draw conclusions about the overhead caused by P2P topology information. To this end, we will compare the required load with the evaluated load necessary for state-of-the-art information provisioning.

References

1. European commission / directorate general energy and transport (2007), *Road safety evolution in the eu.* CARE (EU road accidents database) or national publications.
2. M. Torrent-Moreno, P. Santi, and H. Hartenstein (2005), *Fair sharing of bandwidth in vanets.* Proceedings of the 2nd ACM international workshop on vehicular ad hoc networks, September, pp. 49–58.
3. C. E. Perkins and E. M. Royer (1999), *Ad-hoc on-demand distance vector routing.* Second Workshop on Mobile Computing Systems and Applications, February, pp. 90–100.

4. E. Natsheh and T. C. Wan (2008), *Adaptive and fuzzy approaches for nodes affinity management in wireless ad-hoc networks*. Mobile Information Systems, Vol. 4 (No. 4), pp. 273–295.
5. C. Lochert, H. Hartenstein, J. Tian, H. Fussler, D. Hermann, and M. Mauve (2003), *A routing strategy for vehicular ad hoc networks in city environments*. Proceedings of the Intelligent Vehicles Symposium, June, pp. 156–161.
6. E. Royer and C. Perkins (2000), *Multicast ad hoc on- demand distance vector (maodv) routing*. IETF MANET Working Group Internet Draft.
7. C. Schwingenschlogl and T. Kosch (2002), *Geocast enhancements of aodv for vehicular networks*. SIGMOBILE Mobile Computing and Communications Review, Vol. 6 (No. 3), pp. 96–97.
8. ASTM International (2003) *ASTM E2213, Standard Specification for Telecommunications and Information Exchange Between Roadside and Vehicle Systems – 5 GHz Band Dedicated Short Range Communications (DSRC) Medium Access Control (MAC) and Physical Layer (PHY) Specifications*.
9. Dedicated short range communications (DSRC) home. Available from World Wide Web: `http://www.leearmstrong.com/DSRC/DSRCHomeset.htm`, cited 27 February 2009.
10. L. Wischoff, A. Ebner, H. Rohling, M. Lott, and R. Halfmann (2003), *Sotis - a self-organizing traffic information system*. 57th IEEE Vehicular Technology Conference, April.
11. L. Wischof, A. Ebner, H. Rohling, M. Lott, and R. Halfmann (2003), *Adaptive broadcast for travel and traffic information distribution based on inter-vehicle communication*. Proceedings of the Intelligent Vehicles Symposium, June, pp. 6–11.
12. X. Yang, L. Liu, N. H. Vaidya, and F. Zhao (2004), *A vehicle-to-vehicle communication protocol for cooperative collision warning*. The First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services, August, pp. 114–123.
13. L. Briesemeister, L. Schafers, and G. Hommel (2000), *Disseminating messages among highly mobile hosts based on inter-vehicle communication*. Proceedings of the Intelligent Vehicles Symposium, October, pp. 522–527.
14. N. Chand, R. C. Joshi, and M. Misra (2007), *Cooperative caching in mobile ad hoc networks based on data utility*. Mobile Information Systems, Vol. 3 (No. 1), pp. 19–37.
15. B. Williams and T. Camp (2002), *Comparison of broadcasting techniques for mobile ad hoc networks*. Proceedings of the 3rd ACM international symposium on Mobile ad hoc networking & computing, June, pp. 194–205.
16. T. Imielinski and J. Navas (1996), *Gps-based addressing and routing*. Tech. rep., The Internet Society.
17. C. Maihöfer, W. Franz, and R. Eberhardt (2003), *Stored Geocast*. Proceedings of Kommunikation in Verteilten Systemen (KiVS), February, pp. 257–268.
18. A. Oram (2001) *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*. O'Reilly Media, Inc.
19. Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, and S. Shenker (2003), *Making gnutella-like p2p systems scalable*. Conference of the Special Interest Group on Data Communication, August, pp. 407–418.
20. I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan (2003), *Chord: A scalable peer-to-peer lookup service for internet applications*. IEEE/ACM Transactions on Networking, Vol. 11 (No. 1), pp. 17–32.
21. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim (2005), *A survey and comparison of peer-to-peer overlay network schemes*. IEEE Communications Surveys & Tutorials, Vol. 7 (No. 2), pp. 72–93.
22. G. Gehlen, F. Aijaz, Y. Zhu, and B. Walke (2007), *Mobile p2p web services using sip*. Mobile Information Systems, Vol. 3 (No. 3–4), pp. 165–185.
23. A. Tveit (2001), *Peer-to-peer based recommendations for mobile commerce*. Proceedings of the 1st ACM international workshop on Mobile commerce, pp. 26–29.
24. B. N. Miller, J. A. Konstan, and J. Riedl (2004), *Pocketlens: Toward a personal recommender system*. ACM Transactions on Information and System Security, Vol. 22 (No. 3), pp. 437–476.

25. W. Wahlster (2004), *Advances in Artificial Intelligence*, chap. Smartweb: Mobile applications of the semantic web, pp. 50–51. Springer Berlin (Heidelberg).
26. Q. Huang, R. Miller, P. McNeille, D. Dimeo, and G. Roman (2002), *Development of a peer-to-peer collision warning system*. Tech. rep., Washington University, Ford Research Lab.
27. C. Lochert, J. Rybicki, B. Scheuermann, and M. Mauve (2008), *Scalable data dissemination for inter-vehicle-communication: aggregation versus peer-to-peer*. it - Information Technology, Oldenbourgh, Vol. 50 (No. 4), pp. 237–242.
28. J. Rybicki, B. Scheuermann, W. Kiess, C. Lochert, P. Fallahi, and M. Mauve (2007), *Challenge: Peers on wheels – a road to new traffic information systems*. Proceedings of the 13th ACM international conference on Mobile computing and networking, September, pp. 215–221.
29. G. Muehl, L. Fiege, and P. R. Pietzuch (2007), *Distributed Event-Based Systems*. Springer Berlin (Heidelberg).
30. M. Castro, P. Druschel, A. Kermarrec, and A. Rowstron (2002), *Scribe: A large-scale and decentralized application-level multicast infrastructure*. IEEE Journal on Selected Areas in Communications, Vol. 20 (No. 8), pp. 100–110.
31. A. Gupta, O. D. Sahin, D. Agrawal, and A. Abbadi (2004), *Meghdoot: content-based publish/subscribe over p2p networks*. Proceedings of the 5th ACM/IFIP/USENIX international conference on Middleware, October, pp. 254–273.
32. R. Baldoni and A. Virgillito (2005), *Distributed event routing in publish/subscribe communication systems: A survey*. Tech. rep., Dipartimento di Informatica e Sistemistica, Universit di Roma.
33. A. Carzaniga, D. S. Rosenblum, and A. L. Wolf (2000), *Achieving scalability and expressiveness in an internet-scale event notification service*. Proceedings of the nineteenth symposium on Principles of distributed computing, July, pp. 219–227.
34. G. Banavar, T. Chandra, B. Mukherjee, J. Nagarajarao, R. E. Storm, and D. C. Sturman (1999), *An efficient multicast protocol for content-based publish-subscribe systems*. Proceedings of the 19th International Conference on Distributed Computing Systems, June, pp. 262–272.
35. R. S. Filho, C. R. D. Souza, and D. F. Redmiles (2003), *The design of a configurable, extensible and dynamic notification service*. Proceedings of the 2nd international workshop on Distributed event-based systems, June, pp. 1–8.
36. T. Sivaharan, G. Blair, A. Friday, M. Wu, H. Duran-Limon, P. Okanda, and C. F. Sorensen (2004), *Cooperating sentient vehicles for next generation automobiles*. Workshop on Applications of Mobile Embedded Systems, June.
37. T. Sivaharan, G. Blair, and G. Coulson (2005), *GREEN: A configurable and reconfigurable publish-subscribe middleware for pervasive computing*. On the Move to Meaningful Internet Systems 2005: CoopIS, DOA, and ODBASE, pp. 732–749. Springer Berlin (Heidelberg).
38. L. Mottola, G. Cugola, and G. P. Picco (2008), *A self-repairing tree topology enabling content-based routing in mobile ad hoc networks*. IEEE Transactions on Mobile Computing, Vol. 7 (No. 8), pp. 946–960.
39. G. Cugola and G. P. Picco (2006), *Reds: a reconfigurable dispatching system*. Proceedings of the 6th international workshop on Software engineering and middleware, November, pp. 9–16.
40. J. Luo and J. Hubaux (2006), *Embedded Security in Cars*, chap. A Survey of Research in Inter-Vehicle Communications, pp. 111–122. Springer Berlin (Heidelberg).
41. K. Matheus, R. Morich, and A. Lübke (2004), *Economic background of car-to-car communication*. Proceedings of the 2. Braunschweiger Symposium Informationssysteme für mobile Anwendungen, October.
42. The vissim traffic flow simulation software. Available from World Wide Web: `http://www.ptvamerica.com/vissim.html`, cited 27 February 2009.
43. S. Assenmacher, M. Killat, F. Schmidt-Eisenlohr, and P. Vortisch (2008), *A simulative approach for the identification of potentials and impacts of v2x-communication*. Proceedings of the 15th World Congress on Intelligent Transport Systems, November, World Congress on ITS.
44. The network simulator - ns-2. Available from World Wide Web: `http://www.isi.edu/nsnam/ns/`, cited 27 February 2009.