

# Challenge: Peers on Wheels – A Road to New Traffic Information Systems

Jedrzej Rybicki

Björn Scheuermann

Wolfgang Kiess

Christian Lochert

Pezhman Fallahi

Martin Mauve

Computer Networks Research Group  
Heinrich Heine University  
Düsseldorf, Germany

{rybicki, scheuermann, kiess, lochert, fallahi, mauve}@cs.uni-duesseldorf.de

## ABSTRACT

In the context of vehicular ad-hoc networks (VANETs), a number of highly promising convenience applications have been proposed. These include collecting and distributing information on the traffic situation, distributed monitoring of road and weather conditions, and finding available parking places in a distributed, cooperative manner. Unfortunately, all of these applications face major problems when a VANET is used as a means to distribute the required information. In particular a large number of vehicles needs to be equipped with dedicated VANET technology before these applications can provide a useful service. Even if customers were willing to purchase a system which is not immediately useful, it would still take quite some time until the required density of equipped cars is reached. In contrast, affordable always-on mobile Internet access is already mainstream. Such Internet connectivity could be used to build the proposed applications in a different fashion: by using peer-to-peer communication, essentially creating a peer-to-peer network of cars sharing traffic information. This allows to overcome the limitations of VANETs, while it preserves their key benefits of decentralization and robustness. In this paper, we describe the technical challenges that arise from such an approach, point out relevant research directions, and outline possible starting points for solutions.

## Categories and Subject Descriptors

C.2.1 [Computer System Organization]: Computer - Communication Networks—*network architecture and design, distributed networks*; C.2.4 [Computer Systems Organization]: Computer - Communication Networks—*distributed systems, distributed applications*; H.4.3 [Information Systems]: [Communications Applications]

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom'07, September 9–14, 2007, Montréal, Québec, Canada.  
Copyright 2007 ACM 978-1-59593-681-3/07/0009 ...\$5.00.

## General Terms

Algorithms, Design

## Keywords

Peer-to-peer Networks, Traffic Information Systems, Car-to-car Communication, VANET

## 1. INTRODUCTION

The application of mobile communication technology to support road traffic constitutes a challenging, but at the same time very promising working area for research and development. A whole community has formed around the questions that vehicular communications and, in particular, vehicular ad hoc networks (VANETs) pose. Consisting of public authorities, academia, and car manufacturers [4, 11, 31], this community fosters the use of communication technology to enhance driving security and comfort. Proposed applications reach from the reduction of road casualties by means of brake warning, intersection assistance, or collision avoidance systems [33] to offering guidance to available parking lots [5], discovering the traffic situation on a planned route [32], and coordinating car flow and traffic lights [9, 31].

In this paper, we focus on non-safety related applications that can be subsumed under the term distributed *traffic information systems* (TIS). In these systems, the participating cars are not only consumers of information but at least some of them also produce information by sharing their observations.

Traffic information systems require communication among many participants over relatively large distances that can span some ten kilometers in the case of a city scenario up to some hundred kilometers on highways. Thus, the communication requirements of TIS applications are quite challenging: continuously updated data spread over a high number of network nodes is to be made available to many vehicles in a relatively large area.

Undoubtedly, many of the proposed applications are highly useful and very desirable, but market introduction and technological hurdles of VANET-based TIS solutions are high. VANETs, where cars communicate directly with Wi-Fi-like equipment [19], are well suited if local communication is needed. However, when it comes to communication between many partners over longer distances, this kind of networks suffer from very bad connectivity until a significant amount of vehicles is equipped with this technology. This has been demonstrated, e.g., in [21], where upper bounds on the

speed and reliability of information dissemination in a VANET city environment with delay tolerant network-like opportunistic data exchange have been studied. The conclusion is that sufficient performance is simply not possible with penetration ratios that are realistic within the near future—at least not without additional infrastructure support. So, the number of necessary network nodes for quick and reliable information *distribution* exceeds the number of participating cars required to *collect* enough information for a working service by far. Moreover, even given a sufficiently high equipment ratio, there are also the inherent capacity limits of wireless multihop communication, first formalized by Gupta and Kumar [15]. The limited transport capacity of wireless multihop networks most likely impedes distributing detailed data continuously to many interested parties in a large surrounding. This is further aggravated by the fact that a substantial fraction of the bandwidth has to be kept free in order to guarantee working safety applications [34].

Thus, it becomes clear that the bottleneck in terms of connectivity and capacity is the VANET. However, the decentralized character of VANETs for distributing traffic information is very appealing: a cooperative approach where every participant contributes, distributes, and consumes information, where no central institutions and no central infrastructure are necessary, where, consequently, the absence of single points of failure promises a robust service, without recurring fees for the users—all this is obviously highly desirable. But do these features really require VANETs, do they require wireless multihop networks?

Cheap mobile Internet access, be it via 3G, GPRS, WiMax, WiFi, or any other technology is already widespread. UMTS flatrates, for instance, are available in many countries and are rapidly getting cheaper. In short: always-on mobile Internet access will soon be common, long before VANET technology in cars will be deployed. In harsh contrast to wireless multihop networks, when using infrastructure-based communication the connectivity, latency, and bandwidth are almost independent from the physical distance. There are no separate network partitions, and differences in bandwidth or latency, if relevant at all, will only depend on the access technologies.

In this paper, we discuss how infrastructure-based communication might be leveraged to build a distributed, cooperative TIS—only requiring every participant to maintain an Internet connection, but preserving all the benefits of the proposed VANET solutions: a decentralized, scalable, and cooperative approach. We envision such a solution based on distributed hash tables, i. e., on peer-to-peer technology: a huge peer-to-peer network of vehicles. Such a system, implementing well-designed distributed data structures and algorithms allows to build the TIS applications that are discussed for VANETs without being impeded by VANET insufficiencies when it comes to communicating over longer distances. This may result in a radically shorter market introduction time for car-to-car distributed traffic information sharing.

Using infrastructure-based communication of course also yields the possibility of a centralized system, with all the well-known advantages and drawbacks of such an approach. Compared to a peer-to-peer approach, a centralized system poses different technical challenges (see, e. g., [18]). However, we argue that a distributed approach is preferable if it is able to deliver the same service in a comparable quality—not least because it avoids the effort for setting up and maintaining the central components. In this paper, we focus on the peer-to-peer approach.

Our approach poses completely new challenges, very different from what has so far been examined in the car-to-car community. Up to now, a central aspect of almost all considerations was the physical position and the movement of the nodes, the resulting network topology, and the implications for the protocols and appli-

cations. This is no longer important when an infrastructure-based peer-to-peer network is used. However, the arising requirements of our approach are also very different from the focus of existing research on peer-to-peer systems and the corresponding distributed data structures. Data provided and requested by nodes will typically be structured and highly correlated, such as the traffic situation on a sequence of roads from a starting point to a destination. Furthermore, the update and request frequency is likely to be much higher than in existing peer-to-peer applications. As a consequence it is necessary to tailor the distributed data structures to the unique environment created by traffic information systems.

This paper is organized as follows. In Section 2, we give a short introduction to peer-to-peer systems and the underlying data structures. Section 3 sketches a first approach to the problem that already reveals some of the major challenges that come with the proposed change of paradigms. We extend this first system in Section 4. Finally, we outline further possible applications, some of them significantly exceeding the scope of a “classical” traffic information system in Section 5. The paper is concluded in Section 6.

## 2. PEER-TO-PEER NETWORKS

Peer-to-peer overlay networks are virtual communication structures logically established over a physical network such as the Internet. They emerge by self-organization of peers in absence of a central supervising entity. Network nodes cooperate and share resources, like data or computational power. In the following, we will use the terms car, node and peer interchangeably.

Peer-to-peer systems distribute the resources among participants, in contrast to a client-server architecture where they are hosted centrally. Thus, the main challenge in such networks is to locate the resources, i. e., to find a file identified by a name or to find a car that knows about the traffic situation on a given road. The first peer-to-peer systems used index servers, but this has obvious drawbacks in terms of performance and robustness [29]. Modern systems are so-called structured networks, and are often also referred to as *distributed hash tables* (DHT). Well-known examples are Chord [30], CAN [27], or Tapestry [35]. In these peer-to-peer systems, resources are mapped to peers by means of a hash function. Each peer knows a subset of the other peers, forming an overlay network along which lookups can be routed. In most of the proposed structured networks, such a lookup has logarithmic complexity. For a more detailed overview of peer-to-peer networks see, e. g., [2, 22].

To make the discussion a little more specific, let us look at the concrete example of Chord [30]. Chord’s basic structure is a circular ID space, the Chord ring, sketched in Figure 1. Nodes and resources are both mapped uniformly to this ID space by means of a hash function. Each node hosts the resources with keys in the range between its predecessor’s and its own ID. Every node knows its successor, so queries can be routed clockwise around the ring. In addition, the nodes maintain further links, the so-called *fingers*. For one of the nodes, with ID 69105, these are indicated by dashed arrows in the figure. The fingers are shortcuts over  $1/2$ ,  $1/4$ ,  $1/8$ , etc. of the ID space, pointing to the respective nodes with the next higher ID. This structure allows to reduce the distance to the destination by at least one half in every step, resulting in the above mentioned logarithmic lookup complexity.

Assume, for instance, that the peer with ID 69105 in our example figure is looking for a resource identified by key 4004. It uses the nearest known node in anti-clockwise direction from 4004 as the next hop. This is, through the  $1/4$  finger, 42 in our example. Node 42, following the same principle, will forward the request to node 1138, which in turn knows that its successor, 4711, hosts the sought-after key.




Figure 1: A Chord ring.

Some research on enabling peer-to-peer in cellular networks has already been undertaken. E. g., in [16] it has been shown that existing peer-to-peer filesharing systems work in mobile networks (3G, GPRS) and exhibit satisfactory performance.

### 3. A FIRST SYSTEM: TRAFFIC INFORMATION SHARING

Let us now sketch a first, naive peer-to-peer based traffic information system that allows cars to exchange information about the current traffic status, e. g., in a city or in a network of highways. Thus, this system is very similar in intention to VANET-based systems like SOTIS [32] or TrafficView [24].

This basic system is schematically outlined in Figure 2. The cars participate in some distributed hash table. Roads are divided into road segments, each with a unique ID. These IDs are used as keys in the DHT. Each node is responsible for a certain part of the ID space. It stores the information about the current traffic status on the respective road segments. When a car passes a certain road section, it shares its observations regarding the local traffic situation by sending a report to the node that hosts this road segment. To acquire information about the roads that lie on a possible route, a car may then query the nodes responsible for the respective roads.

Obviously, if built that naively, the system will not work satisfactorily. There are a number of issues involved that require careful consideration and lead to challenging research questions. These algorithmic issues appear in a very similar way in many peer-to-peer TIS applications. In the following, we discuss these issues and sketch some potential solutions.

#### 3.1 System load and scalability

In existing peer-to-peer systems entries in the DHT are usually added or updated in order to announce the presence of a file on a certain peer. Since this information is rather static, the update rate of the DHT is much lower than in a traffic information system where each peer constantly adds new or updated information to the system. Just in order to illustrate the load that such a system would have to endure it might be helpful to compare it with a well known Internet service. Google USA processed about 3.1 billion queries in November 2006 [25]. A comparable number of queries would be generated by 72,000 cars continuously sending updates or requests once per minute—that is only 0.03 % of the vehicles registered in the U.S. [3]. Even these very rough figures show that the load of




Figure 2: Traffic information sharing system.

queries and updates on a peer-to-peer TIS would be challenging. In particular it will be necessary that updates and queries are routed in an efficient manner through the DHT.

One possible approach to solve this problem is to exploit the specific structure of the data. The road segments are not as independent as, for example, different files in a peer-to-peer file sharing system but are topologically connected. Thus, the lookups and updates that are performed in the system will follow a pattern. If a car's navigation system is interested in the traffic situation on one road, it is likely that it is interested in adjacent roads as well. Similarly, if a car reports measurements on one road segment, it is very likely that it has previously passed some other, close-by road segment, and reported measurements on it. Simply hashing the road segments' IDs independently to nodes in the distributed hash table would require a full lookup for each road segment, which is, under the given circumstances, surely not optimal. Ideally, the underlying peer-to-peer system should understand and support the interdependencies to reduce the update and lookup overhead. This could be achieved by maintaining additional pointers in the peer-to-peer system, providing "shortcuts" to nodes responsible for adjacent road segments. Following such pointers, a node can look up close-by road segments with constant effort, when starting from some already known responsible peer. Maintaining the consistency of these pointers, however, might turn out to be an issue.

The problem could also be tackled more fundamentally by choosing or designing the distributed hash table in a way that respects the locality of the information and the dependencies between close-by places. This might be challenging since most existing peer-to-peer systems use a one-dimensional key space whereas the road network has essentially a two-dimensional topology. However, one particular distributed hash table already supports the notion of multiple-dimensions: CAN [27] employs multi-dimensional "universes", where nodes and keys are mapped into an  $n$ -dimensional ID space. Therefore, if close-by road segments are mapped to close-by positions in this space, CAN might provide fast local lookups in a very "natural" way: the respective pointers are already part of the design of this DHT. Thus, CAN could be a good starting point for developing DHTs that are well suited for traffic information systems.

#### 3.2 Fairness and reliability

Another main problem is the fairness of the workload distribution. A node that happens to be responsible for, e. g., a highway intersection will not only receive reports at a very high frequency,

it will also be constantly queried by a large number of navigation systems. Closely related is the problem of reliability as it is surely not optimal if all the cooperatively collected data on a road segment is lost if the responsible node gets disconnected from the network.

To deal with these issues, some form of load balancing and redundancy is necessary [6]. It is conceivable that not one single node, but groups of nodes are responsible for each road segment. Report messages could then be distributed to the members of this group, e.g., in a randomized or round-robin fashion. The size of the groups should be adjusted dynamically, based on the network traffic load associated with the road segment. The group members might periodically exchange aggregated data to share their knowledge. Such a scheme avoids losing valuable information if a node fails or is temporarily not reachable, and it also provides a simple form of load balancing: to learn about the current situation on a road segment, it is sufficient to ask one of the group members.

### 3.3 Bootstrapping

In order to join the network, a new user needs to know at least one peer which is already in. In existing DHTs this is accomplished by providing a list of well-known always-on nodes. This is not as easily possible in a system where all the nodes are cars, and therefore are most likely not “always-on”. But still an approach where the IP addresses of some currently online peers are stored at a well-known location seems viable. The system might of course also include some non-car peers on the Internet.

Ideally, the joining procedure takes into account the problem of workload balancing introduced above. New nodes could be integrated in the ID space where their help is most needed. Some work towards a network supporting such functionality is presented in [1].

### 3.4 In-network aggregation

In most cases, cars will travel multiple consecutive segments of a road. Typically, they are mainly interested in the traffic situation along a whole possible route. It might thus be a good idea to store an aggregated description of longer parts of roads, or of often-used routes spanning multiple roads. This reduces the effort for all cars interested in these larger building blocks: they would otherwise all have to query each segment separately, finally all ending up with doing the same processing. Such a preprocessed aggregate must of course stay up-to-date. The simplest solution would be to let the nodes responsible for some part of the aggregate update it from time to time. There is obviously a tradeoff involved between traffic and computational effort for proactive preprocessing of aggregates and savings in query and response traffic.

A different approach could be a form of distributed caching. If both the query and the response are routed along the DHT, each node can remember routed answers, generate aggregates on its own, and use them to “shortcut” answers to reoccurring queries. The effectiveness of such an approach will mostly depend on the structure of the distributed hash table, in particular on how likely it is that similar queries from different sources go through the same intermediate nodes. Again, there is a tradeoff, now between local storage space for caching and reduced query traffic.

### 3.5 Trustworthiness and privacy

So far, we have naively assumed that availability and up-to-date-ness are the only important properties of the data in the proposed system. In a real-world implementation, the data’s trustworthiness should, however, also be taken into account. This problem has already been recognized and studied in the VANET context [7, 13, 17]. In peer-to-peer networks, there are at least two possible points where trustworthiness might be injured. A node might (on purpose

or not) insert incorrect data into the DHT. With sufficient alternative measurements present in the system and provided resistance against Sybil attacks [10], this seems to be a marginal problem—abnormal data will be easy to notice. Misbehavior of the node hosting the data is harder to cope with. The above mentioned redundancy of stored data is probably not sufficient. To keep the data consistent, the nodes responsible for the same road segment will exchange information. Thus, the malicious node might be able to “poison” all descriptions stored in the DHT. Solutions proposed for VANETs to counter this problem usually take advantage of public key cryptography with central key certification. Since in our proposed approach Internet access is used, such a system may even be simpler to deploy than in VANETs.

Note that a central key signing instance is still very different from a centralized traffic information system, especially also with regard to user privacy. A mere key signing authority is comparable to a number plate issuing authority, which does not know about the movements of the registered vehicles. Sending periodical updates to a central instance allows the tracking of all movements and thus requires a whole different level of trust. An interesting feature of our proposed decentralized approach is thus that spreading the data among the users will in fact improve the protection of each single user’s privacy: there is no single instance where all data is available or a history of events could be collected.

### 3.6 Connectivity

Finally, a peer-to-peer system of highly mobile devices over the mobile Internet has to face external adversities. In particular, the underlying distributed data structures must be able to deal well with intermittent connectivity. There are two somewhat different cases. There may be abrupt loss of connectivity which is unforeseeable for the device and its neighbors in the overlay. This could, e.g., happen due to varying quality and level of deployment of the access network used. But there is also the case where the disconnection happens voluntarily, or at least in a predictable way. The system should be robust with respect to both.

In the first case, this will most likely be a question of the level of required redundancy in the overlay. The information should always be available on currently connected devices, such that queries can be answered within reasonable time. For the latter case, it may be possible to delay the disconnection slightly, in order to trigger restructuring of the respective part of the DHT and to transfer valuable data to other nodes.

## 4. PUBLISH/SUBSCRIBE TIS

Assuming that the so far discussed issues are solved a peer-to-peer system would be able to collect and provide traffic information in a robust and scalable fashion. A participating car’s navigation system is thus able to request information on the current traffic situation along a possible route. What is not easily possible, however, is keeping track of changes of the underlying data. To check whether the currently chosen route is still the best option, a car would have to query the peer-to-peer system periodically. Even with continuous connectivity to the system this seems inappropriate. Although well-designed querying schemes that avoid many unnecessary lookups seem feasible, they do not fundamentally solve the problem.

So far, the peer-to-peer system has been used like a distributed database. Measurement data is inserted, stored in the distributed hash table, and subsequently queried by other nodes. This is undoubtedly a viable basis and an important building block for many TIS applications. The additional requirement of keeping track of significant changes, however, motivates substantial extensions.






Figure 3: Publish/subscribe in a traffic information system.

Pull-based information retrieval initiated by the interested car's navigation system can be complemented by a push-based elements. This is called “continuous queries” in the database context. Here, it can be understood as a publish/subscribe architecture, where the system knows which clients are interested in some incoming information and actively distributes it to them.

The main idea of the publish/subscribe paradigm is to provide a message distribution service where subscribers register their interest for certain types of messages without knowing the sources of these messages in advance. The sources on the other hand will only push their messages to the service, which manages the subscription database and is then responsible to deliver the content. A more detailed introduction to the publish/subscribe concept is provided in [12].

A matter of particular interest is that publish/subscribe services can be implemented in a completely distributed way, using peer-to-peer networks. Such systems have already been designed, examples are Mirinea [8] or Scribe [28]. Both subscriptions and dissemination are handled in a decentralized manner. This is accomplished by using interest definitions as keys, peer IDs as values in a DHT, and building a multicast-like message distribution tree upon generic overlay routing. In a peer-to-peer TIS, this might be integrated with the DHT used for storing the information.

Based upon the publish/subscribe paradigm, such an approach could serve as a basis for a more sophisticated peer-to-peer traffic information system. The general idea is visualized in Figure 3. The responsibilities of such a system boil down to three main tasks: 1) detecting changes that should trigger a notification, 2) managing subscriptions, and 3) distributing the information.

#### 4.1 Detection of changes

Obviously, only a significant change of the traffic situation is relevant for the drivers. Recognizing that such a change has occurred might belong to the duties of the peers responsible for the respective road segment. Since this is where the updates arrive, these nodes will have the necessary knowledge—the history of observations.

After detecting an event important for cars traveling the road, like a rapid change of the reported average speed in some region, the second step would be to bring this information to the corresponding peers. This includes finding out which cars are interested in receiving a notification about the event on the one hand, and distributing the message to these cars on the other.

#### 4.2 Managing the subscriptions

Generally speaking, a car will be interested in updates concerning its route, and possibly also in significant changes on alternative routes. More specifically, *deteriorations* on the chosen route are interesting, as well as significant *improvements* on alternative routes that have not been chosen due to a previously determined adverse traffic situation.

To be able to identify interested cars, their routes need to be stored and indexed in the peer-to-peer system. This essentially corresponds to the subscription database of a classic publish/subscribe system. Support for more specific subscriptions, defining more precisely which updates the car is interested in, are not necessary for a working basic system, but might provide additional functionality. The whole range from a simple, purely road ID based subscription management to a sophisticated system like in fully fledged publisher/subscriber systems is conceivable here—similar improvements as for the basic system in the previous section come to mind.

#### 4.3 Dissemination service

Having the data and the list of subscribed users available, the information finally has to be “pushed” towards the group of subscribers. Again, a problem of workload balancing arises. It is certainly not appropriate if the originator of a notification, e.g., the responsible node for a road segment, alone has to bear the burden of sending the notification to each subscriber. Existing peer-to-peer based publish/subscribe systems use a distribution tree to deal with that, essentially they build a multicast overlay based on the peer-to-peer system. The publisher sends the message to a number of nodes, and those will forward it further.

But distribution along a tree structure like in the above mentioned solutions has a negative property, especially in dynamic networks with frequent restructuring and a relatively high node failure rate: if one node fails, then a whole subtree will not get the information. It seems that redundancy in the dissemination structure might contribute positively to robustness and reliability. Such redundancy is conceivable in the whole range from partly overlapping subtrees to multiple parallel, fully separated distribution trees. The cost of the reliability improvement, namely the multiplication of some messages, does not appear to be too high—in the end, this tradeoff might also be influenced by the importance of the information for the respective receivers.

Another aspect that might be worth to consider in this context is the acceptable delay. In many cases, it seems to be smaller than the information age tolerable for route planning: cars need to be informed early enough to be able to react, i.e., to change to an alternative route. One might think about a system that assigns the receivers of a notification a higher priority if they need to react sooner, e.g., because they are closer to the region the notification refers to—or closer to a junction of two alternative routes.

#### 4.4 Privacy issues

The proposed system requires its users to upload their route planning and IDs to a more or less publicly accessible storage space. The system must also maintain the relation between the specific participant and the intended route. From a privacy perspective, this is much more problematic than sending reports on single traversed road segments to varying peers. Thus, special care should be taken to protect the privacy of users in such a system. Not only is it advisable to store the data no longer than necessary, but also other means of privacy protection might be considered. A respective measure might be to distribute sensitive data over many nodes, such that none of them has all the information on a particular participant

available. Obviously, there is a tradeoff between performance of the system and protection of the user's privacy, which might be of more general interest for distributed publish/subscribe systems, and probably deserves careful investigation.

## 5. FURTHER ADVANCED SYSTEMS

In the previous sections we have established a “baseline” of peer-to-peer based traffic information systems. Now, we may move the research horizon further and look at different, more speculative applications of peer-to-peer systems for car-to-car information exchange.

### 5.1 Processing specific requests

In all approaches described so far, the cars actively push certain information (e. g., observed traffic density, parking lot occupancy, or congestion warnings) into the distributed database. This is reasonable as long as the information is of general interest. However, each car is able to collect a large number of parameters via its sensors. It is very likely that only a fraction of these meets enough interest to be actively pushed to the system. On the other hand, this does not mean that these parameters are of no interest at all: there may be some peers for which some specific information *is* relevant. Therefore, one might look into a system that allows to query relevant sensor readings on demand.

An example may be a driver of a convertible who has three possible routes and wants to select the most sunny one. Data about the current sunshine intensity is most likely not of general interest, thus it might be better to let the system play the role of a broker that just refers requests to the cars that have the necessary information, or are able to obtain it easily. To provide this functionality, techniques developed for sensor networks and for content addressable networks [14] can be employed and adapted.

In some sense, such a service is similar to an “inverse location service”: while location services for geographic routing in mobile ad hoc networks [23] are used to find the position of a node with a given ID, this service is intended to find the ID of a node at or close to a given position.

### 5.2 Traffic coordination

One interesting, advanced application envisioned for car-to-car communication is traffic coordination [20, 31]. E. g., during afternoon rush hour, a lot of people will drive from the city to their homes. It seems that in this case providing users only with the current traffic situation is not optimal. Once a majority of traffic participants use advanced communication techniques for route planning and information sharing, a service coordinating the routes to optimize the traffic flow suggests itself. Putting this in the context of the subscription system discussed above, the system might be extended to keep track of the number of vehicles *intending* to use a specific route in the future.

In such a traffic coordination system, cars will likely no longer query and analyze data locally, but rather cooperate for a decentralized analysis. In some sense, they share computational power to optimize the global traffic flow jointly.

### 5.3 Distributed computation

Distributed calculations and cooperative data analysis are also possible in a different context: even the calculation of the optimal route to be used by some specific vehicle can possibly be turned into a problem that can be solved efficiently in a distributed way.

In order to gain the full knowledge about the road conditions along possible route alternatives, one has to query sequences of nodes, each providing information on part of the intended route.

So, instead of first collecting and then evaluating all data in the car that performs route planning, the query for the optimal route might travel through the network, being partially processed in the nodes where information of interest is locally available, until it finally returns carrying the answer. Essentially, the query itself could travel along the possible routes in the virtual DHT space and virtually “visit” the places of interest. This can be understood as a specific case of a “mobile agent” system [26]. Again, techniques as they are discussed for in-network data processing in sensor networks might also provide a starting point for research in this direction.

## 6. CONCLUSION

In this paper, we have presented a new paradigm to implement traffic information systems, using an infrastructure-based peer-to-peer network made up of vehicles. This approach has several advantages over traditional, VANET-based systems. Since the approaches are of course non-exclusive, it might also be possible for both VANET and peer-to-peer based systems to coexist and complement each other.

Maybe the most beneficial feature of the peer-to-peer approach is the low penetration rate that is required: starting from two cars in the system communication is possible, and the only bottleneck is the number of data contributors. Thus, the system offers an advantage already for the first buyers, while the usefulness rapidly increases with the number of equipped cars.

## 7. REFERENCES

- [1] K. Aberer, P. Cudré-Mauroux, A. Datta, Z. Despotovic, M. Hauswirth, M. Puceva, and R. Schmidt. P-Grid: A self-organizing structured P2P system. *ACM SIGMOD Record*, 32(3):29–33, Sept. 2003.
- [2] S. Androutsellis-Theotokis and D. Spinellis. A survey of peer-to-peer content distribution technologies. *ACM Computing Surveys*, 36(4):335–371, 2004.
- [3] Bureau of Transportation Statistics. Number of U.S. Aircraft, Vehicles, Vessels, and Other Conveyances. [http://www.bts.gov/publications/national\\_transportation\\_statistics/html/table\\_01\\_11.html](http://www.bts.gov/publications/national_transportation_statistics/html/table_01_11.html).
- [4] Car 2 Car Communication Consortium. <http://www.car-2-car.org/>.
- [5] M. Caliskan, D. Graupner, and M. Mauve. Decentralized discovery of free parking places. In *VANET '06: Proceedings of the 3rd ACM International Workshop on Vehicular Ad Hoc Networks*, pages 30–39, Sept. 2006.
- [6] G. Chiola. An empirical study of data redundancy for high availability in large overlay networks. In *HotP2P '05: IEEE 2nd International Workshop on Hot Topics in Mobile Peer-to-Peer Systems*, pages 43–50, July 2005.
- [7] J. Y. Choi, M. Jakobsson, and S. Wetzel. Balancing auditability and privacy in vehicular networks. In *Q2SWinet '05: Proceedings of the 1st ACM international workshop on Quality of service & security in wireless and mobile networks*, pages 79–87, Oct. 2005.
- [8] Y. Choi and D. Park. Mirinae: A peer-to-peer overlay network for large-scale content-based publish/subscribe systems. In *NOSSDAV '05: Proceedings of the International Workshop on Network and Operating Systems Support for Digital Audio and Video*, pages 105–110, June 2005.
- [9] F. Doetzer, F. Kohlmayer, T. Kosch, and M. Strassberger. Secure communication for intersection assistance. In

- WIT '05: *Proceedings of the 2nd International Workshop on Intelligent Transportation*, May 2005.
- [10] J. R. Douceur. The sybil attack. In *IPTPS '01: Proceedings of the 1st International Workshop on Peer-to-Peer Systems*, pages 251–260, Feb. 2002.
  - [11] W. Enkelmann. FleetNet – Applications for Inter-Vehicle Communication. In *IV '03: Proceedings of the IEEE Intelligent Vehicles Symposium*, pages 162–167, June 2003.
  - [12] P. T. Eugster, P. A. Felber, R. Guerraoui, and A.-M. Kermarrec. The many faces of publish/subscribe. *ACM Computing Surveys*, 35(2):114–131, June 2003.
  - [13] P. Golle, D. Greene, and J. Staddon. Detecting and correcting malicious data in VANETs. In *VANET '04: Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks*, pages 29–37, Oct. 2004.
  - [14] R. Govindan. Data-centric routing and storage in sensor networks. In C. S. Raghavendra, K. M. Sivalingam, and T. Znati, editors, *Wireless sensor networks*, pages 185–205. Kluwer Academic Publishers, Norwell, MA, USA, 2004.
  - [15] P. Gupta and P. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, Mar. 2000.
  - [16] T. Hossfeld, K. Tutschku, F.-U. A. de Meer, and J. Oberender. Simulative performance evaluation of a mobile peer-to-peer file-sharing system. *IEEE Next Generation Internet Networks*, pages 281–287, Apr. 2005.
  - [17] J.-P. Hubaux, S. Capkun, and J. Luo. The security and privacy of smart vehicles. *IEEE Security and Privacy*, 2(3):49–55, 2004.
  - [18] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden. Cartel: A distributed mobile sensor computing system. In *SenSys '06: Proceedings of the 4th International Conference on Embedded Networked Sensor Systems*, pages 125–138, Oct. 2006.
  - [19] IEEE P802.11, Task Group p. WAVE: Wireless access for the vehicular environment. [http://grouper.ieee.org/groups/802/11/Reports/tgp\\_update.htm](http://grouper.ieee.org/groups/802/11/Reports/tgp_update.htm).
  - [20] J. LeBrun, J. Anda, C.-N. Chuah, M. Zhang, and D. Ghosal. VGrid: Vehicular ad hoc networking and computing grid for intelligent traffic control. In *VTC '05-Spring: Proceedings of the 61st IEEE Vehicular Technology Conference*, pages 2905–2909, May 2005.
  - [21] C. Lochert, B. Scheuermann, M. Caliskan, and M. Mauve. The feasibility of information dissemination in vehicular ad-hoc networks. In *WONS '07: Proceedings of the 4th Annual Conference on Wireless On-demand Network Systems and Services*, pages 92–99, Jan. 2007.
  - [22] E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications Surveys and Tutorials*, 7(2):72–93, Mar. 2005.
  - [23] M. Mauve, J. Widmer, and H. Hartenstein. A survey on position-based routing in mobile ad-hoc networks. *IEEE Network Magazine*, 15(6):30–39, Nov. 2001.
  - [24] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Ifode. TrafficView: Traffic data dissemination using car-to-car communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 8(3):6–19, 2004.
  - [25] Nielsen//Netratings. U.S. search share rankings, press release. [http://www.nielsennetratings.com/pr/pr\\_061219.pdf](http://www.nielsennetratings.com/pr/pr_061219.pdf), Dec. 2006.
  - [26] V. A. Pham and A. Karmouch. Mobile software agents: An overview. *IEEE Communications Magazine*, 36:26–37, July 1998.
  - [27] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *SIGCOMM '01: Proceedings of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 161–172, Aug. 2001.
  - [28] A. I. T. Rowstron, A.-M. Kermarrec, M. Castro, and P. Druschel. SCRIBE: The design of a large-scale event notification infrastructure. In *NGC '01: Proceedings of the 3rd International Workshop on Networked Group Communication*, pages 30–43, Nov. 2001.
  - [29] S. Saroiu, P. K. Gummadi, and S. D. Gribble. A measurement study of peer-to-peer file sharing systems. In *MMCN '02: Proceedings of Multimedia Computing and Networking*, pages 156–170, Jan. 2002.
  - [30] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: a scalable peer-to-peer lookup protocol for Internet applications. *IEEE/ACM Transactions on Networking*, 11(1):17–32, 2003.
  - [31] U.S. Department of Transportation. ITS: Federal intelligent transportation systems. <http://www.itsoverview.its.dot.gov>.
  - [32] L. Wischhof, A. Ebner, H. Rohling, M. Lott, and R. Halfmann. SOTIS – a self-organizing traffic information system. In *VTC '03-Spring: Proceedings of the 57th IEEE Vehicular Technology Conference*, pages 2442–2446, Apr. 2003.
  - [33] Q. Xu, T. Mak, J. Ko, and R. Sengupta. Vehicle-to-vehicle safety messaging in DSRC. In *VANET '04: Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks*, pages 19–28, New York, NY, USA, 2004.
  - [34] J. Yin, T. ElBatt, G. Yeung, B. Ryu, S. Habermas, H. Krishnan, and T. Talty. Performance evaluation of safety applications over DSRC vehicular ad hoc networks. In *VANET '04: Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks*, pages 1–9, Oct. 2004.
  - [35] B. Y. Zhao, L. Huang, J. Stribling, S. C. Rhea, A. D. Joseph, and J. D. Kubiatowicz. Tapestry: A resilient global-scale overlay for service deployment. *IEEE Journal on Selected Areas in Communications*, 22(1):41–53, Jan. 2004.