4th Annual IEEE/IFIP Conference on Wireless On Demand Network Systems and Services (WONS), Obergurgl, Austria, January 2007

Vehicular Ad-Hoc Networks: From Vision to Reality and Back

Holger Füßler Sascha Schnaufer Matthias Transier Wolfgang Effelsberg
Department of Mathematics and Computer Science,
University of Mannheim, Germany
{fuessler, schnaufer, transier, effelsberg}@informatik.uni-mannheim.de

***Abstract*— VANETs, or Vehicular Ad-Hoc Networks, have recently gained scientific and commercial interest. In fact, they have drawn many people from neighboring fields like general Mobile Ad-Hoc Networks (MANETs), into their wake. In this invited paper, we discuss the history of Vehicular Ad-Hoc Networks from our perspective. In detail, we will show the early vision of creating a huge MANET that would facilitate cheap and ubiquitous communication on the ISM bands, and how this vision was reduced to cars sending emergency information in a geographically limited area. Also, we will describe how new challenges emerge from these new constraints, and then argue that VANETs are still an interesting research area.**

I. Introduction

Vehicular Ad-Hoc Networks (VANETs) are wireless networks between vehicles. Recently, they have attracted increasing scientific interest from the wireless networking community. This is, on one hand, due to the car manufacturing industry's determination to roll out vehicle-to-vehicle communication in the near future and, on the other, to the increasing disillusionment concerning the need for the vast number of protocols developed for general Mobile Ad-Hoc Networks (MANETs) in the past few years. In the case of VANETs, industry pressure has created a situation in which an overwhelming interest in solutions to problems leads to a preference for real-word research as opposed to fancy theory.

In this paper, we talk about our own part of the history of (Vehicular) Ad-Hoc Networks. Moreover, we will speak of the initial proposition of the "FleetNet—Internet on the Road" [1], [2] project, a German government and industry initiative envisioning roads as huge ad-hoc networks running mainly Internet applications. Over the years, industry interest and research reality has limited this vision to running specialized (non-Internet) applications with a very limited hop scope.

The remainder of this paper is organized as follows: The next section describes our view on VANETs emerging from general MANETs. Section III shows how the initial work was confronted by reality, in terms of both real-world protocol experience and changing industry needs. Section IV discusses how this change resulted in a crisis concerning network architecture. Finally, Section V briefly addresses evolving challenges.

II. From MANETs to VANETs

When we entered VANET research with the FleetNet project in mid-2001, ad-hoc research was largely dominated by efforts to standardize MANET protocols in the same-named IETF working group. Consequently, these protocols were tailored to transporting IP unicast datagrams, enabling the variety of IP applications to be run transparently over these networks. Thus, early MANET research focused on the network layer, since (a) wireless network hardware including layers 1 and 2 could already be bought off the shelf, and (b) everything above IP was already there. So the ultimate challenge seemed to lie in the problem of how to reach nodes not directly within radio range by employing neighbors as forwarders.

MANET research itself tried to treat the networks as generally as possible, following the idea that worst-case engineered protocols also fit special-case scenarios. E.g., the widely used Random-Waypoint Mobility model [3] modeled node movement by means of randomized sequences of linear movement, combined with randomized periods of standstill. Being able to parametrize the random distributions, this created a powerful tool to create scenarios with various degrees of mobility. However, these scenarios do not exist in reality, so the basic assumption was that a protocol able to cope with worst-case (because unpredictable) node movements of a random movement model would also be able to cope with the more correlated movements occurring in reality.

At this point, there seemed to be an obvious chance to do some original research. If we regarded street-bound vehicles as a MANET, we could at least study well-known methods, or, even more, create modifications or even a novel proposal for these scenarios. And, as a matter of fact, we partially succeeded by exploring and extending position-based routing methods for vehicular scenarios. Actually, we were so successful that when observing the route-finding capability of position-based methods on highways, the theoretically achievable number of hops seemed to be almost unlimited[1]. However, informal probing of the IETF MANET people quickly showed that getting a protocol proposal on the standards track would create a tremendous effort and would encounter significant opposition, since they were fiercely trying to reduce the number of candidate protocols. Nevertheless, since the FleetNet agenda called for a proof-of-concept demonstration platform, we felt fit to implement position-based routing for Linux with the ultimate goal to perform some real-world measurements on the road.

III. Hitting Reality

A. Implementation

In the process of implementing greedy position-based unicast forwarding, we quickly realized how assumptions made in protocol simulations can easily create problems that are hard to overcome in reality:

1) In protocol simulations, the processing time for the routing decision is assumed to be negligible, i.e., zero. In reality, one has to be careful about algorithmic complexity, and even calculations consuming a constant time (like trigonometric functions) might be too expensive if they have to be done per packet.
2) Although a discrete event simulator [5] easily handles concurrency while still being able to run single-threaded itself, the handling of parallel, but dependent processes on a real system increases the complexity of an implementation considerably.
3) The simulator allows for arbitrary protocol modifications on any layer, whereas on a real system, one usually cannot get any lower than the interface offered to the protocol parts that run inside a network interface card. E.g., in a simulator, it is quite easy to notify the network layer if a link-layer transmission has failed, while on a real system, it depends on whether or not the driver offers this feature.
4) The probably most damaging differences between simulation and reality are the inadequate modeling of the radio channel and the electrical circuitry directly dealing with it. In this context, not only the probabilistic nature of the channel but also the over-simplification of MAC protocol elements like carrier sensing or the real behavior of a network interface when a "packet collision" occurs [6] creates an enormous difference between the model and reality, often not only on a quantitative scale.

[1]Of course, this limited view ignores the scaling limit discovered by Gupta and Kumar in [4]. However, it is the usual approach to measuring routing performance.

Despite the problems listed above, we have managed to build a communication platform capable of unicast packet forwarding in a VANET [7]. Still, protocol performance is significantly lower in reality than in simulation, obviously due to the best-case assumptions listed above that do not hold true in reality. Consequently, the simulations have been useful and important, keeping in mind that reality is worse.

B. Communication Requirements

In addition to the technical and performance issues listed above, working with the real system taught us that while we were happy evaluating our unicast protocols, the application developers required almost everything but this form of packet forwarding. In contrast, the FleetNet demonstration applications were mostly based on geocast or scoped broadcast forwarding, with only one application called "virtual paper chase" using unicast to let specific other cars know about the junctions where they had turned. Then again, our primary goal of enabling standard TCP/IP communication in our VANET lacked (a) the usefulness of the majority of Internet applications when there is no connection to the Internet and (b) the availability of a TCP-like transport protocol that is capable of dealing with the difficult delay and loss situation.

Alongside this technical reasoning stands the question of the "added value" of TCP/IP applications on the road. Especially in first-world countries, where infrastructure-based Internet access is the killer technology for such a network, ad-hoc networking can hardly grow better—only cheaper.

As a direct consequence, a "FleetNet – Internet on the Road" scenario is not likely to happen. On the contrary, ad-hoc technology will probably deal with specialized car-to-car communication applications, while Internet access will be handled by infrastructure or will use the car-to-car communication system only in (static) single-hop mode.

However, the requirements for the car-to-car communication part still pose challenges to the ad-hoc networking community because the networking system has to make sure that in a low-density situation, the available bandwidth is utilized so as to maximize the spread of information while in a high-density situation the network does not choke on packet load. The latter can be especially critical regarding time- or reliability-critical delivery modes.

IV. Architecture Crisis

The previous section described the changes that have occurred to VANET research over the years. As a consequence, VANET research has conformed to the shifting demands. Long-distance unicast networking has drifted out of focus, and people tend towards applications that are more likely to be integrated into a future VANET system. Fortunately, we have managed to rescue some of our results by showing their similarity in concept. E.g., on a highway, efficient unicast forwarding is algorithmically very close to broadcast or multicast.

An entirely new dimension of problems arises from the revocation of the end-to-end packet payload integrity. That means that for multi-hop VANET communication, the communication endpoint is practically on every system. While in the Internet the payload is transferred without alteration until finally delivered, every node in a VANET may contribute to the payload content, crucially reducing information redundancy between packets. E.g., a car might sense an icy road, and the communication system should be used to distribute this information in order to allow preventive action. A communication system would then have the mission to distribute this message reliably within a certain region. However, environmental sensor information is usually highly correlated. Thus, it is very likely that many cars will sense the same danger, thereby resulting in many redundant messages. These messages may likely choke the network if the information redundancy is not properly exploited. An alternative to letting the communication system distribute the unaltered information is of course to only tell your radio neighbors about the danger, then let them aggregate the information and let them tell their neighbors etc. Through this modus operandi, information redundancy can be perfectly reduced since the respective applications deal with information they *understand* as opposed to information they only forward. Obviously, this is also reducing communication system complexity down to single-hop operation, which some people consider very favorable for a quick system roll out.

However, this alternative burdens every application protocol developer with the task to build protocols that efficiently use available bandwidth and do not choke the network. In our experience, application protocol designers do not care enough about these issues, while network and MAC protocols mostly work on unaltered payloads, creating an architectural dilemma because the layer separation is not as easy as in classical network stacks. Also, the purely single-hop system would require every node to run every application to enable communication beyond the line of sight.

Figure 1 depicts the two different approaches as described above. On the left side, a classic TCP/IP-like stack is shown, where the application uses a transport layer to send payloads addressed to, e.g., a geographic region. The network layer then distributes the packet to all nodes in that region. On every node inside that region, the packet is also handed upwards to the respective application.

The right figure shows the single-hop solution, where the only usable network layer is the link layer that is used to send broadcasts to its neighbors. On every node, the application handles received broadcasts directly, deciding for itself whether or not the information contained in the packet should trigger a new broadcast[2].

V. Current Challenges

The above section described the architectural development process of VANETs in our scope. In this section, we will deduce some challenges which we think will play a major role in the ongoing process of VANET research.

A. VANET Protocols

VANET protocols that are able to make it to the product will stage the need to work under very different conditions. For the first couple of years, a car equipped with a VANET system will find hardly any other cars with which to productively exchange messages. Thus, the first task of the protocols will be to operate under these conditions. They will—in the beginning—not care very much about channel usage to maximize utility. I.e., in the beginning, the probabilistic channel will be used frequently to increase the utility range of VANET messages. However, as system penetration increases, the scarcity of the radio channel as a resource will increase. Paying

[2]For a deeper discussion of this topic and our current approaches to tackling it, see [8].

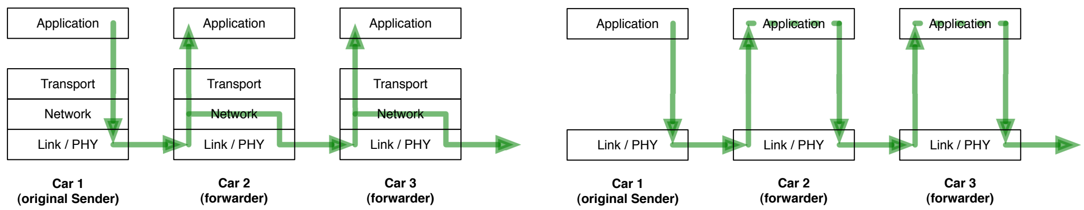

(a) with Network Layer

(b) without Network Layer

Fig. 1. VANET Multi-hop Network Stacks

respect to the increasing likelihood of packet collisions, this will imply the absolute necessity of minimizing its usage acknowledging the increasing likelihood of packet collisions. Consequently, a significant challenge lies in building protocols that work in both cases, and a great danger lies in building protocols that are hastily tailored to cope with the low-density situation. The high-density situation, however, creates the greater challenge of seeing the multi-hop effect of single-hop broadcasts. This means that whenever information triggers broadcasts, the subsequent message exchange is, in fact, part of a multi-hop protocol, which has to be evaluated on a non-local scope by people with knowledge in ad-hoc networking. It is—in our view—simply not enough to provide single-hop broadcast to application developers and then let them worry about the rest.

B. Architecture and Cooperation

Following Section IV's arguments, the architecture of a future VANET system is still not clear, at least for the projects we were involved in. While many people consider this an academic discussion, it has some impact on how protocol development can be separated and cooperation can be stimulated. Obviously, the Internet's end-to-end paradigm has to be reconsidered, since there is technically no backbone vs. end system structure, but every node is both end system *and* router. For the cooperation part, the know-how lies in the hands of different groups: On the one hand, there are network researchers holding knowledge about multi-hop protocols, retransmission timing, broadcast redundancy etc., and on the other, there are people with an understanding of traffic flow, time-criticalness of information distribution and so on. In our opinion, since both groups directly influence any resulting protocol's "radio profile", a stronger interaction is necessary to avoid a system with protocols that either will not really work in the beginning or choke the channel later. Especially the extreme high-density situation in a congested highway under the assumption that every car is running a VANET system will create the ultimate protocol challenge. Any available protocol design trick will have to be used to tackle these problems: from using infrastructure to control the channel or to coordinate information gathering, over the usage of classical algorithmic methods to increase scalability like hierarchization, up to methods involving the electrical properties of the signals like power control or different physical codings to stabilize transmissions. Moreover, the problems will have to be tackled quickly since the car manufacturing industry is eager to roll out a car-to-car communication system, and the consequences in rolling out a closed-box system have a long reach.

REFERENCES

[1] "The FleetNet Project Homepage," http://www.et2.tu-harburg.de/fleetnet/, 2006, [Online: Accessed at 2006-04-07].

[2] W. Franz, H. Hartenstein, and M. Mauve, Eds., *Inter-Vehicle-Communications Based on Ad Hoc Networking Principles—The FleetNet Project*. Karlsruhe, Germany: Universitätsverlag Karlsruhe, November 2005.

[3] D. B. Johnson and D. A. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks," in *Mobile Computing*, T. Imielinski and H. Korth, Eds. Kluwer Academic Publishers, 1996, vol. 353, pp. 153–181.

[4] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388–404, March 2000.

[5] A. M. Law and W. D. Kelton, *Simulation Modeling and Analysis*, 3rd ed. McGraw-Hill, 2000.

[6] A. Kochut, A. Vasan, A. U. Shankar, and A. K. Agrawala, "Sniffing Out the Correct Physical Layer Capture Model in 802.11b," in *Proceedings of the 12th IEEE International Conference on Networking Protocols (ICNP '04)*, Berlin, Germany, October 2004, pp. 252–261.

[7] H. Hartenstein, H. Füßler, M. Mauve, and W. Franz, "Simulation Results and Proof-of-Concept Implementation of the FleetNet Position-Based Router," in *Proceedings of the IFIP-TC6 8th International Conference on Personal Wireless Communications (PWC '03)*, Venice, Italy, September 2003, pp. 192–197.

[8] S. Schnaufer, H. Füßler, M. Transier, and W. Effelsberg, "Vehicular Ad-Hoc Networks: Single-Hop Broadcast is not enough," in *Proceedings of 3rd International Workshop on Intelligent Transportation*, Hamburg, Germany, March 2006, pp. 49–54.