On the implementation, deployment and evaluation of a networking protocol for VANETs: the VARON case☆

M. Isabel Sanchez[a,b], Marco Gramaglia[c], Carlos J. Bernardos[a,], Antonio de la Oliva[a], Maria Calderon[a]

[a]*Departamento de Ingeniería Telemática, Universidad Carlos III of Madrid*
[b]*Institute IMDEA Networks*
[c]*Istituto Superiore Mario Boella*

Abstract

Research on vehicular communications has been quite extensive over the past few decades. Most of the initial studies were theoretical and research has just recently moved to more experimental works. Conducting real field operational tests is extremely challenging due to the number of vehicles required, the lack of control over the environment and the cost of the necessary equipment and personnel. However, simulation tools may not reflect properly the highly dynamic and complex characteristics of the vehicular scenario. This article explores why practical research in the field of vehicular communications is so demanding, by reporting on our experience in prototyping and experimentally evaluating VARON. Published in 2008, VARON is a multi-hop wireless vehicular communication protocol which was already validated via extensive simulations. In this work, we have fully implemented it, first on a lab-based environment, and then in a real-life testbed. This long and exhausting process has shown that some common assumptions do not necessarily hold when evaluated under real situations, as well as taught us valuable lessons on how to design and conduct experiments with real vehicles. We believe that the experience and lessons learned during this process

☆This work has been funded by the European Community through the CROWD project (FP7-ICT-318115). The work of A. de la Oliva has been supported by the MICINN research grant TIN2010-20136-C03. At the time when this work was developed Marco Gramaglia was affiliated to Institute IMDEA Networks and University Carlos III of Madrid.

Email addresses: mariaisabel.sanchez@imdea.org (M. Isabel Sanchez), gramaglia@ismb.it (Marco Gramaglia), cjbc@it.uc3m.es (Carlos J. Bernardos), aoliva@it.uc3m.es (Antonio de la Oliva), maria@it.uc3m.es (Maria Calderon)

do not only apply to VARON, but also to other multi-hop wireless vehicular communication solutions, and that therefore these lessons are helpful for other researchers willing to validate their protocols in a real scenario.

Keywords: VANET, network mobility, route optimization, 3G, WiFi, experimentation, lessons learned, prototyping.

---

1. Introduction

Vehicular communications are no longer just a research topic. Governments, car manufacturers and telecommunication players have been working towards the definition of a communication architecture to improve road safety and efficiency that benefits from vehicles having communication capabilities. Among the candidate architectures that could provide connectivity in a vehicular environment, vehicular ad-hoc networks (VANETs) are probably the most popular ones due to their decentralized nature, which supports unmanaged operation without infrastructure involvement.

Over the last decade, a considerable amount of effort has been devoted to the analysis of technical solutions that could be applied to the vehicular scenario. This effort has mainly targeted road safety and traffic efficiency services, although Internet alike applications have also been recently considered, given their importance in terms of users' demand. As one very simple example of the impact of vehicular communications in research, a search for the keywords *vehicular communications* in *Google Scholar* returns more than 275k results. However, most of this existing research has not been experimentally validated because of its high cost and the complexity required to deploy and maintain a vehicular testbed.

In 2008 we published a VANET-based mechanism, called VARON [1], which consists of a route optimization for vehicular ad-hoc networks. The key use cases targeted by VARON involved Internet alike applications of interest for automobile users (e.g., personal communications or entertainment services) or vehicular specific services, such as traffic information or car diagnosis activities. VARON was validated via extensive simulations and the results proved not only the feasibility of the solution, but also showed that VARON was able to provide quite interesting performance gains. Many other research works show the same kind of positive feedback in a simulation-only environment, without exploring the issues that potentially arise in a real world experimental scenario. We believe in system research based on real life implementations, and this drove us to go all the way down to prototyping and testing our solution. We soon found out that many of the problems we faced during the process and the solutions we adopted can also give very helpful insights for other researchers in the area, so we do not focus in this article just on the performance of our particular solution. Therefore, this article reports on the whole process of taking a conceptual solution, implementing it (first in a closed lab environment), and then deploying it on a real-life testbed.

The main contributions of this article are: *i)* the identification of the main procedures commonly found in VANET solutions, so we can perform the mapping against VARON procedures, and then relate our findings to them; *ii)* a review of the related work in terms of real-life implementation experience; *iii)* a detailed description of our prototype implementation of VARON in a laboratory controlled environment, and the results of a first set of conducted experiments; *iv)* a walk-through description of the deployment of VARON in a close-to-real-life scenario, involving real vehicles and performing experiments on the street; *v)* a report of the adjustments and improvements made to the original VARON design derived from the experience gained in the implementation and deployment process, relating them to the common VANET mechanisms previously identified; and *vi)* a summary of the lessons learned in the process, as well as those aspects that may go unnoticed but should not be overlooked in experimental vehicular communications research.

The rest of this article is organized as follows. Section 2 presents some additional background information, useful to understand the operation of VARON. Section 3 overviews related work on experimental research with real vehicular networks. Then we describe the process followed to bring VARON from a simulation environment to become operative on real cars. The initial implementation efforts and tests in a lab-controlled environment are detailed in Section 4. Section 5 reports how this early prototype evolves to the one used in real vehicles, paying special attention to the problems we had to face due to the unexpected, but real, conditions found when we hit the road. We also analyze the performance of VARON under real conditions. Section 6 summarizes the lessons learned during the whole process of setting up the vehicular testbed and conducting the experiments. Finally, Section 7 concludes this work.

2. Background

The main goal of this article is to report on the challenging experience of implementing and experimenting with a vehicular communication protocol. In order to do so we have selected a particular solution we are familiar with, VARON, published in [1]. This section provides some background information on vehicular communications in general, in order for the reader to better understand our experimental work on VARON. This is required to follow the explanation on the problems faced during the prototyping phase, as well as the solutions we designed to tackle them, which will be presented throughout this article.

2.1. IP vehicular communications

The are two main types of communication in a vehicular scenario: between a vehicle and the infrastructure (V2I) or among vehicles (V2V). In the near future several devices within a vehicle, such as internal sensors, on-board computers or infotainment back-seat boards will likely benefit from having Internet connectivity, and we have to consider also external devices carried by passengers, such as laptops or smartphones. In this vehicle-to-Internet scenario, it is commonly assumed that a specialized node called *mobile router* provides external connectivity to these devices in the vehicle, which form a *mobile network*. A mobile router is in charge of providing connectivity to the intra-vehicle network, also managing transparently its mobility, that is, without any additional requirements to the attached devices. These mobile routers are also expected to have multiple access technologies available, so they can take advantage from this heterogeneity to forward the traffic through the most appropriate interface (e.g., 3G/LTE, WLAN).

Besides Internet access, there are several applications which involve vehicle-to-vehicle (V2V) communications, such as semi-automatic driving or gaming in platooned vehicles. In this case, the role of ad-hoc networks is even more clear, as they naturally enable V2V communications without involving any external infrastructure.

It is commonly assumed that the mobile router deployed in each vehicle has at least three network interfaces: one *ingress* interface to communicate with the nodes inside the vehicle that belong to the mobile network (e.g., WLAN, Ethernet), one or more *egress* interfaces to connect to the Internet (e.g., 3G/LTE), and an additional ad-hoc interface (e.g., WLAN) to communicate with neighboring cars and to set up multi-hop ad-hoc networks.

Another common assumption is that vehicles can always communicate with other vehicles through the Internet. In addition, they may communicate directly if a multi-hop route can be set up in the VANET. In our prototype, the access through the Internet is provided by the NEMO Basic Support protocol and the multi-hop ad-hoc route in the VANET is set up by VARON.

2.2. Network Mobility

The Network Mobility (NEMO) Basic Support protocol [2] was proposed by the IETF[1] to enable mobility of complete networks. This, for example, allows a set of devices deployed in a vehicle (e.g., a car, bus or train) to benefit from Internet connectivity. To do so, the NEMO Basic Support protocol extends the basic end-host mobility solution, Mobile IPv6 [3], to provide mobility management to complete networks. In this solution, a mobile network (known also as *network that moves* – NEMO[2]) is defined as a network whose point of attachment to the Internet varies with time. A specialized device, called the mobile router (MR), connects the NEMO to the Internet. It is assumed that the NEMO has a *home network*, connected to the Internet, where it resides when it is not moving. Since the NEMO is part of the home network, the mobile network nodes (MNNs) use IP addresses that belong to one or more address blocks assigned to the home network: the mobile network prefixes (MNPs). These addresses remain assigned to the NEMO even when it is away from home[3]. Thus, when the NEMO is attached to another network, packets addressed to the mobile network nodes will still be routed to the home network, and redirected by the home agent (HA) to the current location of the MR. When the NEMO is connected to a visited network, the MR acquires an address from the visited network, called the care-of address (CoA), where the routing architecture can deliver packets without any additional mechanism (see Figure 1).

2.3. Commonly found procedures in VANET solutions

Vehicular Ad-hoc Networks are a particular kind of mobile ad-hoc networks, characterized by a high mobility of nodes, short-lived links and an unstructured nature. VANETs are very dynamic and lack a pre-established topology, and they operate in a fully distributed way without control nor

[1] `http://www.ietf.org/`
[2] NEMO can mean NEtwork MObility or NEtwork that MOves according to the context.
[3] These addresses only have topological meaning when the NEMO is at home.

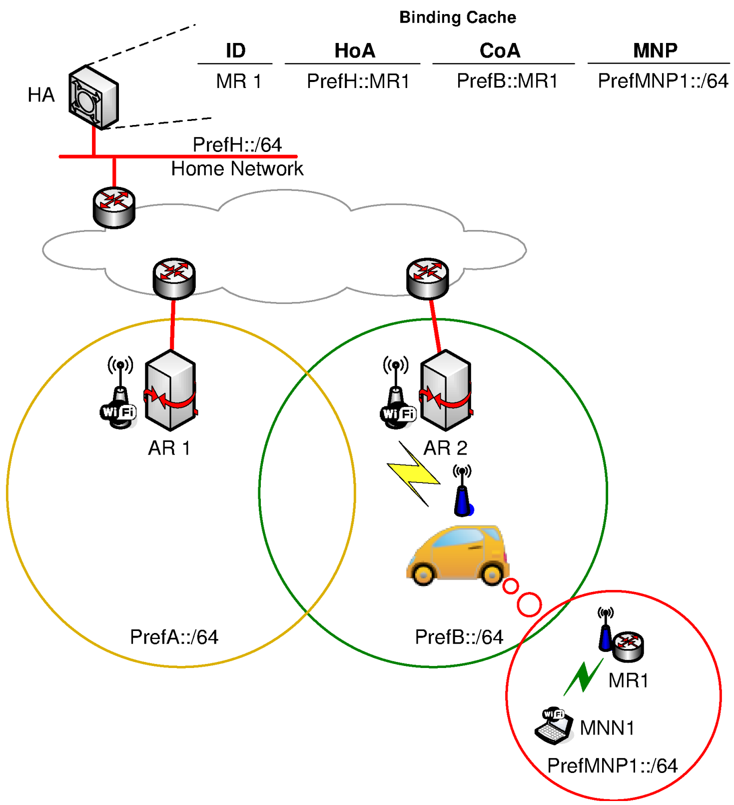

Figure 1: NEMO Basic Support protocol operation overview.

monitoring from a centralized entity. This makes more complex the design of solutions for vehicular protocols, as there are many different variables defining the nature of VANETs. They can be very dense, as in a urban area, or very sparse, as in a remote highway, and these conditions can vary in a few minutes. Moreover, the protocols designed for VANETs have to provide communications with reliability and deal with different sources of interference.

These particular characteristics of VANETs have fostered the publication of routing protocols tailored for them. These protocols adapt the mechanisms in ad-hoc networking to the vehicular environment and the mobility patterns of vehicles. In this section, we review the main mechanisms commonly present in VANET protocols and match them to their equivalent in VARON. Some of these mechanisms were not included in the original definition of VARON, but added as found necessary during our experimental process, as described in Section 5.2. The main procedures commonly found in VANET solutions are identified next:

- **Self-organization and discovery**. Vehicular communications occur in a highly dynamic and variable environment with no entity in charge of the network management. To deal with the unstructured nature of VANETs it is essential to include a mechanism for node discovery, so that vehicles are aware of the presence of neighbors. In addition, vehicles have to be able to distribute information among a group. The scope of this group varies depending on the protocol and the scenario: it can be a message to be distributed only to nodes playing a special role in the network, such as cluster heads, or to every single node in a given area in the case of geocasting protocols. Most of the solutions designed for VANETs include a flooding or broadcast mechanism by which nodes can announce their presence, position, speed or direction among other information. These messages are sent periodically in order to effectively deal with the very dynamic nature of these networks. Depending on the density of the network, the distribution of these messages can originate a so-called *broadcast storm* if the transmission is very frequent. On the other hand, if the transmission interval is too long, the discovery may be less effective, so there is a trade-off between signaling overload and effectiveness of the mechanism.

- **Reliable signaling**. After performing the node discovery, multi-hop ad-hoc protocols may use signaling messages to set up paths between selected nodes. This signaling may have different scopes, may be multicast or unicast and can be sent at different rates. Reliable delivery is critical as signaling often involves creating or updating the protocol state. VANET solutions might make use of confirmations or acknowledgments to make network nodes aware of potential failures, and avoid in this way state inconsistency in the network.

- **Use of cellular communication**. A cellular communication channel is sometimes present in the VANET scenario as a reliable always-on connection, providing backup to the unstable and intermittent ad-hoc network. Many solutions transfer critical or time-constrained information via the cellular access network. Interestingly, this cellular connection has been one of the main causes of failure in our experiments, having considerable packet losses and very variable delay. Therefore, we claim that this common assumption on the reliability of the cellular connection cannot be taken for granted in real world scenarios.

- **Mechanisms to deal with link quality variability**. The existence of a symmetric communication channel is also frequently assumed when designing a VANET protocol. However, VANETs conditions are very dynamic, and the reception of a message from a node cannot translate into a subsequent successfully transmission on the way back. Moreover, the use of multicast and unicast signaling messages, which may be sent at different rates in the WLAN, impose different transmission distance ranges. In addition, in the real world, the communication range and the link quality are affected by many different and variable factors, such as vehicles' speed, taller vehicles passing by or just the weather conditions. Link conditions impact strongly the lifetime of a multi-hop route and therefore, it becomes essential to monitor link status in order to maintain the communications quality as long as possible.

2.3.1. VARON: a VANET solution for NEMO route optimization

This section provides a brief overview of the solution experimentally evaluated in this article: *Vehicular Ad-hoc Route Optimization for NEMO* (VARON). The interested reader is referred to the original design of VARON [1] for further details. Note that the present article introduces some modifications to the original design, based on the key knowledge acquired by implementing a real prototype, as it will be described in the next sections.

VARON enables the optimization of vehicle-to-vehicle communications in a secure way by combining a network mobility approach – that supports vehicle-to-Internet communications via the 3G interface – and a vehicular ad-hoc approach – used when a multi-hop network becomes available (i.e., communication takes place between vehicles that are close enough to communicate through a VANET formed by the mobile routers deployed within those vehicles, and perhaps within other vehicles in their surroundings). We do not consider intra-vehicle communications, as we manage the mobile network as a whole, represented by the entity of the mobile router. The route optimization process takes place as follows:

1. *Self-organization and discovery of reachable networks.* Each mobile router needs to find out which other MRs are available within the VANET, that is, which mobile network prefixes (MNPs) are reachable through its ad-hoc interface. To that purpose, every MR pe-

riodically broadcasts a message called *Home Address Advertisement* (HoAA), which contains its home address and an associated lifetime. These messages are announced through the ad-hoc interface using a hop-limited flooding, so every MR becomes aware of the MNPs that can be reached through the VANET. This mechanism makes a node visible to other nodes potentially interested in establishing a direct communication through the VANET. In our case, in order to avoid missing optimization opportunities and not to flood the network with signaling, these messages are sent at a configurable interval, which was set to 10 seconds in our experiments[4].

2. *Reliable signaling for the creation and validation of a secure ad-hoc route.* The ad-hoc routing protocol should at least provide the same security level than today's Internet communication. The mechanism used by VARON to set up and maintain a secure ad-hoc route is based on [4], modified and extended to fulfill the requirements of a network mobility-based vehicular scenario:

    - First, once a mobile router has identified an optimization opportunity, this MR (called *originator MR*) has to trigger the ad-hoc route setup by sending to its one-hop neighbors a *Care-of Route Test Init* (CoRTI) message. This message is re-broadcast by intermediate routers using a limited flooding until the message reaches its final destination, which is the mobile router handling the target prefix (called *target MR*).
    - Second, the target MR generates a reply message called *Care-of Route Test* (CoRT) message, and unicasts it back to the MR that triggered the procedure. Note that the route used to deliver the CoRT message is learned by the intermediate mobile routers during the limited flooding of the CoRTI message, and that the reverse route is learned while delivering the CoRT message.
    - Third, this new bi-directional ad-hoc route (called *care-of route*) cannot be used yet to forward packets between the mobile networks managed by the originator and target MRs, as it has to be verified that the MRs are actually authorized to manage those prefixes. This requires exchanging some additional messages, called

[4]Note that in our experiments we tried several configurations, achieving best performance for the 10s interval.

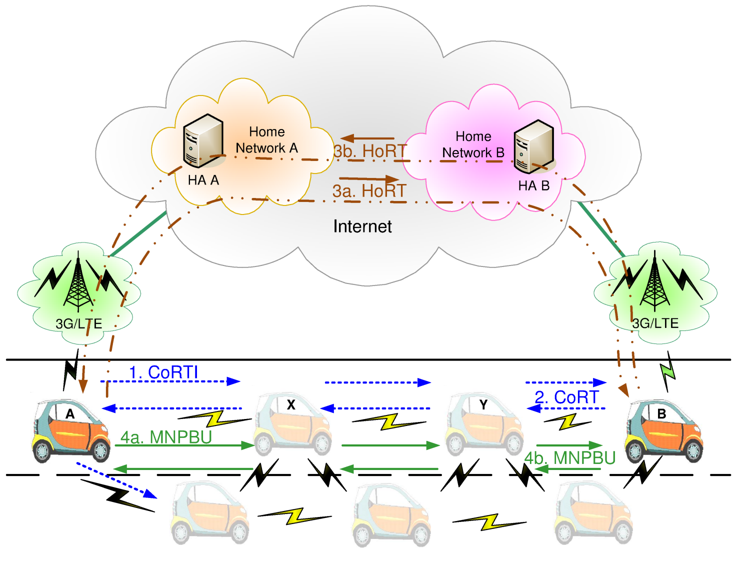

(a) Care-of Route discovery and validation.

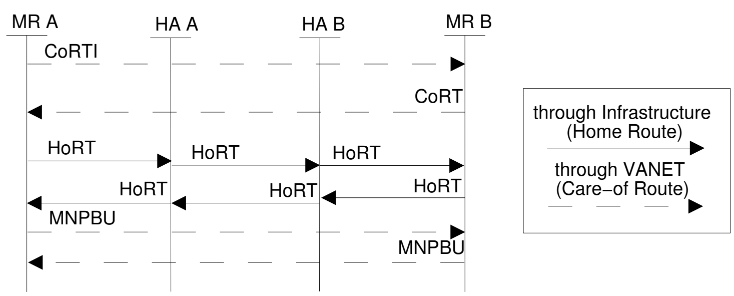

(b) Care-of Route authentication signaling.

Figure 2: VARON signaling.

Home Route Test (HoRT), using the default NEMO route through the 3G interface (called *home route*), and some final messages, called *Mobile Network Prefix Binding Update* (MNPBU), through the VANET.

VARON signaling is secured using cryptographic mechanisms, as extensively explained in [1]. Figure 2 shows a simplified example of the route optimization process.

In VARON, these signaling messages define the state machine of the route optimization process, which needs to be properly completed. The nodes involved have to monitor the current state at every moment, in order to ensure consistency at both ends of the communication. Note that some of the signaling messages are transmitted through the cellular communication channel, which proved not to be very reliable and sometimes introduced a considerable delay. A simple but effective way to perform this monitoring is the use of timeouts to control that these messages are received within a certain interval. In order to avoid wasting time trying to optimize a route when the link conditions are not favorable, VARON checks the quality of the link during the initial signaling stage and aborts the optimization in case the quality is not good enough. In this way, we deal with the link quality variability that may cause a failure in the optimization process or the creation of routes that can be used for a very short period of time (which make them practically useless).

If an ad-hoc route in use becomes invalid because it expires or it is broken, a *Care-of Route Error* (CoRE) message is sent (and forwarded) by each MR in the path to the originator MR. For example, let us consider the scenario in Figure 2(a) where MR A and MR B have an ongoing communication using a multi-hop route in the VANET (MR X and MR Y are intermediate hops). If MR Y is forwarding data from MR A to MR B, and detects that the link to the next hop in the path (MR B itself) is broken, then MR Y sends a CoRE message towards the source (MR A). This message is received by MR X and forwarded to MR A, who notifies MR B the withdrawal of the care-of route with a CoRE message sent through the Internet connection. Then, MR A and MR B switch back to the home route, though they may start a new route discovery procedure to set up a new optimized care-of route within the VANET.

Due to the link quality variability commonly found in VANETs and the need to deal with several sources of interference, it is very important to continuously monitor the link quality. As soon as a quality degradation is detected, the optimized route is withdrawn and all the communications take place via the cellular interface. A significant part of our experiments focused on the configuration of the most adequate link quality thresholds, both to establish a multi-hop route in the VANET and to withdraw this route, falling back to the default Internet connection. The design of these thresholds is extremely important in order to avoid short-lived routes, which additionally incurs useless signaling overhead in the network, which might also disturb other users' communications.

Some of the aforementioned mechanisms are included in our implementation as a result of the experimental learning process, as we will explain in more detail in Section 5.2.

2.4. VARON adaptability to current VANET standards

In the last decade, many standardization bodies have been working on standard communication architectures for vehicular networking. Among them we mention the ISO CALM[5], the European Telecommunications Standards Institute (ETSI) TC ITS[6] and the IEEE 1609 standards for Wireless Access in Vehicular Environments (WAVE)[7]. All of them leverage on a particular amendment of the IEEE 802.11 standard (i.e., IEEE 802.11p) which has been included into the 2012 version of the document [5]. All the aforementioned standards are particularly focused on safety purposes, being it the primary reason to introduce short range wireless communications in vehicular environments. The ETSI, for example, defines two kind of messages for this purpose: Cooperative Awareness Messages (CAMs) and Decentralized Environmental Notification Message (DENM). In order to integrate VARON-related information into the already standardized cooperation messages, the VARON original protocol would need to undergo a heavy modification. Because of that, and the fact that the VARON protocol semantics does not really overlap with the cooperation awareness messages, we decided not to overload these existing cooperative awareness messages to also convey VARON related information. Moreover, the initial setup envisioned for VARON did not include a positioning device. Thus, we implemented VARON using plain IPv6, leaving geographic improvements for a future extension.

3. Related Work

Research on vehicular communications has extensively addressed many different aspects, from routing protocol design to location privacy or peer-to-peer file sharing. So far, only a minority of the existing research includes experimental results, due to the considerable challenges posed by real-life experimentation.

[5] `http://calm.its-standards.info/`

[6] `http://www.etsi.org/technologies-clusters/technologies/intelligent-transport/`

[7] `http://vii.path.berkeley.edu/1609_wave/`

In order to maximize the efficiency of an experimental vehicular setup, some renowned institutions have opted for deploying their own stable vehicular testbed, as we describe next. The Campus Vehicular testbed at UCLA or *C-Vet* [6] is a comprehensive testbed deployment, open to external researchers too, which includes mobile nodes with several wireless interfaces and a wireless mesh infrastructure. The *VanLan* testbed at Microsoft campus in Redmond [7] is formed of eleven access points (APs) distributed all over the campus, and two vans equipped with mobile nodes that cross the campus several times a day, so the research group can take advantage of this for research purposes. The *DieselNet* and *UMass DOME* testbed at Amherst[8] [8] are delay-tolerant networks deployed in 40 buses and a mesh network built in cooperation with the city of Amherst. The deployment of a stable vehicular testbed makes possible to run experiments frequently, to evaluate and compare networking protocols as well as to design models for mobility and traffic patterns validation, among many other applications. In addition, researchers can foster the experimental evaluation of vehicular communications by opening their testbeds to other research groups and sharing their valuable experience in the testbed deployment [9], [10]. However, apart from C-Vet, which also provides V2I communication, these platforms cover only single-hop V2I communication scenarios, not looking into multi-hop V2V/V2I communications, such as the ones considered by VARON.

A more moderate and frequent approach present in the literature is the use of a wireless network infrastructure already deployed. Vehicles roam within this network in order to validate a particular vehicular networking protocol and/or assess its performance. In this way, researchers may extract significant insights from real-world experiments at a lower implementation cost. For example, Deshpande et al. [11], benefit from a metro-scale WiFi deployment provided by an ISP to download files during test drivings of different lengths; and Giannoulis et al. [12] focus on the evaluation of channel quality and design an access point quality scoring mechanism to evaluate WiFi performance in the vehicular environment, thanks to the urban wireless mesh network deployed by the TFA (Technology For All) in Houston in cooperation with Rice university[9].

Following this approach, most of the existing works in the literature are

[8] https://dome.cs.umass.edu/
[9] http://tfa.rice.edu/

centered on delivering data to or from moving vehicles by means of WiFi APs opportunistically accessed along their way. *Cabernet* [13] has been deployed in 10 taxis in Boston and presents a transport protocol to avoid the shortcomings of TCP when dealing with 802.11 networks, and a scanning mechanism to reduce delay in the wireless association process. Note that in VARON, ongoing communications are not delayed by the latency in the wireless association, as data is being transferred by the cellular connection. Authors in [13] also implement timeout optimizations to avoid losses. Similarly, VARON switches to using the WLAN route only if the network is reliable. In relation to that, several works study the most appropriate hand-off technique and try to predict WiFi connectivity to avoid losses in the data transfers from or to a vehicle, but very few tackle the issue of vehicle to vehicle communications. Authors in [14] measure packet delivery ratio and packet inter-arrival time between vehicles that travel together in a 960-km long test drive. However, their measurements are based on the transmission of beacons and their successful reception, whereas we look for a more general solution. *ViFi* [15], tested in VanLan and DieselNet, modifies the wireless driver to evaluate different hand-off strategies in an ad-hoc network and implements a relaying mechanism with a main AP (namely the anchor) and several auxiliary APs to forward data to its final destination. However, this mechanism can only operate in a deployment where all the access points are working on the same channel and requires vehicles to send beacons at regular intervals. In addition, they evaluate different handover strategies, based on Received Signal Strength Indication (RSSI), beacon reception ratio and performance history. These last two mechanisms cannot be adopted by VARON, since in a wireless ad-hoc network not all the nodes are sending beacons, nor it is feasible to maintain a history record of performance. Based on that, we concluded that RSSI monitoring was the only possible choice for VARON.

Authors in [16] also play with the 3G-WiFi interaction and present *Wiffler*, which postpones transmission of delay-tolerant data until a stable WiFi connection becomes available, and switches back to 3G if the packet cannot be transmitted fast. This approach uses WiFi as an auxiliary tool to improve 3G transmission and avoids switching to 802.11 unless it can provide quality communications. However, the offload of the cellular connection translates into delaying transmission of non-critical applications, which may not satisfy the user. VARON switches from using 3G connectivity to using a wireless multi-hop ad-hoc route between neighboring vehicles, taking advantage of the locality of communication end-points. Moreover, a potential enhancement considered for future work is to offload some flows, instead of all the traffic, from one route to the other, taking into account traffic characteristics and mobility patterns. In a similar way, authors in [17] present a network stack, *CafNet*, which lets the application decide which data to send when a WiFi connection becomes available, instead of transmitting "outdated" information buffered by the link layer when there is a connectivity change. Contrarily, VARON does not defer or discard any transmission, but it modifies the routing when the conditions in a wireless multi-hop path are favorable.

In line with the work we present, [18] conducts tests using off-the-shelf devices with Linux OS and only two vehicles for introducing a rate adaptation mechanism for the vehicular environment (namely, CARS). They perform tests in scenarios with different mobility requirements: from cars being static to high speed or intermittent connectivity. They also ran their tests in the 802.11a frequency band and experienced similar RSSI variations when the vehicles are moving.

In summary, there are recurrent issues in vehicular experimental works: *i)* the connection establishment latency and the association time in 802.11 networks, *ii)* the choice of a proper handover strategy, and *iii)* the data transfer methodology, as the network dynamics involve frequent disconnections and lossy links. In order to enhance communications and avoid these problems, most of the proposals design predictive methods, based on caching or keeping history of the signal strength, link quality or performance associated to an access point in a certain location and rely on familiarity of routes and paths [7], [15], [16], [19]. VARON aims at optimizing inter-vehicle communications. Therefore, keeping track of the signal quality at a certain geographical point is meaningless, since the vehicles may not roam around that point in the future and even in that case, there is no guarantee that the wireless link between them would keep the same conditions.

Vehicular testbeds are also used to carry on research on issues related to the physical (PHY) layer. A very comprehensive work is the one performed by Shivaldova et al. in [20] and references therein, in which several aspects of the wireless channel issues are evaluated, showing patterns similar to the ones we obtained with our experiments. The same authors also proved the feasibility of multi-hop wireless communications in [21], showing that the coverage of a Road Side Unit (RSU) can be extended (in space) up to 30%. This result supports the VARON approach, as it shows the benefits provided by the extension of the wireless coverage using a multi-hop path.

Antennae performance is of fundamental importance during the development and test process of a vehicular networking protocol as VARON. In [22] the authors provide a broad evaluation of the achieved throughput for vehicle to RSU communications using different antennae placed at different locations of the vehicle (the rooftop and the windshield). The results show substantial difference in the performance depending on the chosen setup. Similarly, [23] evaluates the impact of obstructing obstacles (even just another vehicle) and how this non-line-of-sight (NLOS) communication decreases performance. They show experimentally the importance of line-of-sight (LOS) conditions and how the effect of the attenuation is highlighted the closer the obstacle is to the sender. In this paper we wanted to focus on the implementation aspects of VARON rather than considering antennae-related issues in detail. So we chose an omni-directional antenna mounted to an elevated structure in the center of the rooftop, in order to guarantee line-of-sight connectivity. More details about the experimental setup are detailed in Section 5.1.

4. From simulation to the lab: Experiments in a controlled environment

VARON was initially validated by extensive simulations using OPNET Modeler[10]. The results showed an impressive performance improvement compared to existing approaches. These performance gains appear as a direct result of opportunistically exploiting multi-hop communications via the dynamic set-up of the VANET. As reported in this study, these simulation-based experiments turned out not to fully resemble reality, missing some important aspects of real-life wireless behavior.

In this section we report about the implementation and deployment of a lab-based VARON prototype, as well as the testbed built for its validation and performance assessment.

4.1. Network scenario

In order to experimentally validate VARON in the laboratory, the network scenario presented in Figure 3 was deployed. During this phase (lab tests), all the network elements were physically deployed in a laboratory at the University Carlos III of Madrid. The implementation of each entity and the

[10] `http://www.opnet.com/solutions/network_rd/modeler.html`

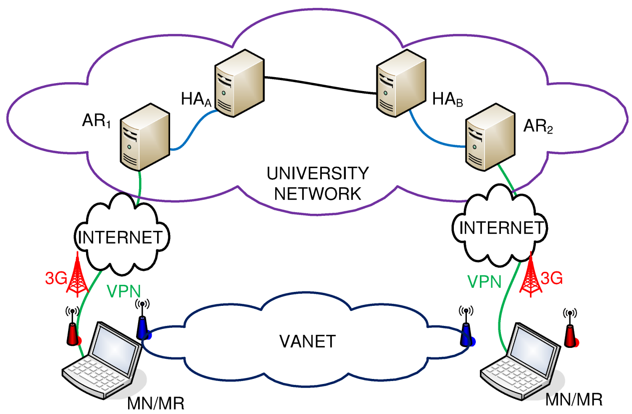

Figure 3: Complete network and VARON prototype scenario.

Table 1: HW platform for the different nodes in the network.

Device		CPU	RAM	OS version
MR	**Router**	266MHz	32MB	OpenWRT (Kamikaze)
	Laptop	Dual core 1.33GHz	2GB	Ubuntu Linux
HA		Dual core 1.66GHz	2GB	Ubuntu Linux
AR		Dual core 1.66GHz	2GB	Ubuntu Linux
MNN		Dual core 1.66GHz	2GB	Ubuntu Linux

description of the main elements that are part of the prototype are described next:

- The mobile router (MR) is the key element in the VARON prototype. In the initial phase of the experiments, this entity was implemented in an Asus WL-500g Premium router, a low cost device equipped with a 266 MHz CPU and 32 MB of RAM. In spite of its limited capabilities, it was initially selected[11] because of its low cost, acceptable performance (as shown in previous works in which we used the same hardware for similar purposes), the possibility of customizing its firmware according to our needs and, most importantly, the flexibility in terms of potential number of network interfaces. The device is equipped with two USB

[11]Later in this article we show that this decision proved to be wrong.

ports that can be used to increase the number of network interfaces, e.g., by connecting a 3G USB dongle. Being able to deploy a heterogeneous network is essential for the functionality that VARON aims to offer. In addition, we replaced the original miniPCI wireless network card with one using an Atheros family chipset, which also supports the 802.11a mode.

- Each vehicle behaves as a mobile network managed by its mobile router, and implements the Network Mobility Basic Support protocol. This requires the deployment of home agents (implemented in two PCs) at their respective home networks, and access routers (ARs) at each of the visited networks. On a first stage, all the tests were performed indoor in the laboratory and WLAN was used as the technology to connect to the infrastructure (instead of 3G). On a second stage, some tests were performed outdoor and 3G was used as access technology. The access routers were physically located at the laboratory and, since we used a commercial IPv4-only 3G network, we set up a Virtual Private Network[12] (VPN) over the public Internet to offer IPv6 connectivity between the access routers and the mobile routers.

- A USB 3G dongle (Huawei E1752C) was used as the additional network interface providing the mobile network access to the infrastructure in order to reach its home network. The MR used the 3G connectivity to reach an IPv6 access router, via a VPN tunnel that hides the traversal of the public Internet and emulates the direct connection (i.e., one-hop distance) between the mobile router and the access router.

- A mobile network node (MNN) was attached to each mobile router. The role of MNN was played by a netbook. Each of these nodes acquires an IPv6 address belonging to the mobile network prefix (MNP) managed by its MR.

4.2. Software implementation

This section is devoted to describing the different software modules of the VARON implementation. The key entity of the prototype is the mobile router, which implements the modules shown in Figure 4.

[12]OpenVPN: Open source VPN, available at `http://openvpn.net/`

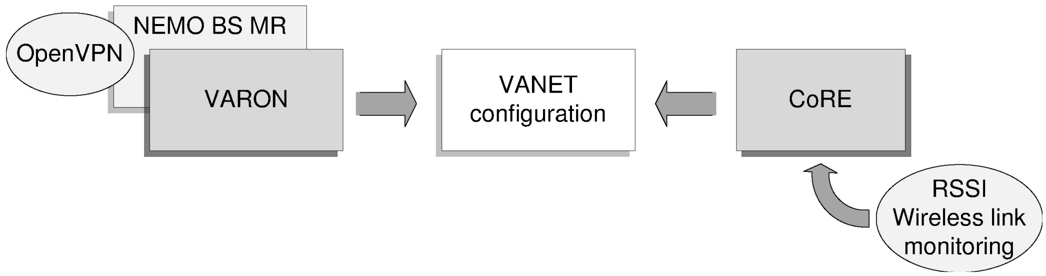

Figure 4: Software modules implemented in the mobile router.

The `NEMO BS MR` module performs the tasks of a mobile router according to the NEMO B.S. protocol: it performs the signaling exchange with the home agent, and configures the bidirectional tunnel used as default data path (home route), keeping the reachability of the mobile network while moving.

The `VARON` module performs all the signaling defined by the protocol, taking care of setting up a care-of route and then establishing a tunnel between the mobile routers involved in the VARON optimization. Once this process is completed, the traffic flowing between the mobile nodes managed by the MRs is forwarded through the care-of route, instead of using the default one to the home network through the Internet, with the subsequent delay reduction and bandwidth gain. Both, `VARON` and `NEMO BS MR` modules are developed in C.

Given the dynamic environment of vehicular networks, care-of routes may not last for a long time. The original VARON design specified the use of Care-of Route Error (CoRE) messages to signal when a care-of route should not be used anymore. However, the mechanisms that could be used to trigger CoRE signaling were not fully specified in [1]. This is actually one of the aspects of the VARON design that we wanted to analyze in more detail and improve as part of this experimental work. The `CoRE` module is in charge of monitoring the wireless links used by active care-of routes, in order to avoid packet losses due to insufficient link quality. The `ath5k` wireless driver[13] used in our prototype provides Received Signal Strength Indication (RSSI) measurements of every node in communication range, which can be used as a metric to detect link quality degradation with the next hop of an optimized (care-of) route. The VARON original design proposed monitoring

[13]`http://wireless.kernel.org/en/users/Drivers/ath5k/`

layer-2 acknowledgments as a possible mechanism to assess link quality but we discarded this mechanism because in our experiments it was not able to detect quickly enough link quality oscillations. We adopted an approach used by other works found in the literature (see Section 3): to take the RSSI as an indicator of link quality, although its very oscillating nature makes its use very challenging.

VARON internally sets an RSSI threshold: values below this threshold indicate that the communications are likely to fail. VARON monitors the RSSI using a sample-averaging algorithm to avoid frequent switching between an optimized care-of route and the default home route. The algorithm used for this purpose leverages on [24], which concludes that the Weighted Mean of 3 Samples works as a suitable low pass filter for high mobility scenarios. Hence, the implemented algorithm uses the following formula:

$$y[n] = \alpha \cdot x[n-2] + \beta \cdot x[n-1] + \gamma \cdot x[n]. \quad (1)$$

The three weights[14] are distributed in order to weight more the most recent sample as compared to the rest. For the exact parameter setup, we have performed an experimental matching, testing for different values until we found the best one for our purposes. The optimal configuration of these parameters for any kind of scenario is not a trivial task and is out of the scope of this work. The RSSI sampling frequency is a configurable parameter to be able to modify it according to the test running and the kind of traffic.

When the value of $y[n]$ resulting from (1) computed by a mobile router goes below the configured threshold, a CoRE message is sent to the other endpoint of the VARON tunnel through the Internet (default) route. The transmission and/or the reception of this message by an MR automatically withdraws the optimized route, falling back to the default route over the Internet.

4.3. Validation in a controlled environment

The first validation and experimental analysis of VARON was conducted indoor, inside a networking lab of our university. As it will be shown later, conclusions obtained from this first validation also have a lot in common with

[14] In the final VARON experimental evaluation, reported in Section 5.3, we used the following values: $\alpha = 0.1$, $\beta = 0.3$ and $\gamma = 0.6$.

Table 2: Time needed for establishing VARON optimized route.

Care-of route length (hops)	**VARON signaling delay [ms]**
1	549.6±26.5
2	853.5±61.7
3	1166.5±76.4
4	1434.96±47.7
5	1759.8±63.5
6	2055.2±85.1
7	2377.5 ±97.6
8	2639.4±74.8

other laboratory-based efforts: observed behavior and performance deviates from the one experienced in a real-life trial.

These initial tests focused on assessing the feasibility of the protocol and validating the prototype, as well as analyzing the impact of the length of the care-of route on the performance in a non-mobile scenario.

The first tests involved only two mobile routers, using also IEEE 802.11 as wireless technology to connect to the fixed infrastructure (instead of 3G), increasing the number of MRs subsequently until reaching a maximum of 9 nodes (i.e., 8 hops). All the nodes were located inside the same room, so there was direct radio connectivity among them. In order to emulate a multi-hop scenario, we used `ip6tables`[15] to limit the reachability of the nodes at the IP layer. It should be noted that with this approach the channel contention is higher than the one occurring in a "real" multi-hop scenario. Due to limitations in terms of available physical space to deploy the mobile routers, we believed this was a good *compromise* scenario, which allowed us to evaluate – in a worst-case scenario – the performance of VARON in multi-hop situations.

Table 2 shows the time needed to complete the establishment of the care-of route for different number of hops. Most of that time is spent by the mobile router performing cryptographic operations (using a key size of 512 bits), essential for security enforcement but very time consuming for a device with limited capabilities such as the nodes we used in these tests. These cryptographic operations mainly consist in signing/verifying the signatures

[15]`ip6tables` is a Linux kernel tool that let the user examine and configure the tables of IPv6 packet filter rules.

of each VARON message, using the RSA algorithm. As described in more detail in [1], VARON uses Cryptographically Generated Addresses (CGAs), and performs hop-by-hop signing of the CoRTI, CoRT and CoRE messages. This allows to authenticate not only the original sender of the message, but also the forwarder. This procedure is used to authenticate the multi-hop route established in the VANET. Note that during the optimized route setup time, the data connection is never interrupted, as traffic is forwarded using the home route, always available.

5. From the lab into reality: On-road experimentation

In the previous sections we have reported on the initial evaluation efforts of VARON, which started with extensive simulations and continued with an experimental validation in a controlled environment, allowing us to check the correctness of the protocol and serving as a proof of concept of VARON. In many vehicular research works, the experimental analysis stops here, i.e., solutions do not *leave* the simulator or a controlled-environment laboratory. This article attempts to go beyond this, by proving the feasibility of the mechanisms proposed by VARON under realistic conditions. This section describes the deployment of the vehicular network prototype, the experiments conducted, as well as the main results and the lessons learned from this experience.

5.1. Scenario and testbed deployment

Starting from the VARON prototype evaluated in a laboratory, we followed an incremental approach to improve and further develop the prototype while validating it under real life conditions[16].

In the first phase of our validation work, conducted in the laboratory, we decided to use commercial off-the-self (COTS) devices. We selected routers like the ones that can be commonly found at home or small-office environments, as they are very convenient in terms of flexibility, cost, and size. This type of hardware usually carries out quite a lot of operations, such as network address translation (NAT) – involving application layer gateway functions that typically require inspecting and processing every forwarded

[16]Some of the first changes aimed to experiment with real vehicles were first tested in our lab, as for example the use of 3G and the VPN setup. This saved us from spending valuable time in the field.

packet – firewalling, P2P downloading, etc. Our initial assumption was that this hardware would be powerful enough to support VARON operations, and that it would better resemble the type of hardware that car manufacturers would be willing to deploy in cars, because of its reduced cost. This assumption held while operating with one single network interface in the lab tests, but as soon as we introduced the external 3G interface, and the use of a VPN connection (involving non negligible cryptographic operations), we realized that we needed more powerful hardware, which was also convenient for debugging purposes. All this led us to upgrade the hardware of mobile routers, wireless cards and antennae. We replaced the router in the field trials phase with a laptop, capable of easily handling all the processes that run in parallel in the mobile router. The laptop was equipped with an 802.11a/b/g Ubiquiti SRC wireless card with an MMCX plug for an external antenna at the outdoor testing phase. At the operating frequency and the transmission data rate, the sensitivity of the wireless card is $-94dBm \pm 1dB$[17]. Additionally, the communication range of the WLAN link was enlarged by attaching an omnidirectional antenna with a higher gain (8 dBi) working in the range of 5 GHz. We did not use an 802.11p card because, at the time of the experimental work, there were few available and they were very expensive. In our work we mostly use low cost COTS devices and we focused on the networking side of the stack. Developing an 802.11p standard MAC requires kernel coding as, currently there is not any open source implementation [25]. Therefore we decided to use 802.11a, which uses the ISM band (i.e., 5 GHz) closest to the one used by 802.11p.

The first round of experiments involved two vehicles, and was aimed at validating the VARON route optimization and evaluating the feasibility of inter-technology handovers. On a second round we added a third vehicle, in order to assess the multi-hop route optimization and estimate the suitability of the multi-hop mechanism.

Each of the vehicles (see Figure 5) was fully equipped with a mobile router (a Linux-based laptop) installed on the car's roof top, next to the antenna, while the mobile network node (a Linux-based netbook), connected to the MR by Ethernet, is placed inside the car. Both devices are powered by means of an AC/DC adapter.

The selection of a proper scenario for the test drives is an important deci-

[17] `http://dl.ubnt.com/src_datasheet.pdf`

Figure 5: A snapshot of the experimental set up.

sion. We chose the Leganes scientific cluster[18] because of two main reasons: *i)* it presents a relatively low traffic load (to prevent us from obstructing other vehicles), and *ii)* it has a regular street design, with roundabouts and straight stretches, which is a key point for the repeatability of the experiments. At the same time, the test course is not isolated and its characteristics keep our experiments close to a real scenario.

From the very beginning, the design of the experiments took into account *repeatability* as one critical requirement. Figure 6 shows the trace of the vehicles involved in the experiments for the 2-node and the 3-node experiments. Each testing round consisted of the following steps:

1. The starting position of each vehicle is pointed by the text boxes labeled MR1, MR2 and MR3. The optimized route in the VANET is established between MR1 and MR2. In the case of the 3-node experiment, MR3 acts as intermediate hop. Note that at this initial position no direct MR-to-MR connectivity is possible between MR1 and MR2.
2. The vehicle labeled as MR2 starts getting closer to MR1.
3. When they are close enough to enable WLAN communication, they start the VARON route optimization signaling, which results in setting up a care-of route if the process ends successfully. Otherwise, the optimization process is aborted, until the next optimization opportunity.

[18] `http://www.leganestecnologico.es/`

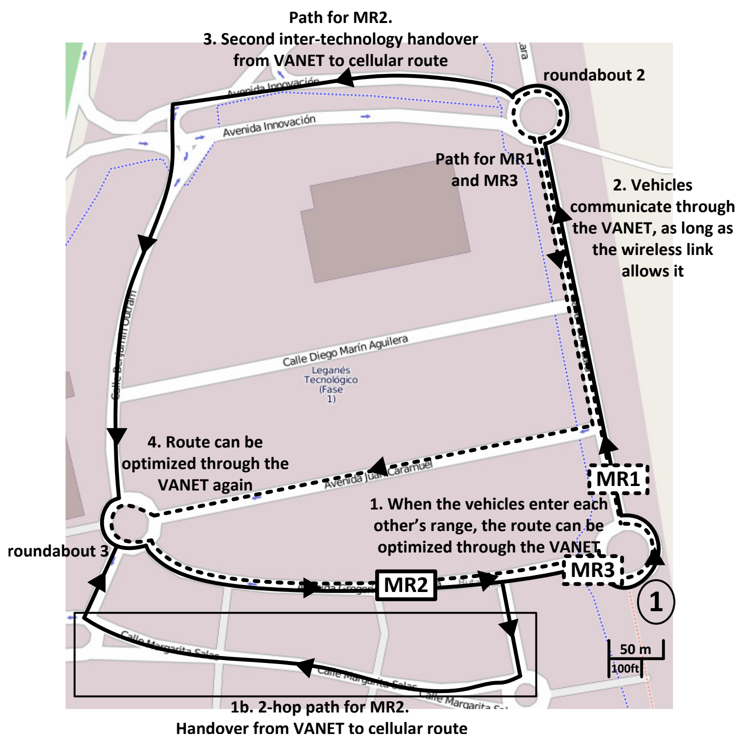

Figure 6: Test itinerary followed in the Leganes scientific cluster.

In the case of the 3-vehicle scenario, MR2 initially follows the path enclosed in a box at the bottom of Figure 6. As MR2 moves away from MR3, the link between them is broken, the route falls back to the 3G connection, until MR2, which continues approaching, is close enough to enable the WLAN communication again.

4. Vehicles move forward together until reaching the intersection at the roundabout 2, where they take separate paths. MR2 continues by the road at the upper part of the map, while MR1 (and MR3 too in the 3-node experiments) take the dash-lined path. This split compromises the wireless link quality and thus forces a handover back to the cellular connection, triggered by a CoRE message.
5. Vehicles continue on their corresponding designated paths, which lead them to meet each other again at the roundabout 3, enabling the establishment of a care-of route again. Note that in the 3-node experiment, MR1 and MR3 follow the same path, leaving some distance between them.

6. Finally, the vehicles continue towards their initial positions, inducing a second handover to the cellular default route because WLAN connectivity is not possible from their starting positions.

Therefore, a complete round accounts for two inter-technology handovers from the cellular network to the WLAN and another two in the other way. This allows VARON to recover from an optimization and to confirm that both MRs are able to start a new route optimization, if required. Each test was repeated a minimum of 20 times, both for the single-hop (2-node) and two-hop (3-node) scenarios. The average speed of the vehicles during these tests was 50 km/h.

To detect when to switch from the care-of route to the home route, VARON considers the link quality metric provided by the RSSI. As mentioned before, the use of RSSI for this purpose is challenging, due to the dynamism of VANETs. In order to fine tune our algorithm, we first analyzed how the WLAN RSSI and the connectivity varied with the distance between nodes. We selected the thresholds for triggering the handover from the WLAN to 3G based on these results, which are shown in Figure 7. The gray areas denote lack of WLAN connectivity due to packet losses and errors. Then, we need to anticipate those losses by sending a CoRE message to the two communication ends, as explained in [1] and Section 2.3.1. Obtained results provide some interesting insights: *i)* the degradation of the link quality due to the increasing distance between nodes is evident, but *ii)* distance is not the only determinant factor. For instance, in light of the differences observed in the measurements when the nodes are moving farther or closer to each other, we can claim the importance of their relative positions, the location of the antennae (i.e., their alignment depends on the actual shape of the car) and the multi-path reception.

5.2. Early testing and implementation feedback

As we will describe next, to fully deploy VARON in a real scenario we had to adjust and improve the original VARON design in order to tackle the different issues we found in the process, which relate to the common VANET mechanisms identified in Section 2.

5.2.1. Effectiveness of the limited flooding

VARON implements a limited flooding mechanism for announcing the presence of a mobile network prefix in the VANET. Every mobile router is

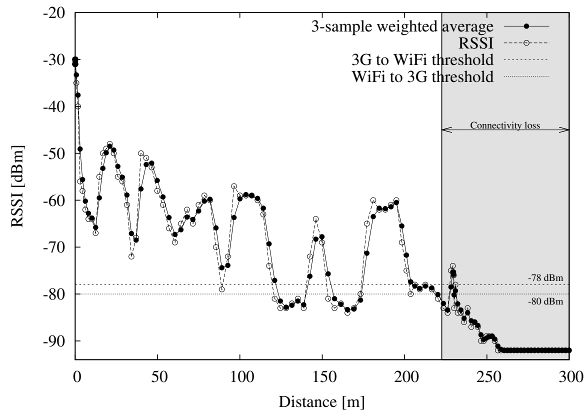

(a) Nodes are getting farther

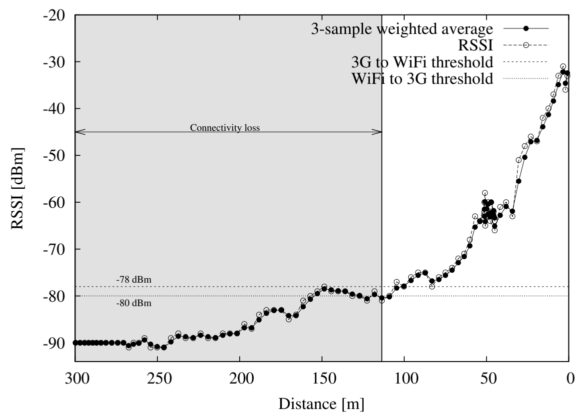

(b) Nodes are getting closer

Figure 7: Experimental RSSI degradation due to distance between nodes.

aware of the prefixes that are available within a configured advertisement scope and can decide whether to start a route optimization process or not. These announcements (i.e., the HoAA messages) are broadcast and forwarded by every node until a hop limit counter, included in the message itself and decreased on each hop, reaches zero (this defines the advertisement scope). Before forwarding it, every mobile router waits for a random uniformly distributed time, in order to reduce the collision probability and increase the effectiveness of the mechanism. In addition, in order to refresh the information, these announcements are transmitted periodically. Therefore, the transmission interval needs to be chosen carefully. The reception of these messages triggers the route optimization process, so, on the one hand, a long interval would reduce the number of opportunities as the presence of some mobile networks may go unnoticed. On the other hand, the interval cannot be set too short neither, not only to avoid a broadcast storm, but also to keep the wireless links between nodes stable. In our prototype this transmission interval is configurable, in order to allow quick changes in the field if necessary, and we conducted tests with values of 5 and 10 seconds. Experiments with a value of 5 seconds showed a lower discovery rate, as there were more errors in the reception of the messages. Besides, if two vehicles were approaching from a zone where they could not reach each other, there was not enough time for the wireless link to provide favorable conditions. The experiments with a 10-second interval showed a higher percentage of optimization attempts that ended up successfully (i.e., the traffic is offloaded to the VANET care-of route), as there were less collisions, the wireless link between the nodes involved had enough time to stabilize and discovery happened within the time that the two vehicles were likely to have connectivity between them. The use of transmission intervals longer than 10 seconds would be inefficient, as the network conditions could dramatically change between transmissions, leading to much less stable VANET connections.

It should be noted that there are other approaches proposed in the literature, such as [26], that could have been adopted by VARON to improve the effectiveness of the flooding mechanisms. However, the goal of the present work is to assess the performance of the original design of VARON, while introducing as few major changes to it as possible, to better compare the performance obtained from the implementation testing to the one obtained from the previous existing simulation work reported in [1].

5.2.2. Use of timers to ensure optimization state consistency

The wireless medium in the VANET may fail during the route optimization process, as the radio link quality depends on many factors and does not remain exactly the same under any given conditions. As a consequence, VARON messages may be lost or delayed during the care-of route setup. In order to ensure consistency in the VARON state machine and avoid stale states, a modification of the original VARON design was introduced, consisting of the *use of timers to deprecate an ongoing optimization attempt*.

In order to clarify on the use of timers, we consider next a VARON specific example. With VARON, a mobile router may complete all the optimization steps on its side, while the other end does not, for example because the last signaling message to complete the process (an MNPBU) is not received. It is not possible to know whether the last MNPBU was lost or was not even sent (because the first MNPBU is the one actually lost). Therefore, a timer is set when an originator mobile router sends the first MNPBU. If no reply-MNPBU is received before this timeout expires, then the originator mobile router sends a CoRE message to the target MR, to ensure that the state is consistent (no optimization is in place) at both mobile routers.

5.2.3. Is the cellular connectivity that reliable? The need for the CoRE ACK

Cellular connectivity is often assumed to be reliable and to offer full coverage. However, this assumption proved out not to be valid according to our experience in the field. Measured 3G connectivity was variable and unstable, delivering data in bursts very often. Additionally, the available bandwidth was very much dependent on the location and, in general, lower than claimed by mobile operators, exhibiting also unacceptable delays. By default, VARON routes some signaling messages through the cellular network, such as the CoRE message triggering to stop using an optimized (care-of) route when it does not meet the minimum required quality. However, our on-field experiments showed that 3G connectivity was not as reliable as we would have expected from a commercial service, and that sometimes this small signaling message (CoRE) is delayed for a long time or even lost. Figure 8 shows the delay experienced by the cellular connection during one of our experiments. In addition, the average throughput obtained in the experiments transmitting UDP traffic was 247.91 Kbps, far below the nominal channel capacity (7.2 Mbps) and also presented high variability (standard deviation 65.07 Kbps, minimum value 9 Kbps, maximum 678 Kbps) which makes difficult a successful data transmission. The research presented in [27] shows similar measurements which support our experimental results.

This finding made us introduce another modification to the original VARON design: the need for *acknowledging CoRE messages*. When a mobile router receives a CoRE, it has to send back a CoRE ACK message. If the acknowledgment is not received within a pre-configured time window, a new CoRE message is sent, to ensure that the other mobile router receives it, and therefore prevents it from using the care-of route anymore. Note that without the CoRE ACK, if the original CoRE message is lost, the resulting routes between the two mobile routers involved would end up being asymmetric: one would use the default home route, while the other would still use the optimized, but likely non-working, care-of route.

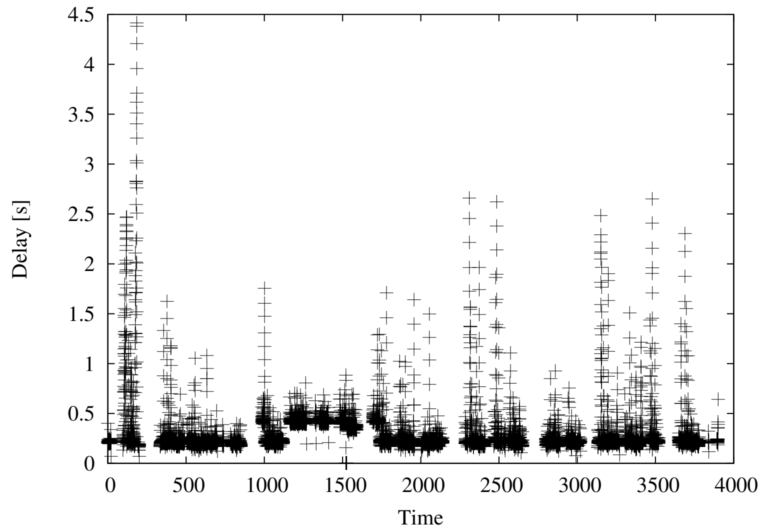

Figure 8: Delay experienced using the cellular connection between two MRs.

5.2.4. Avoiding ping-pong optimization effects and short-lived care-of routes: the use of radio link quality thresholds

In the original VARON design, an optimization attempt was initiated every time a mobile router received a HoAA message from a prefix that was already involved in an ongoing connection. However, our early field experiments showed that VARON was too much reactive under this configuration, leading to a high number of failed optimization attempts and also a high number of short-lived optimized routes. To take into account some of the possible wireless channel issues that may happen in a vehicular environment (e.g.,, weak received signal power or asymmetry of the links) VARON was improved by introducing *radio link quality thresholds* that had to be met before triggering a route optimization setup. It should be noted that HoAA messages are broadcast, so link layer re-transmission techniques are not available. Moreover, the Modulation and Coding Scheme (MCS) used by the network cards when transmitting a broadcast frame has to be set to the lowest (and also the most robust) available. The combination of these two factors may lead to the creation of potentially unreliable links due to the bad channel conditions between two nodes. The choice of the most robust MCS may allow the decoding of data packets even with bad channel conditions, therefore considering a successfully decoded HoAA sender as a valid next hop, whereas the real wireless channel conditions can be far from optimal.

The link quality thresholds, based on the RSSI measurements, are used to decide whether an optimization attempt should be triggered or not. More in detail, a mobile router would not send a CoRTI message if the signal strength towards the transmitter of the corresponding HoAA is not higher than a pre-configured threshold. The same restriction is also applied to the transmission of CoRT messages. This modification makes the VARON route optimization process less reactive, making a node triggering it just when the RSSI associated with the neighbor that sent a signaling message has sufficiently high quality.

This mechanism alleviates the problem of unsuccessful route optimization attempts by discarding those that are likely to fail, although this may still not be enough to increase the average duration of an optimized route. The first experimental results showed a really high correlation between the low quality of the received RSSI during the optimization setup process and the short duration of those routes. Therefore we extended the base mechanism to the whole optimization setup process to further reduce the likeliness of setting up a short-lived optimized route. The mobile router keeps monitoring the RSSI from the transmission of the CoRTI message, which initiates an optimization attempt, until the moment of sending the MNPBU, which completes it. If the average signal strength measured in all the signaling messages received from the mobile router used to reach the target MR is not above the previously defined threshold, the route optimization process is aborted.

This procedure is implemented at both endpoints, as the two mobile routers that are involved in the optimization procedure may discard or not react to received optimization messages if the quality associated to the sender neighbor is not good enough. This kind of optimization can also be applied to the intermediate nodes when setting up a multi-hop route. However, the application of this approach at the intermediate nodes is not as straightforward as at the endpoints. An intermediate node does not recognize it is part of a multi-hop route until it receives a CoRT (CoRTs are the first unicast signaling messages involved in a VARON route optimization procedure). Only in that moment an intermediate node becomes aware that it is being selected to be part of a multihop path being built. Therefore, only at that point an intermediate node would be able to use an effective mechanism (i.e., a CoRE message) to warn the communication endpoints about a possible wireless link disruption. Without the possibility of using CoRE messages, an intermediate node can just silently discard packets, leaving the task of canceling an ongoing optimization process to the timeouts implemented in the endpoints. In a larger setup, having quality thresholds implemented also in the intermediate nodes can help to build better paths but, for the sake of simplicity, we decided to not implement such policies in the intermediate nodes, letting the endpoints to directly manage the transitions in case of link failures.

5.3. Experimental results

In the previous section we have described different improvements to the original VARON design triggered by our implementation experience. We next focus on the actual evaluation of the VARON protocol performance, by presenting the main results and measurements collected during our experiments. While doing so, and although we refer to a particular solution (VARON), the obtained results can also be used to evaluate the suitability of some of the most common networking mechanisms for VANETs, such as multi-hop routing, message signaling flooding or access technology heterogeneity.

Table 3 summarizes the results for the most representative experiments. A first experiment characterizes the VARON signaling delay, which is the time required to send and process all the VARON protocol messages, set up the required tunnels and forwarding, and configure a care-of route. We performed this experiment both in the lab and outdoors, using 3G connectivity and different hardware choices for the mobile router: the Asus WL500g Premium router and a regular laptop. The results show that the limited resources of the COTS devices impact the route establishment delay, which is considerably shorter when the laptop is used. Furthermore, it might be surprising that the delay in the 2-hop case is lower than the average for the 1-hop route establishment. However, most of the time needed for the route establishment is due to the exchange of messages through the cellular interface, which is identical independently of the number of hops in the VANET. The cellular network offers a RTT much higher than the one in the multi-hop ad-hoc network, which occasionally led to slightly lower average delay for the route establishment of a 2-hop route than for the 1-hop scenario. When COTS devices are used, this is harder to notice, due to the limited processing capabilities of these devices.

We highlight that the time required to set up a care-of route does not directly impact the user experience, as traffic is routed normally via the default (home) route while the VARON optimization signaling is taking place. On the other hand, the time required to withdraw an optimized route does

Table 3: Experimental results.

Parameter	Care-of route length (hops)	
	1	2
VARON signaling delay [ms] Indoor, COTS router as MR	908 ±91	1450 ±199
VARON signaling delay [ms] Outdoor, COTS router as MR	1210 ±590	1750 ±850
VARON signaling delay [ms] Outdoor, laptop as MR	550 ±74	514 ±39
Care-of route withdrawal time [ms]	524±233	
Percentage of completed optimization attempts	14.53%	10.84%
Discarded HoAAs	82.90%	55.96%
Aborted optimization attempts	2.56%	13.64%
VARON efficiency	79.3%	60.7%

have an impact, as the quality of a link in the optimized path might become so poor that the path is unusable, and therefore it is critical to anticipate this event and switch to the default route. We measured this “withdrawal time” as the time elapsed since the mobile router detects a link quality degradation triggering the transmission of a CoRE message, until the moment a CoRE ACK is received. Note that since the CoRE and CoRE ACK messages are sent via the 3G interface, the care-of route withdrawal latency is mainly caused by the delay over the 3G network, independently of the number of hops in the care-of route. Figure 9 shows the box-and-whisker plot for these measurements, representing minimum, maximum, median and first and third quartiles for the care-of route set up and withdrawal latencies, for both the single-hop and two-hop scenarios.

Table 3 also shows the percentage of completed VARON optimizations (i.e., the traffic is routed through the VANET) over the total number of optimization attempts. A possible optimization attempt is counted whenever a node receives a HoAA announcing that a prefix which is being used by an ongoing communication through the cellular connection is available through the VANET and therefore, a route optimization is possible. Note that this percentage is quite low because many received HoAAs were simply discarded due to its low associated RSSI (note the high percentage of discarded HoAAs). An optimization process may not be completed because the radio link con-

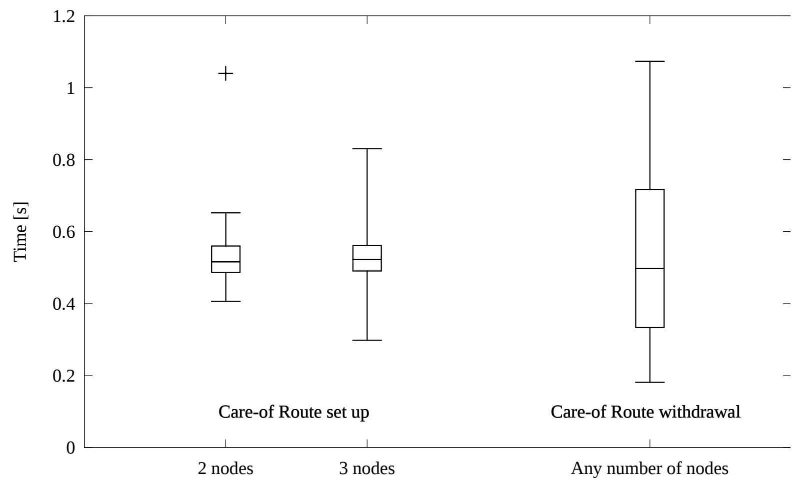

Figure 9: Care-of route set up and withdrawal latencies.

ditions are not good enough, not only at the reception of a HoAA, but also during the whole optimization procedure. It is remarkable that even with this low percentage of completed optimization attempts, VARON achieved a quite high efficiency (around 60-70%). This is actually possible thanks to the use of the RSSI thresholds (not devised in the original design), which avoids attempting to establish a care-of route on a poor-quality wireless (multi-hop) link. This improvement also contributes to saving useful resources in terms of signaling overhead and energy consumption.

Another relevant performance metric is the efficiency of VARON, defined as the amount of time that a vehicle takes advantage of VANET communication instead of using the default cellular connection. The connectivity in the VANET using IEEE 802.11 technologies depends on the speed of the vehicles, their relative positions or their path, among other variables. In order to provide a measurement that is not biased by these factors, we set up IP background traffic over pre-configured routes through the VANET between the two mobile routers involved in the optimization. Note that the static pre-configured routing uses an addressing space independent of the one used by VARON, to avoid any impact on the experiments. We measured the time that an optimized route is in use over the total amount of time that connectivity in the VANET is possible. This experiment also served to confirm the effectiveness of the RSSI threshold that triggers the switch from the home route to the care-of route. Our tests reflect that the route established by

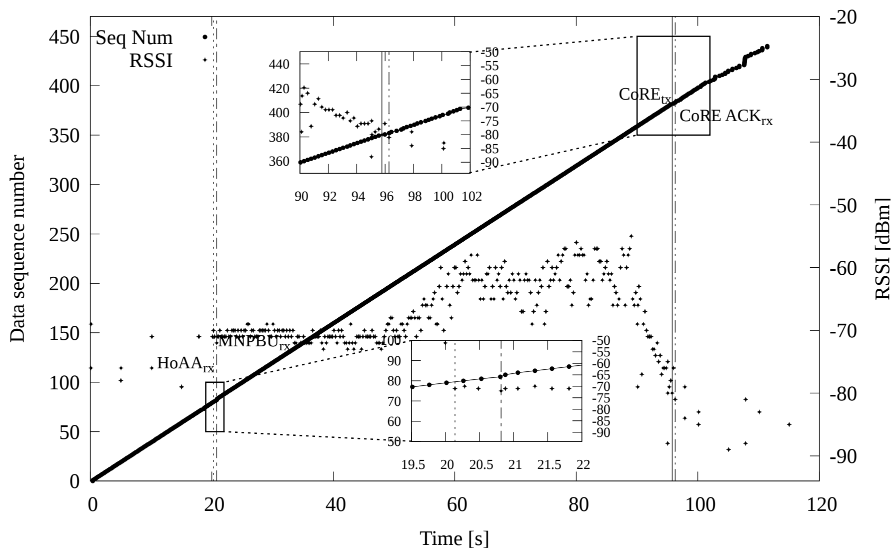

Figure 10: VARON route optimization signaling and data sequence number.

VARON took advantage of the VANET connectivity on average on a 79.3% out the total time with VANET connectivity available for the single-hop scenario, and on a 60.7% for the two-hop scenario.

In order to conclude the experimental validation and provide a complete view of the overall process, Figure 10 shows a snapshot of a complete VARON experiment for the case of a two-hop optimization. The sequence number of user generated traffic is shown in the figure, together with the RSSI measured by the originator mobile router (please note that the RSSI shown is already processed following a weighted average as explained in Section 4.2). Looking at the sequence number allows us to understand the impact of the route used by VARON on the user traffic, while the RSSI value lets us explain the behavior of the protocol based on the signal level detected. At the beginning (left side of the figure), the communication between two mobile networks is using the home (default) route via the 3G network. At some point in time (around $t = 20s$), a mobile router decides that a route optimization is possible, based on the reception of a HoAA message with good RSSI (the first vertical line shown in the figure corresponds to the reception of the HoAA, with a corresponding RSSI of approximately -70 dBm). This triggers the VARON signaling sequence on the mobile router, which is successfully completed with an MNPBU message, depicted in Figure 10 through the second vertical line (as shown in the bottom close-up, this process took around $0.5s$). As can be seen by the sequence number plot, there was no perceptible interruption in the IP packet flow, as the traffic was being forwarded via the home route until the care-of route has been established. The optimized care-of route was used for more than $70s$, until the measured RSSI dropped below the configured threshold, triggering the CoRE (third vertical line) - CoRE ACK (fourth vertical line) signaling, which made the mobile routers revert to the home default route. This procedure took also less than $0.5s$, as shown in the upper close-up in the figure. Like in the previous care-of route setup procedure, there was no perceptible impact on the IP packet flow.

The conducted experimental performance assessment shows that VARON provides interesting advantages by enabling the use of the VANET instead of the cellular network, when available. The benefits of using VARON such as reduced delay, lower cost and increased bandwidth, are exacerbated by the results obtained from the field trials, where the 3G commercial solution was not as reliable as expected. This reinforces the idea of using the VANET, enabling the opportunistic use of wireless multi-hop inter-vehicular communications. Obtained results show that this optimization is feasible, as VARON was able to enable VANET routes for very high percentages (more than 60%) of the time in which multi-hop connectivity was possible between two given vehicles. Even though the time required to setup an optimized route is in the order of half to one second, this does not prevent VARON to provide optimization opportunities. It is also worth highlighting that improvements in the cellular connectivity characteristics (e.g., delay, reliability) will have a direct positive impact on the overall VARON performance. An interesting future work item is to evaluate this performance when LTE becomes widely available.

6. Lessons learned and issues not to be overlooked

This section enumerates the most important lessons learned from the implementation, deployment and evaluation of VARON, as well as lists several specific aspects that may go unnoticed and should not be overlooked. The exercise of bringing a vehicular communication protocol into reality has enriched our knowledge, directing our attention to very specific issues that may go unnoticed otherwise. We think that these findings are very relevant for the design of successful vehicular networking solutions. We divide this section into two different groups: lessons learned and issues not to be overlooked.

6.1. Lessons learned

6.1.1. WLAN driver specifics

The design of link layer mechanisms has to take into account the specific hardware and WLAN driver behavior. For instance, the different wireless drivers show different results under the same testing conditions. Node discovery and synchronization in ad-hoc mode are specially critical. In our experiments, we adapted our implementation to two different wireless driver implementations, namely *madwifi* and *ath5k* as we noticed that the support of ad-hoc mode is quite different from one respect to the other. We found that the implementation of the ad-hoc (or Independent Basic Service Set, IBSS) mode in the ath5k driver is more consistent. We experienced some issues when using madwifi, especially during the initial joining process. While using ath5k this process was immediately performed, allowing a fast start of the data exchange, with madwifi the joining process was more difficult. The interfaces were not able to quickly communicate to each other, and sometimes the communication was not possible at all, requiring even a driver reset. We started our development with madwifi but we switched to ath5k to overcome this problem.

6.1.2. Multi-hop vs single-hop in the vehicular environment

Vehicular communication protocols may use single- or multi-hop communications. This actually makes a significant difference, because in addition to the obvious routing considerations, there are practical implications. For example, vehicular protocols quite often rely on broadcast or flooding mechanisms, which do not perform in practice as expected from simulation results (our experiments show low delivery rates of broadcast packets). Special care should be taken when designing a vehicular communication protocol to ensure that signaling messages meant to traverse a multi-hop network actually reach their destination, and if that is not the case, that there are mechanisms in place to detect it and react accordingly.

Additionally, the implementation of the wireless ad-hoc mode is typically much less developed and debugged as compared to the infrastructure one, which is more widely used. Besides, the differences in ad-hoc mode support and capabilities among the different available hardware and drivers are more disparate than for the case of infrastructure support.

6.1.3. Higher dynamism

Many vehicular protocols re-use concepts and solutions from other multi-hop ad-hoc communication scenarios. As already highlighted, the vehicular environment is very dynamic and suffers from severe radio conditions. This requires a more careful design of the VANET protocols, especially in terms of robustness and redundancy of the signaling. As an example, the reachability of a certain node via a multi-hop VANET route at a given moment does not guarantee that this connectivity will still be there some time after it was tested. Therefore, additional mechanisms should be designed to continuously evaluate this connectivity and detect potential disruptions or situations that indicate a high probability of imminent disconnection.

6.1.4. Cellular networks performance

3G networks are commonly assumed to provide always-on connectivity. Even some proposals make use of 3G networks for the delivery of critical safety messages, instead of using the connectivity provided by the VANET. However, our experiments revealed that 3G may not always be as reliable and stable as expected, but the opposite: measured delay showed great variability, bandwidth fluctuated and was often low and there were non-negligible packet loses. This behavior of course depends on the mobile operator and the location, but we believe that roads are particularly prone to suffer from this because operators do not dimension their networks with the goal of providing full 3G data coverage in roads yet. This is likely to change in the future, once mobile data access from vehicles becomes more popular and also when the new cellular technologies, such as LTE, get deployed.

6.2. Issues that should not be overlooked

6.2.1. When and where tests should be performed

The location and time for the tests are key elements in the deployment of a real vehicular prototype. The location has to be selected taking into account the safety of the people conducting the tests and the people on the road. As repeatability of the tests is critical in order to achieve statistically meaningful results, the test vehicles may not follow exactly the same behavior of a regular vehicle. Similarly, the time of the day has to be carefully chosen, so the impact of undesired external conditions is minimized. Some of these conditions might not be under the control of the people doing the experiments, as for example the quality of the 3G connectivity, which in our case experienced non negligible variations throughout the day.

6.2.2. Selection of the proper hardware

Commercial off-the-shelf (COTS) devices are often chosen to deploy test prototypes because they are conveniently handy and inexpensive. However, due to their limited capabilities, they may not be appropriate to every experimental deployment. In our case, COTS devices turned out to be very convenient for the development and preliminary testing phases, but they turned out not to be adequate for the tests on the road (see Section 5.1). This was caused by two main reasons: management of the additional 3G interface, and need of a VPN connection. Moreover, while running on-field tests, it is essential to be able to debug and reset the software and hardware easily, since resources are limited and time on the field is a very valuable resource. Therefore, it is advisable to take all these factors into consideration in advance, and select the equipment that better resembles the final product, but also considering the need of facilities for testing and debugging.

6.2.3. Power supply

When performing field trials, a long-live and reliable power source might not be available. This has to be taken into account when selecting the hardware, as not all the possible solutions are equally convenient. In our tests, we used both battery-powered uninterruptible power supply (UPS) units and AC/DC inverters plugged to the car to provide power to the nodes. The use of netbooks, which usually present a higher battery life, is also very convenient.

7. Conclusion

Vehicular communications have been extensively researched in the past, with a plethora of different solutions being proposed. However, there is still a lack of experimentation and deployment experience, with some remarkable exceptions of large scale testbed efforts.

With this work, we have tried to fulfill two ambitious goals: *i)* to experimentally validate and evaluate a vehicular communication protocol, VARON, which was initially proposed and extensively simulated in 2008; and *ii)* to report on the many insights and lessons learned throughout the process of fully implementing VARON and deploying it in a real scenario.

The validation and experimental analysis of VARON was conducted incrementally, starting from experiments performed on a controlled laboratory environment, and then moving to tests performed with up to 3 vehicles. While this is a reduced number of vehicles, we argue that the obtained results, and more importantly, the knowledge acquired from prototyping VARON in this scenario, would also apply to larger testbeds, as the issues we encountered have to be faced by each communicating vehicle individually. In terms of performance, VARON has shown to be feasible, enabling the use of opportunistically set-up VANET routes between two communicating vehicles, that would have to make use of a cellular network connection otherwise.

While prototyping and experimenting VARON in real vehicles, we had to face and tackle several challenging issues, resulting in some modifications and enhancements to the original VARON design. Additionally, we learned some interesting lessons in this process, that we consider helpful for the vehicular research community, and that we have tried to summarize in this article.

Acknowledgment

The authors would like to thank P. Gallego, P. Salvador and F. Giust for their help during the testing phase.

References

[1] C. J. Bernardos, I. Soto, M. Calderon, F. Boavida, A. Azcorra, Varon: Vehicular ad hoc route optimisation for nemo, Comput. Commun. 30 (2007) 1765–1784.

[2] V. Devarapalli, R. Wakikawa, A. Petrescu, P. Thubert, Network Mobility (NEMO) Basic Support Protocol, RFC 3963 (Proposed Standard), 2005.

[3] C. Perkins, D. Johnson, J. Arkko, Mobility Support in IPv6, RFC 6275, 2011.

[4] K. Sanzgiri, D. LaFlamme, B. Dahill, B. Levine, C. Shields, E. Belding-Royer, Authenticated Routing for Ad hoc Networks, Selected Areas in Communications, IEEE Journal on 23 (2005) 598–610.

[5] IEEE, IEEE Standard for Information technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007) (2012) 1–2793.

[6] P. Lutterotti, G. Pau, D. Jiang, M. Gerla, L. Delgrossi, C-VeT, the UCLA Vehicular Testbed: An open platform for vehicular networking and urban sensing, in: International Conference on Wireless Access for Vehicular Environments (WAVE 2008).

[7] R. Mahajan, J. Zahorjan, B. Zill, Understanding wifi-based connectivity from moving vehicles, in: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, IMC '07, ACM, New York, NY, USA, 2007, pp. 321–326.

[8] H. Soroush, N. Banerjee, A. Balasubramanian, M. D. Corner, B. N. Levine, B. Lynn, DOME: a diverse outdoor mobile testbed, in: Proceedings of the 1st ACM International Workshop on Hot Topics of Planet-Scale Mobility Measurements, HotPlanet '09, ACM, 2009, pp. 2:1–2:6.

[9] H. Soroush, N. Banerjee, M. Corner, B. Levine, B. Lynn, A Retrospective Look at the UMass DOME Mobile Testbed, SIGMOBILE Mob. Comput. Commun. Rev. 15 (2012) 2–15.

[10] M. Gerla, J.-T. Weng, E. Giordano, G. Pau, Vehicular testbeds - Validating models and protocols before large scale deployment, in: Computing, Networking and Communications (ICNC), 2012 International Conference on, pp. 665–669.

[11] P. Deshpande, X. Hou, S. R. Das, Performance comparison of 3G and metro-scale WiFi for vehicular network access, in: Proceedings of the 10th ACM SIGCOMM conference on Internet measurement, IMC '10, ACM, New York, NY, USA, 2010, pp. 301–307.

[12] A. Giannoulis, M. Fiore, E. W. Knightly, Supporting vehicular mobility in urban multi-hop wireless networks, in: Proceedings of the 6th international conference on Mobile systems, applications, and services, MobiSys '08, ACM, New York, NY, USA, 2008, pp. 54–66.

[13] J. Eriksson, H. Balakrishnan, S. Madden, Cabernet: vehicular content delivery using WiFi, in: Proceedings of the 14th ACM international conference on Mobile computing and networking, MobiCom '08, ACM, New York, NY, USA, 2008, pp. 199–210.

[14] F. Martelli, M. Elena Renda, G. Resta, P. Santi, A measurement-based study of beaconing performance in IEEE 802.11p vehicular networks, in: INFOCOM, 2012 Proceedings IEEE, pp. 1503–1511.

[15] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, J. Zahorjan, Interactive WiFi connectivity for moving vehicles, in: Proceedings of the ACM SIGCOMM 2008 conference on Data communication, SIGCOMM '08, ACM, New York, NY, USA, 2008, pp. 427–438.

[16] A. Balasubramanian, R. Mahajan, A. Venkataramani, Augmenting mobile 3G using WiFi, in: MobiSys '10: Proceedings of the 8th international conference on Mobile systems, applications, and services, ACM, 2010, pp. 209–222.

[17] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, S. Madden, CarTel: a distributed mobile sensor computing system, in: Proceedings of the 4th international conference on Embedded networked sensor systems, SenSys '06, ACM, New York, NY, USA, 2006, pp. 125–138.

[18] P. Shankar, T. Nadeem, J. Rosca, L. Iftode, CARS: Context Aware Rate Selection for Vehicular Networks, in: ICNP.

[19] P. Deshpande, A. Kashyap, C. Sung, S. R. Das, Predictive methods for improved vehicular WiFi access, in: Proceedings of the 7th international conference on Mobile systems, applications, and services, MobiSys '09, ACM, 2009, pp. 263–276.

[20] V. Shivaldova, C. Mecklenbräuker, Real-world measurements-based evaluation of IEEE 802.11p system performance, in: 5th International Symposium on Wireless Vehicular Communications: WIVEC2013, Dresden, Germany.

[21] V. Shivaldova, T. Paulina, A. Paier, C. F. Mecklenbruker, Performance measurements of multi-hop communications in vehicular ad hoc networks., in: ICC, IEEE, 2012, pp. 7162–7166.

[22] G. Gavilanes Castillo, M. Reineri, D. Brevi, R. Scopigno, M. Gallo, M. Pannozzo, S. Bruni, D. Zemberlan, Comparative characterization

of four antennas for vanets by on-field measurements, in: IEEE International Conference on Microwaves, Communications, Antennas and Electronic Systems (COMCAS 2013), Tel Aviv, Israel.

[23] R. Meireles, M. Boban, P. Steenkiste, O. Tonguz, J. Barros, Experimental study on the impact of vehicular obstructions in VANETs, in: Vehicular Networking Conference (VNC), 2010 IEEE, pp. 338–345.

[24] T. Melia, A. de la Oliva, I. Soto, C. Bernardos, A. Vidal, Analysis of the effect of mobile terminal speed on WLAN/3G vertical handovers, in: Global Telecommunications Conference, 2006. GLOBECOM '06. IEEE, pp. 1 –6.

[25] C. Ameixieira, J. Matos, R. Moreira, A. Cardote, A. Oliveira, S. Sargento, An IEEE 802.11p/WAVE implementation with synchronous channel switching for seamless dual-channel access (poster)., in: VNC, IEEE, 2011, pp. 214–221.

[26] V. Naumov, R. Baumann, T. Gross, An evaluation of inter-vehicle ad hoc networks based on realistic vehicular traces, in: Proceedings of the 7th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc '06, ACM, New York, NY, USA, 2006, pp. 108–119.

[27] C. Wewetzer, M. Caliskan, K. Meier, A. Luebke, Experimental evaluation of UMTS and Wireless LAN for inter-vehicle communication, in: Telecommunications, 2007. ITST '07. 7th International Conference on ITS, pp. 1–6.