

HarborNet: A Real-World Testbed for Vehicular Networks

Carlos Ameixieira, André Cardote, Filipe Neves, Rui Meireles, Susana Sargento, Luís Coelho, João Afonso, Bruno Areias, Eduardo Mota, Rui Costa, Ricardo Matos, João Barros

Abstract—We present a real-world testbed for research and development in vehicular networking that has been deployed successfully in the sea port of Leixões in Portugal. The testbed allows for cloud-based code deployment, remote network control and distributed data collection from moving container trucks, cranes, tow boats, patrol vessels and roadside units, thereby enabling a wide range of experiments and performance analyses. After describing the testbed architecture and its various modes of operation, we give concrete examples of its use and offer insights on how to build effective testbeds for wireless networking with moving vehicles.

Index Terms—Vehicular networks, testbeds, IEEE 802.11p, cloud-based network management, performance results

I. INTRODUCTION

After more than a decade of research and development, which culminated in the successful standardization of the IEEE 802.11p norm, vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication technologies have reached a stage of maturity in which real-world deployment is both feasible and desirable. Deployment is feasible, because specific radio interfaces that operate in the 5.9 GHz frequency band reserved for vehicles are already commercially available. It is further desirable, because the potential benefits of using V2V and V2I systems to increase vehicle safety, improve fleet management, reduce traffic jams and offer new services are by now widely recognized. A key challenge that remains is how to build a reliable wireless mesh network of moving vehicles that forward each other's data packets in a multi-hop fashion until they reach the Internet via the closest roadside unit (RSU) or access point.

Although numerous research contributions have addressed important issues related to mobile ad-hoc networks (MANET), of which vehicular ad-hoc networks (VANET) are a particular case, very few of them offer actual experimental results to support their claims. On the contrary, the large majority of existing references rely on computer simulations and theoretical models, whose accuracy and realism is yet to be determined.

Seeking to overcome the limitations of simulation-based research, we set out to build a real-world testbed for vehicular networks that could offer (a) high density of vehicles in a manageable space, (b) continuous availability and frequent

mobility (close to 24 hours a day), (c) fiber optical backbone for the roadside infrastructure, and (d) internet access for remote experimentation. By establishing a fruitful partnership with the trucking company responsible for all container transport inside the sea port of Leixões in Porto, Portugal, and engaging with the local port authority and port operator for backhaul support, we were able to achieve this goal. In particular, we currently operate 35 communication devices installed in container trucks, tow boats, patrol vessels, administration vehicles and roadside infrastructure, which together form a wireless mesh network of mobile nodes (also called HarborNet, see Figure 1) that spans the entire area of the port (about 1km²), and offers a unique facility for testing and experimenting with new communication protocols and security mechanisms for connected vehicles.

By sharing our insights, challenges and choices in designing and implementing a real-world testbed for vehicular mesh networking, we aim to motivate and support other research teams that wish to engage in real-life measurements and experimental research. Our main contributions are as follows:

- *Testbed architecture:* We give a detailed description of the onboard units and roadside units, which were custom made for the testbed, and explain how they were integrated in the existing infrastructure of the sea port. The testbed is further connected to a cloud-based system that supports data analytics and remote control of experiments.
- *Testbed operations:* We specify how code deployment and




Fig. 1: HarborNet at the Seaport of Leixões in Porto, Portugal.

Susana Sargento, André Cardote, Luís Coelho, Bruno Areias and João Afonso are or were with Instituto de Telecomunicações and the University of Aveiro. João Barros, Rui Meireles and Eduardo Mota are with Instituto de Telecomunicações and the Universidade do Porto. Carlos Ameixieira, André Cardote, Filipe Neves, Rui Costa, Ricardo Matos, Susana Sargento, and João Barros are with Veniam'Works Inc.

experiment control can be run efficiently from the cloud, both in real-time and in a delay-tolerant fashion.

- *Examples and field trials:* To illustrate the testing capabilities of HarborNet, we share the results of some of the experiments that were recently run there and the conclusions they offer.

The rest of the paper is organized as follows. Section II describes the devices and network infrastructure of the testbed. This is followed by a detailed account of its features, operations and functionalities, which is given in Section III. Practical examples and experimental results are offered in Section IV, followed by an overview in Section V of related work, highlighting the key differentiators between our testbed and other experimental facilities. The paper concludes in Section VI with lessons learned and recommendations.

II. TESTBED ARCHITECTURE

HarborNet is a vehicular mesh networking testbed, consisting of (i) on-board units (OBUs) installed in trucks, administration vehicles, tow boats and patrol vessels, (ii) roadside units connected to the optical-fiber backbone of the sea port, and (iii) cloud-based data and control systems. As shown in Figure 2, the vehicles are connected among each other via standard IEEE 802.11p/ WAVE links and can be reached from the cloud via the Internet. This technology is able to meet key requirements for these environments, because its range can go of up to 900m in line-of-sight, connections can be established in 10-20msec and vehicles can operate in a DSRC based mesh to overcome cellular dead zones between containers. Moreover, DSRC/WAVE works on an unlicensed band that is free of cost and reserved for intelligent transportation systems. Cellular backhaul is also available as a backup. WiFi is included in the OBUs and RSUs so that mobile employees can connect to the network using handheld devices. In the following, we explain the testbed components in greater detail.




Fig. 2: Testbed architecture.

A. On-Board Units and Road Side Units

Each truck is equipped with an OBU, which we denote as NetRider, with multiple wireless interfaces, which enable the vehicle to communicate both with other vehicles circulating inside the port and with RSUs that are integrated in the port infrastructure. OBUs and RSUs have a similar hardware, except for the antennas, which have higher gains in the RSUs. The NetRider OBU was designed for this testbed and includes the following elements:

- Single-Board Computer (SBC)
- Dedicated Short Range Communication (DSRC) wireless interface (IEEE 802.11p)
- Wi-Fi interface (IEEE 802.11a/b/g/n)
- 3G Interface
- GPS receiver
- Antennas for each device

The SBC contains the processing unit and is responsible for coordinating the various interfaces and access technologies. Moreover, it provides an in-vehicle WiFi hotspot for the occupants of the vehicles, cranes or vessels. A mini-PCI 802.11p compliant wireless interface is connected via one of the mini-PCI slots. This interface uses the Atheros AR5414 chipset, which supports the use of the *ath5k* driver. Table I gives an overview of the configuration. The frequency of operation is 5.850 GHz - 5.925 GHz, which has been reserved for intelligent transportation systems in various parts of the world including the EU and the US. A standard 802.11a/b/g wireless interface is connected to one of the USB ports of the mainboard to provide communication between the OBU and other user devices. This interface can also be used to connect the vehicles to any available Wi-Fi hotspot. A cellular interface is connected to another USB port, being used when required to control the OBU operations (or whenever no other connection type is available).

The GPS receiver is integrated with the IEEE 802.11p interface of the SBC to provide multi-channel synchronization. Synchronization to Universal Time Coordination (UTC) is mandatory for DSRC devices that switch between channels. The channel interval boundaries are derived from the GPS signals.

TABLE I: Configuration of the IEEE 802.11p interfaces

Parameter	Value
Channel	175
Center frequency (f)	5.875 GHz
Bandwidth	10 MHz
Setup TxPower	23 dBm / 18 dBm
Measured TxPower	14.58 dBm / 12.51 dBm
Receiver sensitivity	-95.2 dBm
Antenna Gain	2 dBi

The OBUs are running a tailored Linux distribution based on Buildroot [?]. The kernel was customized to include new features such as clock synchronization, as required by IEEE 802.11p. As Linux Wireless [?] does not provide support for the IEEE 802.11p / WAVE norm, the *ath5k* driver was modified to accommodate that norm within the AR5414A-B2B Atheros chipset. The driver was further extended to meet

the requirements of IEEE 802.11p/WAVE [?]. Synchronous channel switching makes it possible to switch seamlessly between the Control Channel (CCH) and a Service Channel (SCH) every 50 ms. This way, the OBU can listen to two wireless channels at the same time with one single radio. This allows it to receive emergency communications in the CCH, even though the SCH may be overloaded with traffic. The WAVE Short Messaging Protocol enables the exchange of short messages between nodes at the MAC layer, offering a fast way for vehicular nodes to safety-critical information.

B. Network Infrastructure

Once inside the port, clusters of vehicles are able to setup their own network autonomously. Whenever one of the nodes is within the range of a roadside unit, the vehicle cluster can connect to the Internet via the fixed infrastructure. In other words, each truck can access the Internet either directly via an IEEE 802.11p/WAVE connection to an RSU or via multi-hop communications with neighboring trucks. The latter can be viewed as an overlay network. Cellular communication is only used when real-time communication is required and no path to an RSU is available. Inside the trucks (and also inside boats, vessels and administration vehicles), each OBU disseminates an IEEE 802.11g network and functions as a mobile hotspot, thereby enabling the driver and the passengers to access the Internet via the vehicular mesh network or with cellular backhaul.

Because of the high mobility of harbor vehicles, their connectivity is hard to maintain. From our experiments, we concluded that the major factors that affect connectivity are the vehicle speed and heading, the number of hops to the infrastructure, the distance to the nodes, and the received signal strength (RSSI). To ensure that vehicles are always connected to the cloud systems for data analytics and remote control, each OBU runs an advanced connection manager [?], which was developed by us and takes into account the aforementioned factors. The connection manager weights them according to their network relevance and decides at each moment, location and network configuration the default wireless interface over which the next packets should be sent. This new connection manager exploits the characteristics of moving vehicles, such as speed and direction of nodes, number of hops to the infrastructure, and the dynamics of vehicular networks based on IEEE 802.11p links. The testbed currently supports a number of routing protocols, including GPSR, LASP, BATMAN and Babel, whose operation can be easily integrated with the connection manager.

C. Cloud Systems for Data Analytics and Remote Control

Data gathered by the vehicular networking testbed is sent to a cloud-based backend system for further processing and analysis. The backend system also deploys the software code for each experiment remotely and controls the experiments in a real-time or delay-tolerant fashion, as explained in Section III. Due to the dynamics of the vehicular mesh network and its constantly changing conditions (location of the vehicles, active connections, routes within clusters, etc.), experiments

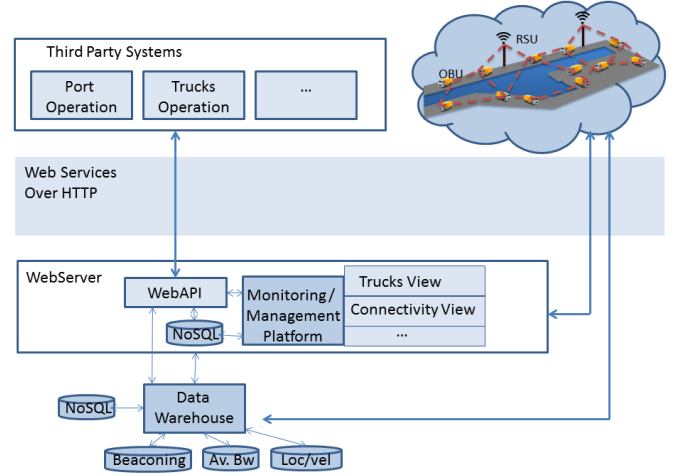


Fig. 3: System architecture.

generate about 8 Mbytes of logging information data per vehicle in each day, which can best be handled with an efficient cloud-based data management system. The architecture of our data management system is depicted in Figure 3. To speed up the data processing we keep the data stored in a central repository (data warehouse) and opted for noSQL [?] databases. Knowledge and information extracted from the data sets is then made available to the researchers via a Web Server, which was implemented with a Web API. We further implemented a number of monitoring applications (or views), that allows us to track the location and movement of the vehicles and assess their connectivity, among other parameters. Open data formats and data accessibility through WebServices and HTTP interfaces allow third parties to leverage the data for other applications. Case in point, the trucking company is currently using data from the testbed to better coordinate the efforts of the truck drivers and reduce both fuel consumption and CO₂ emissions.

III. TESTBED OPERATIONS

Building on the components described in the previous section, we will now address the features, operations and functionalities of the testbed, with special emphasis on the key tasks of deploying software code, controlling experiments and running measurements from the cloud-based backend system, both in real-time and in a delay-tolerant fashion.

A. Control of Experiments

To be able to deploy, monitor and manage network experiments remotely, we modified the open-source testbed controller and manager OMF [?]. From the available testbed controllers, OMF was considered the most appropriate, because it gives a simpler and more scalable way to describe experiments. Our version allows us to control the testbed nodes via IEEE 802.11p links and/or the cellular backhaul. This is made possible by several extensions, one of which addresses the support of dynamic IP addresses. Since vehicles use the connection manager to switch between different wireless

technologies, the IP address may change every time a new connection is established. A new service we added to the server side of OMF enables it to know the IP address of each node at all times, which is a key requirement for the server to be able to control every node.

Another extension is related to fail-safe recovery. This mechanism becomes necessary because nodes can fail and connections can be terminated, be it an IEEE 802.11p link or a cellular connection. To ensure that a node can recover from an unexpected state, we set up a timer that triggers an automatic reboot of the node whenever a prescribed event occurs. The trigger events can be configured through the daemon configuration file. In the current implementation, the node reboots every time it freezes or when the connection to the server is lost for more than a prescribed time. The latter is detected by means of a daemon that sends periodic ICMP messages and listens to the server replies.

To ensure further that nodes recover from failures, even when the failure is caused by one of the experiments they are running, we opted for a dual-boot system. The solution relies on two operating system partitions on each node: the first one is a base system that is never affected by any software operating on an experiment; the second one is the partition in which the software images of the experiment are decompressed after being downloaded from the cloud. When performing an experiment, each node boots on the first partition for normal operation and uses the second one to perform the experiment. A disk image upload is performed via a Secure SHell (SSH) connection to a node or a set of nodes. After the upload concludes, the node unpacks the image file to the second partition, copies the specific node configurations and reboots in the second disk partition. The total time for this process is about 15 minutes.

B. Delay-Tolerant Operation

Since vehicles are constantly moving, the communication links among them and to the RSUs are generally unstable and highly dependent on the relative locations of the mobile nodes and the RSUs. It is therefore reasonable to offer delay-tolerant support for applications that do not require real-time operation. As a first step towards this end, we implemented the network services that are required to download and upload data to and from the vehicles in a mesh network with unstable links. Typical uses include (a) downloading an image file for an experiment with a specific setup for the OBUs and their protocols, and (b) uploading log files with the data collected during the experiments. The delay-tolerant network services are implemented in the application layer, where they can control the data exchange with the cloud. However, the aforementioned services operate closely with the connection manager at the lower layers of the protocol stack, thereby ensuring that the downloading and uploading processes continue via the cellular backhaul if the other wireless interfaces are not available after a timeout event.

In the upload scenario, the application keeps a queue of all the files to be uploaded. These are sorted out by priority level. When the OBU is connected to the infrastructure (directly

or through multi-hop links), the application receives a signal from the connection manager and starts sending the data. If the connection is lost, the application receives another signal from the connection manager. The upload then stops, and the file remains in the queue to be resumed later in time. If the timeout of a file is reached before the transmission is complete, the remainder of the file is sent via the cellular backhaul.

Every transmission from any of the wireless interfaces is recorded for further analysis. This allows us to make very precise statements on how often the OBUs use IEEE 802.11p, WiFi or cellular, how much traffic is sent over each of the wireless interfaces, and how much load is placed on each of the RSUs, among other network metrics that are monitored on a daily basis.

C. Measurements and Tests

To assess the network performance, we implemented several measurement mechanisms, which span from the characteristics of UDP and TCP connections over vehicular networks to less intrusive methods to determine the available bandwidth, such as WBest [?]. At the application level, our performance measurement applications are able to measure the Received Signal Strength Indicator (RSSI) and Packet Delivery Ratio (PDR) of a link between vehicles, the quality of a stream of video, and the bandwidth in continuous channel or alternate channel mode. As an example, we measure the available bandwidth of the V2V and V2I links by dispersing probe packets. The basic idea is to enable the mobile nodes to evaluate periodically how much bandwidth they can use when communicating with their neighbors. The nodes thus maintain a table with their one-hop neighbors, which is updated periodically from the beacons nodes exchange among themselves. For each link, we sample the following metrics: estimated capacity, available bandwidth, available bandwidth with losses, round-Trip Time (RTT), and jitter. The data is timestamped with a Global Positioning System (GPS) time reference and then transmitted to the cloud for further processing.

IV. EXAMPLES AND FIELD TRIALS

The examples in this section illustrate the type of results that can be obtained from real-life experiments run in Harboret. Our field trial contains 25 trucks, 2 boats, 2 vessels, 1 harbor administration vehicle and 5 road side units. The latter were placed at a distance between 300m and 500m from each other.

The first one is shown in Figure 4 and concerns the network coverage offered by the vehicular mesh network and three roadside units (with locations indicated by the red drop icons with a dark dot). Over the course of several days, we collected extensive connection data sets and parsed the events in which the paths from vehicles located in each of the squares shown in Figure 4 to the Internet gateway consisted of one direct hop or multiple hops. In the colored map of Figure 4, the green squares mark the locations that can only be reached by means of multi-hop communication. The red squares indicate the locations that are always out of range, thereby offering relevant insights on where to place new roadside units.

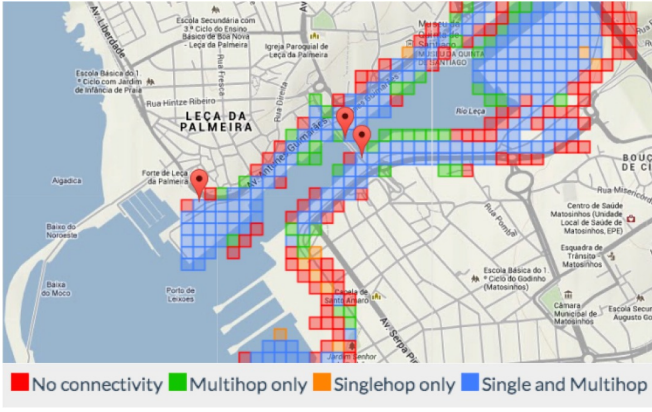


Fig. 4: Network coverage at the Seaport of Leixões.

The second example concerns the dynamics of the connections established by trucks via the vehicular mesh network. The histogram in Figure 5 shows the duration of the time intervals in which the trucks are in fact disconnected. Vehicles are disconnected when there is no network path from their location to a roadside unit. From our experiments, we conclude that in 70% of the time vehicles are disconnected for less than 1 min. Nearly 95% of the times vehicles are able to reconnect in less than 5 minutes. It follows that cellular support is not necessary for applications that tolerate a 5-minute delay (e.g. for updating the position of the trucks) and 95% reliability. Moreover, in one day of operation we can offload 60% of all traffic from cellular to IEEE 802.11p, which corresponds to about 8 GBytes of data per day. As a specific example, the trucks were able to upload large files (in the order of 5 MBytes) in delay-tolerant mode with mean rates exceeding 1Mb/sec (even when taking disconnected times into consideration). When connected to the IEEE 802.11p network, vehicles experience a data packet delay of less than 100 msec when the network is fully loaded with TCP connections. The handovers between IEEE 802.11p road side units are also performed in less than 100 msec, which enables seamless mobile communication within the premises of the harbor.




Fig. 5: Time intervals in which trucks are not connected to the Internet via the vehicular mesh network.

V. RELATED TESTBEDS

The goal of this section is to offer a non-exhaustive list of other existing testbeds for vehicular networking research that are complementary to our experimental setup:

- Cabernet [?] was a testbed deployed in Boston with 10 taxis, which addressed the viability of information sharing among drivers by using Wi-Fi (IEEE 802.11b/g) access points vehicles encountered opportunistically during everyday trips.
- CarTel [?] comprises 6 vehicles equipped with sensors and communications units that feature Wi-Fi (IEEE 802.11b/g) and Bluetooth. The testbed has been active in Boston and Seattle. This testbed provided an important insight on how to handle intermittent connectivity, and how feasible this kind of connectivity is to explore a class of non-interactive applications.
- C-VeT [?] is composed by 60 vehicles that circulate in the University of California at Los Angeles (UCLA) campus. A MobiMESH network has been deployed in the UCLA campus, through the installation of fixed nodes on the rooftops of the UCLA buildings, which covers the whole area of operation. IEEE 802.11g is used as the access technology, whereas IEEE 802.11a is used in the mesh core.
- SUVnet [?] is an emulated testbed in the city of Shanghai, China. 4000 taxis are equipped with GPS devices and report their positions in real-time to a central server. In this testbed, wireless communication is simulated on top of real-world mobility.
- SAFESPOT [?] is a testbed that was run for 4 years in six cities across Europe. It uses vehicles equipped with OBUs, RSUs and Traffic Centres (communicating through Wi-Fi) to centralize traffic information and forward safety-critical messages.
- DieselNet [?] is a testbed composed by 40 transit buses, 26 stationary mesh APs, thousands of organic APs (belonging to third-parties willing to participate), and 6 nomadic relay nodes. All nodes, except the organic APs, are equipped with Wi-Fi, 3G, General Packet Radio Service (GPRS) and 900MHz modems.

Neither of the aforementioned instances uses real-world implementations of V2V communications or IEEE 802.11p protocols. More recently, the Department of Transportation and the University of Michigan conducted a safety pilot deployment in Ann Arbor [?], Michigan. Here, more than 2800 vehicles were equipped with IEEE 802.11p-enabled OBUs and several infrastructure points were deployed in a given area of the city. The pilot trial used mostly private cars and was focused on driver assistance and safety applications. To the best of our knowledge, multi-hop communications and vehicular mesh networking were not yet addressed.

VI. LESSONS LEARNED AND FUTURE WORK

We presented a real-world testbed for vehicular networking that operates nearly 24 hours a day and allows for extensive testing of networking protocols for connected vehicles. The effort and investment required to set up such a network is very considerable. First, one must find owners of fleets who are willing to make their vehicles available for experimentation. Secondly, the communication devices available in the market can easily be inadequate or too expensive, which forces the

testbed promoters to invest in developing their own devices or at least customizing those they can afford. Finally, code development, network integration and cloud support must be well planned and executed, so that experiments can be run in the network with a small time investment on network updates and without compromising the current network operation, and data can be collected and analyzed with as little human intervention as possible, while delivering the desired outcomes and experimental results. Having gone through this entire process ourselves, we strongly believe that the quality and real-world relevance of the research and development projects that can be carried out on an experimental facility like the one described in this paper strongly justifies the effort and investment.

The testbed is currently in use not only for testing purposes, but also as the communication infrastructure of choice for supporting harbor operations [?]. More specifically, different members of the port ecosystem such as port authority, port operators and trucking companies use the vehicular network to share critical information and improve the coordination and safety of key operations involving vessels, vehicles and cranes.

The described infrastructure is ready to be reused in similar port environments and can be extended to larger fleets and networks, for example to support smart city services and applications.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of Sardão Transportes, Administração do Porto de Leixões (APDL), and Terminal de Contentores de Leixões (TCL). Part of this work was funded by Silicon Valley Community Foundation through Cisco University Research Program Fund under gift number 2011-90060 and by the Fundação para a Ciência e Tecnologia under various scholarship grants. Support was also received from the European Commission (EU FP7) under grant number 316296 (Future Cities).

BIOGRAPHIES

Carlos Ameixieira received the M.Sc. degree in Electronics and Telecommunications in December 2011 from the Department of Electronics, Telecommunications and Informatics of University of Aveiro. From July 2011 and December 2012 he joined the Institute of Telecommunications, located in University of Aveiro Campus, as a researcher. His research interests includes embedded systems, Linux kernel and vehicular communications. Since January 2013 he has joined the startup company Veniam Works.

André Cardote is currently an engineer at Veniam Works. He received his Ph.D. in Electronics and Telecommunications Engineering from the University of Aveiro in 2014, under the advisory of Prof. Susana Sargento (University of Aveiro) and Prof. Peter Steenkiste (Carnegie Mellon University). In 2011, he spent 6 months at Carnegie Mellon University as a visiting researcher. While his main expertise is computer networks, he is also interested in electronics and other related fields. He has participated in international research projects, one of them DRIVE-IN, a CMU—Portugal research project.

André has been part of the team that developed an inter-vehicle communication solution for Veniam Works from the beginning, and has since then been involved in the deployment of multiple vehicular networks in harbors and cities, involving fleets of taxis, buses and trucks. For more information, visit linkedin.com/in/acardote.

Filipe Neves studied Electronic and Telecommunications Engineering at University of Aveiro, Portugal, and received his M.Sc. degree in 2011. Between 2011 and 2012, he was a Researcher Grant Holder at Institute of Telecommunications - Aveiro, working on DRIVE-IN project (a CMU-Portugal research project). He has published several papers in the area of wireless networking and he works on ad-hoc networks, especially in vehicular networks. Recently, Filipe joined the Veniam Works team.

Rui Meireles is a dual-degree Computer Science PhD student at Carnegie Mellon University and University of Porto. He received an Engineering degree in Informatics from the University of Porto in 2006. His research interests revolve around vehicular network protocols and testbeds.

Susana Sargento Susana Sargento (<http://www.av.it.pt/ssargento>) received her PhD in 2003 in Electrical Engineering in the University of Aveiro (with a 7 months stay in Rice University in 2000 and 2001). She joined the Department of Computer Science of the University of Porto in September 2002, and is with the University of Aveiro and the Institute of Telecommunications since February 2004, where she is leading the Network Architectures and Protocols (NAP) group (<http://nap.av.it.pt>). She is also a Guest Faculty of the Department of Electrical and Computer Engineering from Carnegie Mellon University, USA, since August 2008, where she performed Faculty Exchange in 2010/2011. In March 2012, Susana has co-founded a vehicular networking company, Veniam Works, which builds a seamless low-cost vehicle-based internet infrastructure. Veniam Works won the ISCTE and MIT Portugal Program "Building Global Innovators" Ventures Competition in 2013. Susana has been involved in several national and international projects, taking leadership of several activities in the projects. She has been TPC-Chair and organizing several international conferences, has also been a reviewer of numerous international conferences and journals, such as IEEE Wireless Communications, IEEE Networks, IEEE Communications. Her main research interests are in the areas of Next Generation and Future Networks, more specifically QoS, routing, mobility, self- and cognitive networks. She regularly acts as an Expert for European Research Programmes.

Luís Coelho studied Electronic and Telecommunications Engineering at University of Aveiro, Portugal, and received his M.Sc. degree in 2011. Between 2011 and 2012, he was a Researcher Grant Holder at Institute of Telecommunications Aveiro, working on DRIVE-IN project (a CMU-Portugal research project). Since January 2013 until September 2013 he was a Researcher Grant Holder at Faculdade de Engenharia Universidade do Porto (FEUP), working on Project Future Cities. Since October 2013, he started working at Truphone.

João Afonso finished his Master Degree in Computers and Telematics Engineering In July 2012, from the Department

of Electronics, Telecommunications and Informatics of University of Aveiro. From April 2012 to February 2013, Joao Afonso worked as a researcher in Institute of Telecommunications, located in the University of Aveiro, on the monitoring and management of a vehicular network testbed, which was also the theme of his Master Thesis. His research interests included embedded systems, testbed management software and vehicular communications.

Bruno Areias is a student of Information Systems and Technologies in the Department of Electronics, Telecommunications and Informatics of University of Aveiro. Since September 2011, he joined Institute of Telecommunications in University of Aveiro, as a researcher. His research interests include Databases, Web Services, Network Monitoring and Management Systems.

Eduardo Mota concluded in 2011 his five-year Integrated M.Sc on electrical and computer engineering at the University of Porto, and in 2013 he attended the doctoral programme in the same field. Since 2012 he is a M.Sc researcher associated with the Institute of Telecommunications (Porto, Portugal), involved in several projects such as I-City, FutureCities and PredictEV. His main research interests are related to vehicular networks, network management solutions, network coding and cyber-security.

Rui Costa received the M.Sc degree in networks and communications engineering from Técnico Lisboa in 2013. Currently he is working at Veniam Works as Systems Engineer developing solutions for vehicular ad-hoc wireless communication environments based on the cloud, security mechanisms for end-to-end encryption over unreliable channels and in-vehicle applications APIs. He is also the founder and coordinator of the international project, IEEE Academic.

Ricardo Matos received his Ph.D. in Electrical Engineering by the University of Aveiro in January 2013. He is a member of Institute of Telecommunications since September 2008, belonging to the Network Architectures and Protocols (NAP) group (<http://nap.av.it.pt>). In January 2013, he joined a vehicular networking company, Veniam Works, which builds a seamless low-cost vehicle-based internet infrastructure. He has been involved in several national projects, as well as in the FP6 European Project WEIRD, and FP7 European Project EuroNF. His current research interests are related with vehicular networking, M2M communications, future Internet, virtualization techniques, wireless mesh networking, self-management, and context-awareness.

João Barros is an Associate Professor of Electrical and Computer Engineering (ECE) at the University of Porto and Founding Director of the Institute for Telecommunications (IT) in Porto, Portugal. He also teaches at the Porto Business School and co-founded two tech startups, Streambolico and Veniam Works, commercializing wireless video and vehicular communication technologies, respectively. He was a Fulbright scholar twice and has held visiting appointments at MIT, Carnegie Mellon, Cornell and Stanford. Between 2009 and 2012, Dr. Barros served as National Director of the Carnegie Mellon Portugal Program. He has been Principal Investigator (PI) and Co-PI of numerous national, European and industry funded projects, co-authoring one book published by Cam-

bridge University Press and more than 150 research papers in the fields of networking, information theory and security. Dr. Barros has received the 2010 IEEE Communications Society Young Researcher Award for the Europe, Middle East and Africa region, the 2011 IEEE ComSoC and Information Theory Society Joint Paper Award, the 2012 BES National Innovation Award, the 2013 Building Global Innovators Grand Prize (ISCTE-IUL and MIT) and a state-wide best teaching award by the Bavarian State Ministry of Sciences, Research and the Arts. He received his undergraduate education in ECE from the Universidade do Porto (UP), Portugal and Universitaet Karlsruhe, Germany, a performing arts degree in flute from the Music Conservatory of Porto, and a Ph.D. degree from the Technische Universitaet Muenchen (TUM), Germany.