ADAPTIVE HELLO PROTOCOL FOR VEHICULAR NETWORKS

Nathalie Mitton[1], Yasir Saleem[2], Valeria Loscri[1], Christophe Bureau[3]
[1]Inria, France, [2]Aberystwyth University, UK, [3]#dot:dot, France

NOTE: Corresponding author: Yasir Saleem, yasir.saleem@aber.ac.uk

Abstract – *In vehicular networks, the update of car Firmware Over The Air (FOTA) is becoming a challenging issue and it mainly relies on topology discovery of neighbouring nodes. Topology discovery in mobile wireless networks is usually done by using HELLO messages. Due to mobility, topology changes occur frequently and must be quickly discovered to avoid routing failures. Since the optimal HELLO frequency depends on parameters that are subject to changes (* e.g.*, speed of nodes, density of nodes), it must be dynamically adjusted to obtain the best trade-off between the network load and the freshness of routing tables. Existing solutions assume random mobility, constant node density and average speed, which do not hold in vehicular networks because vehicles follow specific trajectory patterns (the roads) and density and speed evolve as a function of time (rush hour vs non-rush hour) and area (urban, rural, highway). In this paper, we first draw the specific features of a vehicular network at different times and spaces by analysing real datasets and then propose a dynamic neighbour discovery protocol, Vehicular Adaptive Neighbour discovery Protocol (VANP). VANPis a fully-distributed protocol that sends beacons at an optimal frequency without knowing it* a priori*. The objective is to reduce the frequency at which HELLO messages are sent to save bandwidth and energy while still preserving the quality of the neighbour discovery. Through extensive simulations run on real datasets, we show that the optimal HELLO frequency can be reached by maintaining a constant optimal turnover, independent from the speed of the nodes and by aiming at this turnover, nodes automatically use the optimal HELLO frequency. Results show that VANPallows the discovery of relevant neighbours by missing at most two neighbours over all scenarios and reducing the number of HELLO messages up to twice, hence saving bandwidth and energy.*

Keywords – Adaptive hello, beacons, hello interval, vehicular networks

1. INTRODUCTION

The update of car Firmware Over The Air (FOTA) is a growing trend that transforms software-related recalls into mere updates, thus considerably reducing the cost of recalls for car manufacturers, while increasing driver safety in quasi-real-time. However, the number and data volume of FOTA updates is expected to increase exponentially with the steep increase in the number of processors per vehicle, combined with the increase in the number of connected and electric vehicles that is projected to be 69 million new connected vehicles in Europe by 2025 [1]. Multiplied by the increasing cost of mobile data subscriptions, this leads to the paradox that by 2030 or even before that, it may be more costly to update vehicles than to move them if FOTA is only relying on mobile network communication (be it LTE/4G or 5G) [2]. In addition, a significant portion (15 to 20%) of the reduction in CO2 footprint brought by (connected) electric vehicles is ruined by the carbon impact of mobile networks and cloud infrastructure [3]. There is thus a need to find a credible alternative to mobile networks for FOTA, from the ones that occur on the parking of the factory to the ones taking place when cars are already on the road. One appealing solution resides in the use of through-car mesh networks: car crowds are “quasi-networks”, only missing a ubiquitous mesh organisation with a radio network providing sufficient throughput. Such mesh networks are collaborative in nature, and much more open to creative business models without subscriptions, hence with the potential to bypass all or part of the cost of data subscription that currently creates an obstacle for FOTA. In addition, a mesh network being intrinsically more efficient than a global LTE network when devices are in close proximity (e.g., vehicles on the road or connected street lighting) [4], one can also expect a significant reduction in the carbon footprint of FOTA using mesh networks rather than mobile-based networks. In such a mesh, Vehicle-to-Vehicle (V2V) communications are dynamically created when the vehicles are close enough to one another, and no pre-established infrastructure is required. Due to the path loss of radio communications, only close vehicles may directly communicate with each other. Long-distance communications require multi-hop routing, where packets are forwarded by multiple intermediate nodes. Many efficient localised routing and data forwarding protocols rely on a local *neighbourhood table* [5]. This table is created by each node u and contains the list of nodes with which u can directly communicate (its *neighbours*). It is generally maintained by letting nodes broadcast HELLO messages to their neighbours.

Moreover, neighbour discovery is required in other applications of wireless networks as well. For instance, establishing a multi-hop routing, identifying a set of common available channels in cognitive radio networks to enable communication [6], data offloading using V2V communications in vehicular networks [7] and establishing cluster-based routing [8, 9] to name a few.

The efficiency of table-driven protocols (*e.g.*, routing, clustering, activity scheduling) obviously relies on the accuracy of these tables. In a mobile environment, neighbourhood tables are subject to change, and in this case, the accuracy depends on the frequency of HELLO messages. If this frequency is too low, nodes may not be detected by their neighbours, leading deprecated neighbourhood tables, and protocol failures are likely to occur. On the other hand, if the frequency is too high, neighbourhood tables are up to date, but then energy and bandwidth are wasted to the detriment of data traffic.
Yet, determining the correct frequency is not obvious. The optimal value, providing a good trade-off between these behaviours, actually depends on dynamic characteristics, such as the speed of nodes, and should thus be dynamically updated. By optimal frequency, we mean the minimum frequency at which HELLO messages can be sent to ensure that nodes correctly detect each other. The most straightforward solution is to suppose that nodes know their speed, but this assumption requires dedicated hardware (*e.g.*, GPS) that may not be available.
In this paper, we first analyse the behaviour of a vehicular network in different environments (urban, highway) at different times (rush hour, non-rush hour) by running extensive real dataset-based experimentations. Thanks to these experimentations, we could observe among others that each node discovers a significant number of neighbours with which it does not stay in range long enough to initiate some data exchanges. In a vehicular network application, we can thus distinguish the "relevant neighbours" from the set of all neighbours. We define the "relevant neighbours" as the neighbours with which we remain in range of each other long enough to exchange data. By adapting and decreasing the HELLO frequency, we consequently may miss neighbours. What we claim in this study is that this is not mandatory to discover all of them, it is even desirable since we mainly need to discover "relevant neighbours".
These important observations have driven our study and the adaptation brought to Turnover-based Adaptive HELLO Protocol (TAP), an existing adaptive neighbour discovery protocol designed for general wireless networks [10]. Our goal is to set proper parameters in our HELLO frequency adaptation algorithm in order to reduce HELLO frequency when possible without jeopardising the efficient discovery of "relevant" neighbours. As a result, we propose VANP, a *vehicular adaptive neighbour discovery protocol*. The key features of VANPare as follows:

- It is localised, and thus topology changes have very limited impact.
- It is fully software-based, and does not require dedicated hardware (such as a GPS).
- It is independent from the communication stack (*e.g.*, routing, MAC).
- It is light on CPU usage as no on-node heavy computation is required.
- It specifically suits vehicular networks and adapts to all road environments.

The contributions of this paper can be summed up as follows:

- a road traffic analysis based on real datasets that shows that it is not useful to discover all neighbours every time,
- the VANPprotocol, an adaptation of the TAP protocol to the vehicular context,
- a complete simulation-based evaluation relying on real datasets showing that VANPallows the discovery of relevant neighbours by missing at most two neighbours over all scenarios and reducing the number of HELLO messages up to twice, hence saving bandwidth and energy.

The outline of this paper is as follows. After discussing related work in Section 2, Section 3 presents the motivation of this work. Section 4 provides a background of TAP. In Section 5, we describe our protocol and discuss a few implementation issues, mainly about the accurate estimation of the turnover. Section 6 provide some experimental results run over real datasets and presents the performance of VANP. We finally conclude and discuss open issues in Section 7.

2. RELATED WORK

HELLO protocols and neighbours discovery have been the focus of several existing studies since networks have existed. In this related work review, we mainly focus on neighbour discovery for vehicular networks [11, 12, 13].
The authors of [14] propose REMR, a position-based routing scheme for emergency messages which adapts the HELLO interval based on the neighbourhood density to minimise HELLO messages congestion. The HELLO interval is initialised to 0.03s. Before sending a HELLO message, a vehicle checks whether the number of its neighbours is greater than the average number of neighbours in the network (considered as a threshold value). If so, the network is considered as high-density and the vehicle adapts the new HELLO interval by adding the ratio of the number of vehicles at current and previous time steps to the HELLO interval. Otherwise, it sets the HELLO interval to 0.03s. The limitation of this work is that there is no means to reduce the HELLO interval except resetting it to 0.03s, which is not always pertinent and assumptions do not always hold.
Li [15] proposed the adaptation of HELLO intervals using a cooperative scheme for vehicular networks based on the acknowledgements of HELLO messages. The acknowledgements are used to identify HELLO messages loss. When a vehicle detects HELLO messages loss, it increases its HELLO frequency to retain the neighbourhood

information. The limitation of this work is that vehicles are required to know their neighbours for receiving acknowledgements and the medium is quickly saturated by multiple acknowledgements.

In [16], rather than adapting the HELLO interval, the authors investigate adapting the transmission power of HELLO messages based on the prediction error of vehicle position. The vehicles increase their transmission power of HELLO messages when prediction errors are high, while they reduce their transmission power on lower prediction errors. The simulations results showed a reduced channel busy time and a higher packet transmission rate compared to a fixed transmission power of HELLO messages but they pursue a different goal.

Lyu *et al.* [17] worked on an adaptation of HELLO intervals to avoid rear-end car crash by defining a danger coefficient which identifies a danger threat of each vehicle to be involved in a rear-end collision. Based on the danger coefficient, the authors proposed a distributed HELLO messages congestion control strategy to adapt the HELLO interval. Each vehicle calculates its danger coefficient and shares this with neighbouring vehicles through HELLO messages. The vehicles with a higher danger coefficient use a lower HELLO interval to allow taking immediate actions by rear vehicles in case of sudden braking to avoid collisions.

Li *et al.* [18] used mobility prediction to adapt the HELLO interval and to reduce channel occupancy due to the frequent broadcasting of HELLO messages. The authors claim that vehicles can predict the position of neighbouring vehicles. The vehicles use Kalman Filter to predict the future possible moving states and use kinematics law to predict the mobility of their neighbouring vehicles. The HELLO interval is adapted by comparing the prediction error to a deviation threshold. If a prediction error is higher than the deviation threshold, the HELLO messages are broadcast, otherwise the HELLO interval is reduced.

Nguyen and Jeong [19] investigated adaptive HELLO broadcast using mobility prediction for Guaranteed Tracking of Degree (GTK) K. The allowed tracking inaccuracy that does not affect the reliability of vehicle collision avoidance is called tracking threshold. Given a tracking threshold, GTK K is defined as the guarantee that a vehicle that receives at least one of the K HELLO messages from a sending vehicle, should be able to estimate the position of the sending vehicle within the accuracy of the tracking threshold. The main objective of the authors' scheme is to postpone the broadcasting of HELLO messages as long as the tracking inaccuracy is less than the tracking threshold. This saves the frequent broadcasting of HELLO messages.

In [20], the authors investigate rear-end collision and proposed BASE for the adaptation of HELLO intervals to save the radio resources without jeopardising the safety requirements. BASE adapts the HELLO interval of each vehicle by guaranteeing that a minimum safety distance should always be maintained between vehicles. The minimum safety distance is defined as the required distance to avoid rear-end collision caused by sudden braking by the vehicle in front. BASE is mainly focused on normal traffic conditions and therefore does not consider complicated conditions, such as crossroads and overtaking.

In [21], the authors worked on congestion control by investigating the adaptation of HELLO intervals and awareness control using non-cooperative game theory and proposed NORAC. NORAC assigns a HELLO interval to each vehicle fulfilling its needs by guaranteeing fairness among vehicles with similar needs. NORAC is distributed and non-cooperative. It guarantees fairness using the fair concept of Nash Equilibrium.

Unlike the above-mentioned works, our approach does not require heavy computing or predictions. It is localised and fully adaptive.

3. MOTIVATION: VEHICULAR NETWORKS STUDY

In order to characterise key features of a vehicular network, we have run a simulation-based study by analysing real traffic datasets. We have used the Simulation of Urban Mobility (SUMO) [22] road traffic simulator integrated with OMNeT++ [23], an event-driven network simulator to simulate the vehicular wireless communications and Veins [24], an open-source framework that provides the functionality of vehicular networks.

Various real-world road traffic scenarios of different cities have been made available by researchers[1] for SUMO [22] that can be imported to various network simulators (e.g., OMNeT++, NS-3). We conducted this analysis on different cities. In this paper, we only present the results for the city of Dublin over different scenarios but results are similar for different cities.

3.1 Scenarios description

In this paper, we mainly focus on the scenarios of rush and non-rush hours of Dublin city centre (referred as Ireland urban scenario hereafter) and Ireland National Road (referred as Ireland national highway hereafter). The rush hour (saturated traffic) is when the level of traffic is high and it starts getting congested. The non-rush hour (free flow traffic) is early morning traffic when the level of traffic is low.

3.1.1 Ireland urban scenario

The Ireland urban scenario is comprised of 5km×3.5km Dublin city centre having 435 signalised intersections. Dublin SCATS dataset[2], a traffic light control system, was used to generate the traffic by counting the vehicles using SCATS sensors every six minutes at 480 locations in Dublin city. It provides traffic for 24 hours. The traffic between 8am—9am and 4am—5am are the rush and non-rush hours traffic, respectively [25].

[1] https://sumo.dlr.de/docs/Data/Scenarios.html
[2] https://data.gov.ie/dataset/traffic-volumes

3.1.2 *Ireland national highway scenario*

The Ireland national highway is N7, a 3-lane two-way road of 17.1 km long with 11 on and off ramps each way. The loop sensors from an open dataset were used to generate the traffic and the Transport Infrastructure Ireland[3] published the data that includes 5-minute aggregated data for each lane and direction for several months. The traffic is averaged to create a workday traffic by excluding weekends and bank holidays. It provides traffic for 24 hours. The traffic between 5am—6am and 4am—5am are the rush and non-rush hours traffic, respectively [25].

3.2 On the pertinence of relevant neighbours

As previously discussed, we claim that it is not useful to detect all neighbours at every time but only *relevant* neighbours, i.e. the ones that allow a contact long enough for a data transmission. To illustrate this, Fig. 1 shows for each environment (urban rush hour, non-rush hour, highway rush hour, non-rush hour) the distribution of the average contact duration of connections between two vehicular neighbours. We can see for instance on this figure that in an urban environment at rush hour time (orange curve) that 50% of connections last for less than 12s and that only 20% of connections last more than $20s$. Applications in vehicular networks generally require the exchange of data and thus a stable contact during a minimum time. Vertical lines on the figure mark the limits for connections that last respectively more than 3, 5 or $10s$. If we define the *relevant* neighbours as the neighbours that remain in contact for at least x seconds, we can see that in an urban environment in rush hour, only 93% of neighbours are relevant for $x = 3$, 90% for $x = 5$ and that this amount drops to 76% for $x = 10$s.

This demonstrates that there is no use in designing a neighbour discovery protocol that has a constant complete view of the neighbourhood since part of the nodes will not be usable.

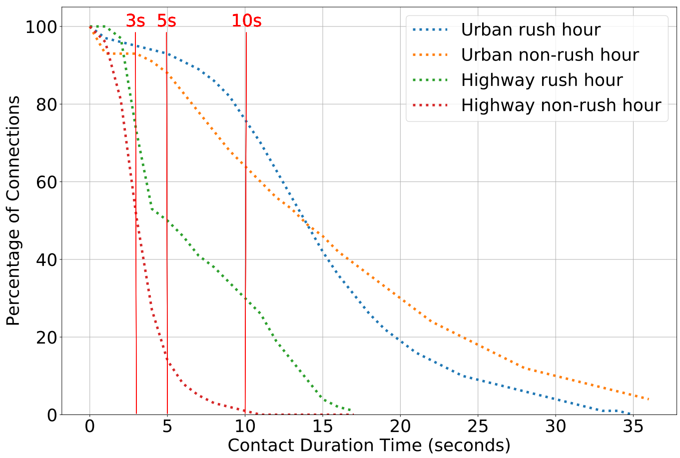

Fig. 1 – Distribution of contact time

[3]https://trafficdata.tii.ie/publicmultinodemap.asp

3.3 On the HELLO frequency adaptation

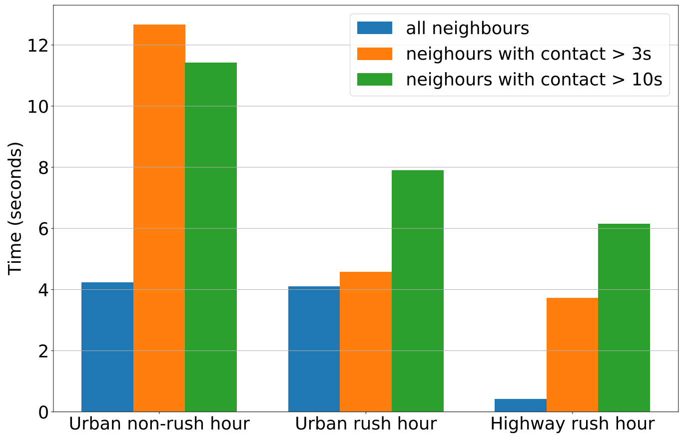

Fig. 2 – Average time before meeting a new relevant neighbour

Fig. 2 shows for each environment the average time before meeting a new relevant neighbour. As expected, it highlights that on average, a node needs much more time before meeting a relevant neighbour than any single neighbour. As such, this means that the delay between two HELLO messages can be increased. This motivates the need to provide an adaptive HELLO protocol specific to the vehicular context.

Table 1 sums up the notations used in this paper.

Table 1 – Notations

Notation	Description
r	Current observed turnover
r_{opt}	Optimal turnover
f_{HELLO}	Hello frequency ($f_{\text{HELLO}} = \frac{1}{d_{\text{HELLO}}}$)
d_{HELLO}	Delay between two Hello messages ($d_{\text{HELLO}} = \frac{1}{f_{\text{HELLO}}}$)
f_{opt}	Optimal Hello frequency
X	Size of history
N_t	Neighbourhood table at time t

4. BACKGROUND: THE TAP PROTOCOL

In this section, we provide a summary of the TAP protocol [10] that lays the basis of VANP. Let us assume that nodes send HELLO messages at the frequency f_{HELLO} (a HELLO message is sent every $d_{\text{HELLO}} = \frac{1}{f_{\text{HELLO}}}$). Whenever a node receives a HELLO message, it updates its neighbourhood table, thus generating a *turnover* r. Given a period of time Δt, we define the turnover $r_{\Delta t}$ to be the ratio between the number of new neighbours (*i.e.*, nodes that were not already neighbours Δt units of time earlier) and the current total number of neighbours. Obviously, the turnover depends on both Δt and f_{HELLO}. Without loss of generality, let us set $\Delta t = 1/f_{\text{HELLO}} = d_{\text{HELLO}}$, that is, the turnover is updated each time a new HELLO message is sent. As is, a high frequency induces a low turnover (nodes frequently update their tables and thus discover a few new neighbours each time), while a low frequency induces a high turnover. Because of that direct relationship between the two values, we can note that acting on one of them directly affects the other one. TAP relies on the idea that the optimal HELLO frequency f_{opt} generates an optimal turnover r_{opt}, and by keeping r close to r_{opt}, f_{HELLO} automatically tends toward f_{opt}. Nodes thus have to slightly adjust their frequency with regards to the observed turnover. Authors of [10] proved that while the optimal frequency depends on the speed of the node, the optimal turnover is fixed and does not depend on any such parameter. The optimal frequency indeed evolves accordingly to the speed (*i.e.*, it increases when the speed increases, and decreases when the speed decreases), but eventually generates a constant turnover r which can be computed offline. TAP is also efficient in static networks.

5. VEHICULAR ADAPTIVE NEIGHBOUR DISCOVERY PROTOCOL

As explained, VANP is inspired by TAP. However, TAP does not specify the way to adapt the frequency based on the observed turnover. Also, TAP has computed the optimal turnover r_{opt} in a general wireless network, assuming nodes are uniformly distributed in space and follow a random way point mobility model with fixed average speed, which is not the case in vehicular networks. In this section, we describe VANP and in particular how it adapts to the HELLO frequency. Then, we compute the optimal turnover to be observed in our environment.

5.1 Implementation

Obviously, accurately estimating the turnover is of prime importance. As we will show in the next section, the value of the optimal turnover may be quite low, and for low densities it can be difficult to effectively estimate the current turnover. For instance, if a node has only two neighbours at a given moment, the turnover is limited to three distinct values (zero if none of the neighbours are new, one if only one of them is new, or two otherwise), which is not precise enough to accurately adjust the current HELLO frequency. A solution to this problem is to take account of the number of new neighbours that appeared at every X previous observations, X being the size of the history. In this case, the value may be obtained by enumerating new neighbours between two successive tables:

$$r = \frac{d_{\text{HELLO}}}{t_X - t_0} \sum_{i=1}^{X} \frac{new_{t_i - t_{i-1}}}{|N_{t_i}|}, \qquad (1)$$

where $new_{t_i - t_{i-1}}$ is the number of new neighbours detected between times t_{i-1} and t_i, and N_{t_i} is the neighbourhood table at time t_i and $|N_{t_i}|$ its size. By doing this, it is possible to obtain a sufficient accuracy of the turnover estimation, as we will show in Section 6.

The adjustment of f_{HELLO} requires the use of an incremental update function based on the observed turnover r. A very simple series like the following one perfectly fits this purpose:

$$d_{\text{HELLO}} = \left\{ \begin{array}{ll} d_{\text{HELLO}} + \frac{d_{\text{HELLO}}}{\alpha} \times g(r) & \text{if } r \leq r_{\text{opt}}, \\ d_{\text{HELLO}} - \frac{d_{\text{HELLO}}}{\alpha} \times g(r) & \text{otherwise.} \end{array} \right\} \qquad (2)$$

Here, α is a constant controlling the speed of the convergence, and $g(r)$ is a function of the observed turnover that returns how much d_{HELLO} should be changed. To design this function, we should first note that the more r and r_{opt} are different from each other, the more likely the current frequency f_{HELLO} and the optimal frequency f_{opt} are different from each other. Thus, the higher the difference between r and r_{opt}, the more quickly f_{HELLO} should "move" towards f_{opt}. We define $g(x)$ as follows:

$$g(r) = \left\{ \begin{array}{ll} (\frac{r - r_{\text{opt}}}{r_{\text{opt}}})^2 & \text{if } r < 2 \times r_{\text{opt}}, \\ 1 & \text{otherwise.} \end{array} \right. \qquad (3)$$

When r and r_{opt} are only slightly different, then the normalised amplitude is close to 0, resulting in a slight variation of f_{HELLO}. Opposed to this, the higher the difference between r and r_{opt}, the higher the amplitude and the faster f_{HELLO} converges toward the optimal frequency f_{opt}. Any other function $g(r)$ such that $g(r_{\text{opt}}) = 0$ may of course be used. However, this one is very straightforward to implement, and a lookup table could actually be used to avoid overhead.

As a result, Algorithm 1 describes VANP.

Algorithm 1 Adaptive HELLO Protocol.

```
Initialise r_opt, d_HELLO, X;
i = 1;
while true do
    new = (1/i) Σ_{j=0}^{t_{i-1}} |N_{t_{j+1}} \ N_{t_j}|
    r = new × d_HELLO / (t_i − t_0)
    Compute g(r) using Eq. (3)
    Compute d_HELLO using Eq. (2)
    if i < X then i++
    end if
    Schedule HELLO message after d_HELLO period;
end while
```

5.2 Determining the optimal turnover r_{opt}

Now that the general design of our protocol has been presented, the principal issue that remains opened is: What is the optimal turnover r_{opt}? Finding the correct value of f_{opt} is indeed not trivial, as it depends on a lot of parameters. To deduce it, we rely on our real dataset analysis.

Fig. 3 shows the proportion of relevant neighbours over time for each scenario. In this figure, we considered as *relevant neighbours* the ones that remain in contact for more than 3 and 10s respectively. Globally we can observe that whatever the scenario and thus the density and speed of nodes, this proportion remains constant over time and is

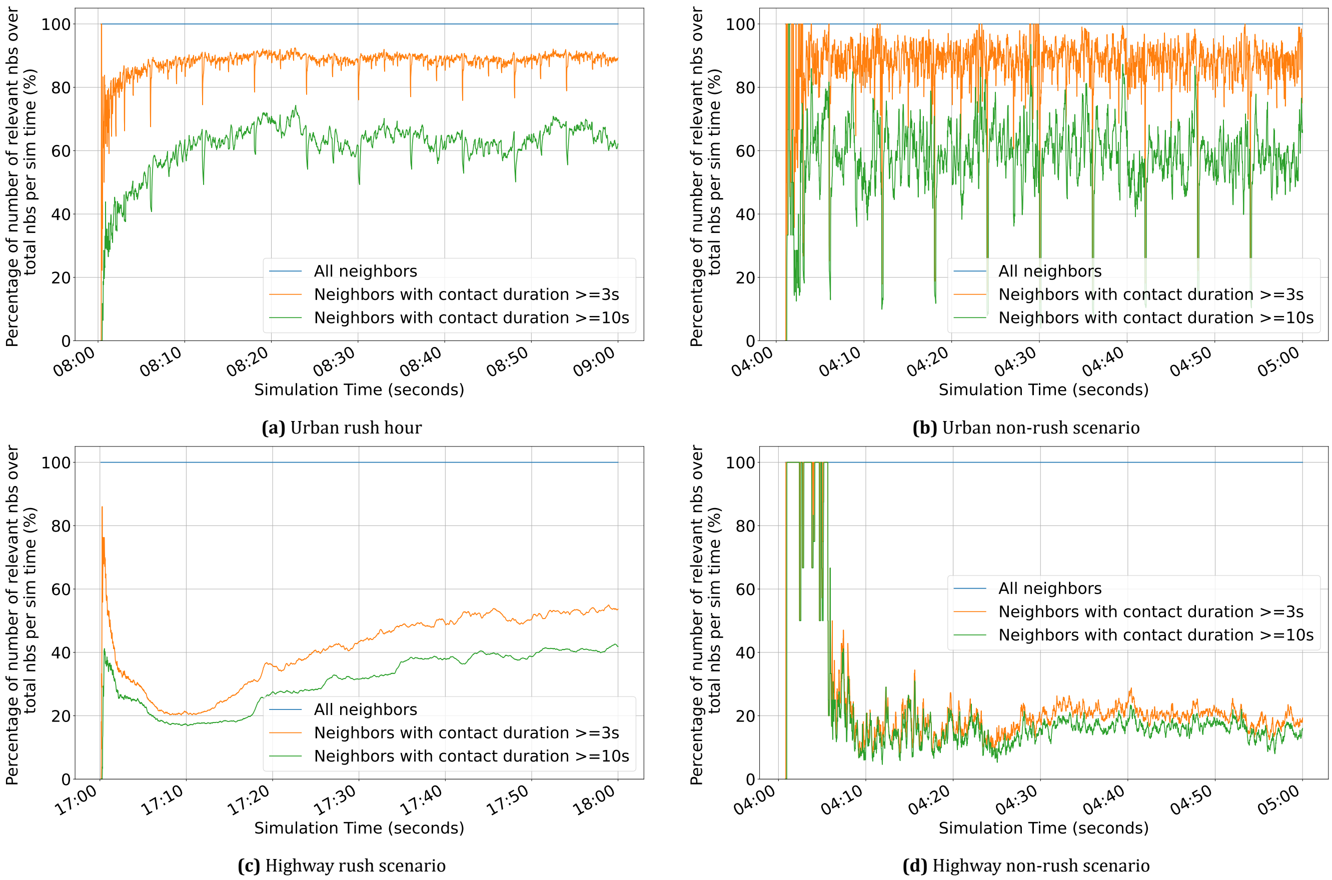

(a) Urban rush hour

(b) Urban non-rush scenario

(c) Highway rush scenario

(d) Highway non-rush scenario

Fig. 3 – Proportion of relevant neighbours for each scenario

Table 2 – r_{opt} values

	Urban		Highway	
	Rush hour	Non-rush hour	Rush hour	Non-rush hour
$> 3s$	0.8	0.8	0.5	0.2
$> 10s$	0.65	0.65	0.4	0.2

the same for different times in the urban scenario (about 85% of neighbours in contact for more than $3s$ and 60% in contact for more than $10s$). A third observation is that in highway scenarios where there are less intersections, the gap between the two categories of *relevant* neighbours is small (in rush hours, 50% of neighbours in contact for more than $3s$ against 40% for more than $10s$ and in non-rush hours, for all kind of relevant neighbours 20% of neighbours only). This proportion allows us to determine the optimal turnover r_{opt} value to be used in the HELLO adaptation process.

Based on these values, we retain the r_{opt} values depicted in Table 2.

6. EXPERIMENTAL RESULTS

In this section, we analyse first the behaviour of VANPand then its performances. We use the same set up and simulation environment as detailed in Section 3.

6.1 Protocol behaviour

Fig. 4 presents the behaviour of the delay between two HELLO messages for each non-rush hour scenario. We can observe that the delay between two HELLO messages dynamically change over time, as expected.

Fig. 5 shows the resulting average d_{HELLO} in every non-rush hour scenario at the end of the simulation. One can observe that for each case, $d_{\text{HELLO}} > 0.5s$ is the usual d_{HELLO} applied in communication protocols. This demonstrates that VANPsends less HELLO messages than classic protocols in urban environments and thus save energy and bandwidth. On the contrary, more packets are sent in high-speed environments (e.g. highway in non-rush hour) in order not to miss important neighbours. We can also observe that, as expected, more packets are sent to discover neighbours that remain in contact for more than $3s$ than for neighbours that remain in contact for more than $10s$ (delay is smaller) since more nodes have to be spotted.

We can also note that the size of history X is not anodyne since using $X = 10$ instead of $X = 5$ allows even more saving and the saving is even more important in high-speed scenarios as in a highway scenario. This is because a greater history size allows for smooth rapid changes.

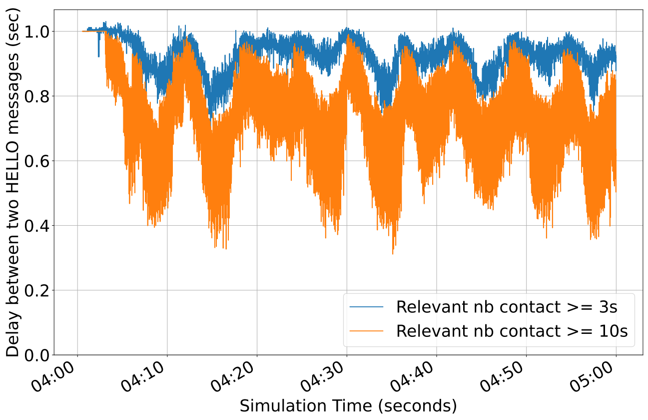

(a) Evolution of d_{HELLO} in non-rush hour urban scenario

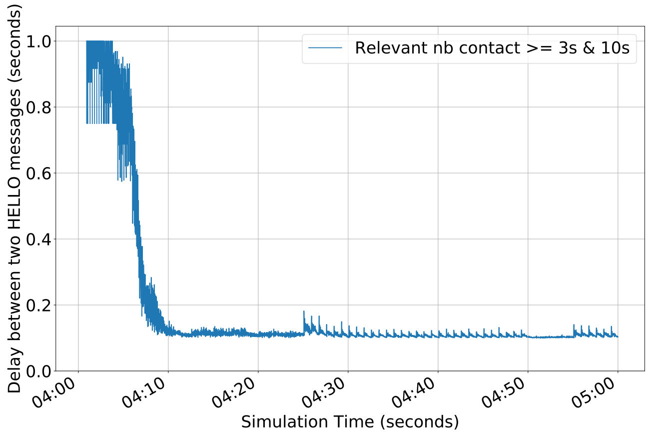

(b) Evolution of d_{HELLO} in non-rush hour highway scenario

Fig. 4 – VANPbehaviour

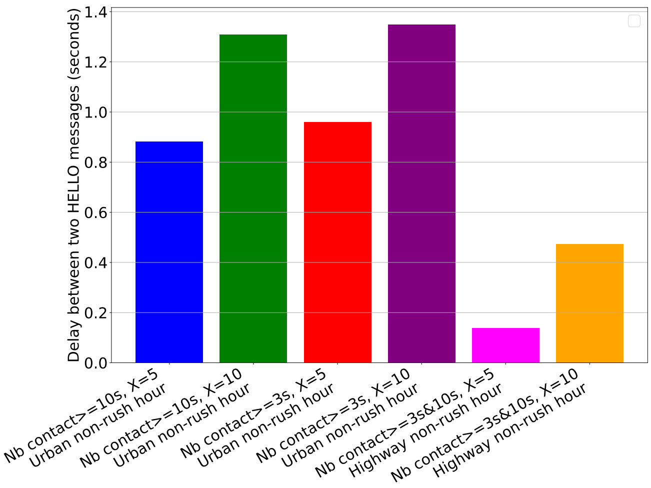

Fig. 5 – Average d_{HELLO} in non-rush hour scenarios

6.2 Protocol performances

Saving energy and bandwidth is paramount but not at the cost of neighbour tables accuracy. Therefore, Fig. 6 presents the amount of false positive relevant neighbours (Relevant neighbours in the neighbourhood table but not in communication range) in different scenarios for non-rush hour. In an urban scenario with a neighbour contact duration greater than 3s in Fig. 6(a), in the majority of cases, there are only 0.1 false positive neighbours. In a few rare cases, the number of false neighbours is 0.2 and 0.5 false positive neighbours in two cases only. The urban scenario with neighbour contact duration greater than 10s in Fig. 6(b) exhibits the similar behaviour as in the previous scenario with the exception of slightly more false positive neighbours having a value of 0.5 when the history size is 10 (i.e., $X = 10$). The highway scenario has 0.2-0.4 false positive neighbours in the first 30 minutes of non-rush hour, while less than 0.1 false positive neighbours in the last 30 minutes.

Whatever the scenarios, we can observe that VANPis very accurate since it never badly tags more than one node at a single time in the worst case and on average 0.2 nodes (so no bad tagging in most cases).
Similarly, Fig. 7 presents the amount of missed relevant neighbours detections (Relevant neighbours in communication range but not in the neighbourhood table) in different scenarios for non-rush hour.
As previously noted, we can observe the high performance of VANPsince it never misses more than two nodes at a single time in the worst case and on average 0.5 nodes (so no missed detection in most cases).

7. CONCLUSION

We have presented an adaptive HELLO protocol for vehicular networks that estimates the neighbourhood turnover and adjusts the HELLO frequency consequently. Besides the fact that our protocol is very straightforward to implement, it is especially well-tailored to standard mobile wireless networks since it does not rely on any specific hardware to aim at an optimal HELLO frequency. We have shown under realistic scenarios using real datasets that VANP is very accurate and allows substantial energy and bandwidth saving. Future work will focus on the automatic detection of the scenario (e.g. urban vs highway, rush hour vs non-rush hour) in order to dynamically adapt the optimal parameter setting to the network density and traffic model.

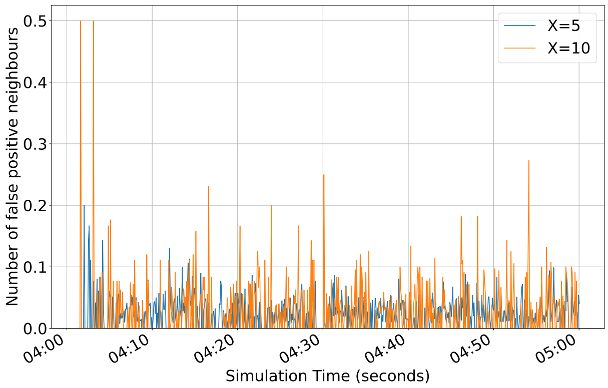

(a) Urban scenario - Neighbours with contact duration greater than 3s

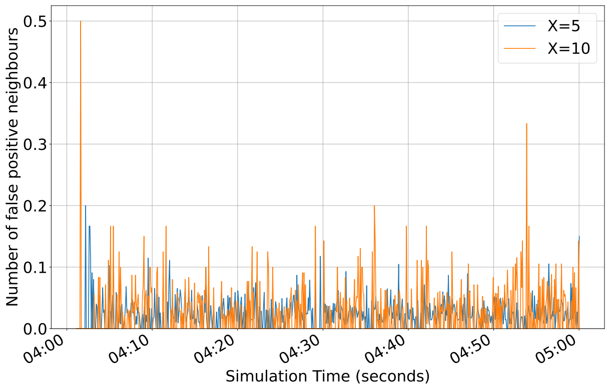

(b) Urban scenario - Neighbours with contact duration greater than 10s

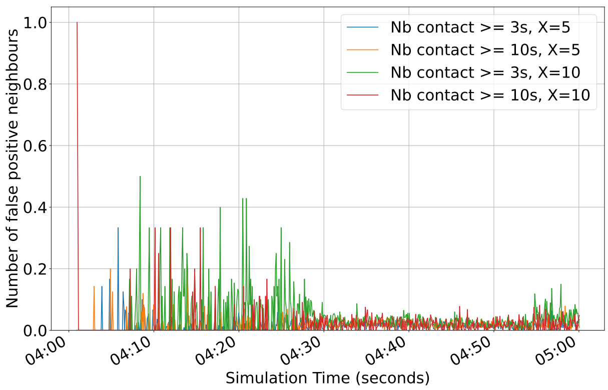

(c) Highway scenario

Fig. 6 – Number of false positive neighbours in different non-rush hour scenarios

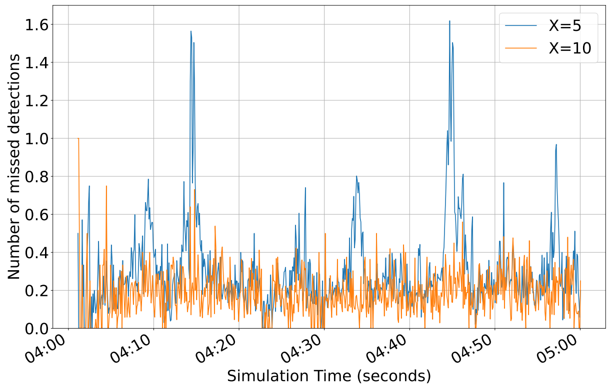

(a) Urban scenario - Neighbours with contact duration greater than 3s

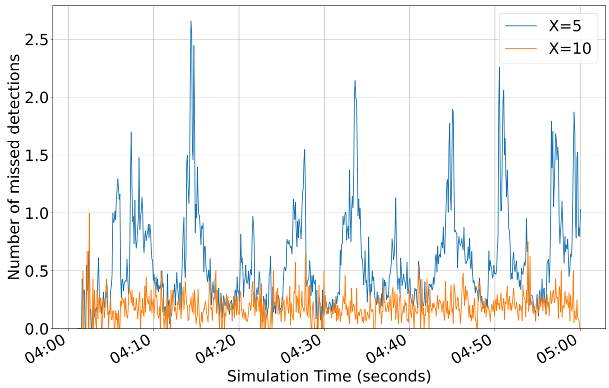

(b) Urban scenario - Neighbours with contact duration greater than 10s

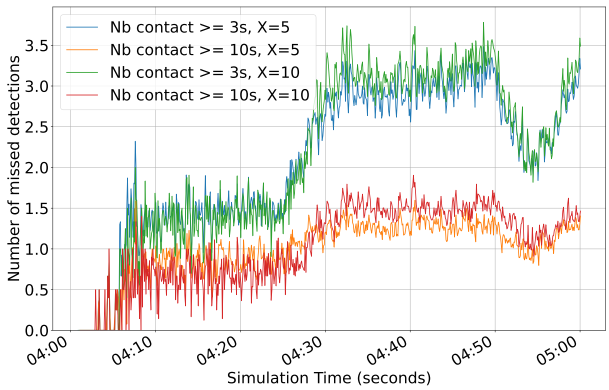

(c) Highway scenario

Fig. 7 – Number of missed relevant neighbours in different non-rush hour scenarios

REFERENCES

[1] *Connected Car Shipments for Europe Hit the Expressway*. Last accessed: December 2023. URL: https://www.counterpointresearch.com/insights/connected-car-shipments-for-europe-hit-expressway/.

[2] #dot:dot. *Car FOTA based on Mobile Networks: A Dead End?* https://beeyond.fr/s/4145071855. 2023.

[3] #dot:dot. *Estimating carbon emissions in mobile networks*. https://beeyond.fr/s/4145074127. 2023.

[4] DIGI. *Mesh Networking Vs. Cellular Technology for IoT Applications*. Last accessed: December 2023. URL: https://www.digi.com/resources/videos/mesh-networking-vs-cellular-technology-for-iot-ap.

[5] Sardar Muhammad Bilal, Carlos Jesus Bernardos, and Carmen Guerrero. "Position-based routing in vehicular networks: A survey". In: *Journal of Network and Computer Applications* 36.2 (2013), pp. 685–697.

[6] Athar Ali Khan, Mubashir Husain Rehmani, and Yasir Saleem. "Neighbor discovery in traditional wireless networks and cognitive radio networks: Basics, taxonomy, challenges and future research directions". In: *Journal of Network and Computer Applications* 52 (2015), pp. 173–190.

[7] Yasir Saleem, Nathalie Mitton, and Valeria Loscri. "DIVINE: Data offloading in vehicular networks with QoS provisioning". In: *Ad Hoc Networks* 123 (2021), p. 102665.

[8] Yasir Saleem, Kok-Lim Alvin Yau, Hafizal Mohamad, Nordin Ramli, and Mubashir Husain Rehmani. "SMART: A SpectruM-Aware ClusteR-based rouTing scheme for distributed cognitive radio networks". In: *Computer Networks* 91 (2015), pp. 196–224.

[9] Yasir Saleem, Kok-Lim Alvin Yau, Hafizal Mohamad, Nordin Ramli, Mubashir Husain Rehmani, and Qiang Ni. "Clustering and reinforcement-learning-based routing for cognitive radio networks". In: *IEEE Wireless Communications* 24.4 (2017), pp. 146–151.

[10] F. Ingelrest, N. Mitton, and D. Simplot-Ryl. "A Turnover based Adaptive HELLO Protocol for Mobile Ad Hoc and Sensor Networks". In: *IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*. 2007.

[11] Byungjo Kim, Inhye Kang, and Hyogon Kim. "Resolving the unfairness of distributed rate control in the IEEE WAVE safety messaging". In: *IEEE Transactions on Vehicular Technology* 63.5 (2014), pp. 2284–2297.

[12] Nader Chaabouni, Abdelhakim Hafid, and Pratap Kumar Sahu. "A collision-based beacon rate adaptation scheme (CBA) for VANETs". In: *IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*. 2013.

[13] Gaurav Bansal, John B Kenney, and Charles E Rohrs. "LIMERIC: A linear adaptive message rate algorithm for DSRC congestion control". In: *IEEE Transactions on Vehicular Technology* 62.9 (2013), pp. 4182–4197.

[14] Ghulam Abbas, Sami Ullah, Muhammad Waqas, Ziaul Haq Abbas, and Muhammad Bilal. "A position-based reliable emergency message routing scheme for road safety in VANETs". In: *Computer Networks* 213 (2022), p. 109097.

[15] Jinru Li. "Cooperative Beacon Rate Adaptation for Vehicular Ad-hoc Network". PhD thesis. KTH Royal Institute of Technology, 2021.

[16] Meng Wang, Tong Chen, Fei Du, Juan Wang, Guanxiang Yin, and Yuejin Zhang. "Research on adaptive beacon message transmission power in VANETs". In: *Journal of Ambient Intelligence and Humanized Computing* (2020), pp. 1–13.

[17] Feng Lyu, Hongzi Zhu, Nan Cheng, Yanmin Zhu, Haibo Zhou, Wenchao Xu, Guangtao Xue, and Minglu Li. "ABC: Adaptive beacon control for rear-end collision avoidance in VANETs". In: *15th IEEE International Conference on Sensing, Communication, and Networking (SECON)*. 2018.

[18] Fei Li and Chuanhe Huang. "A mobility prediction based beacon rate adaptation scheme in VANETs". In: *IEEE Symposium on Computers and Communications (ISCC)*. 2018.

[19] Hoa-Hung Nguyen and Han-You Jeong. "Mobility-adaptive beacon broadcast for vehicular cooperative safety-critical applications". In: *IEEE Transactions on Intelligent Transportation Systems* 19.6 (2017), pp. 1996–2010.

[20] Daniele Barbieri, Ilaria Thibault, David Lister, Alessandro Bazzi, Barbara M Masini, and Oreste Andrisano. "Adaptive beaconing for safety enhancement in vehicular networks". In: *15th International Conference on ITS Telecommunications (ITST)*. 2017.

[21] Forough Goudarzi and Hamid Asgari. "Non-cooperative beacon rate and awareness control for VANETs". In: *IEEE Access* 5 (2017), pp. 16858–16870.

[22] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun Pang Flotterod, Robert Hilbrich, Leonhard Lucken, Johannes Rummel, Peter Wagner, and Evamarie Wiebner. "Microscopic Traffic Simulation using SUMO". In: *IEEE Conference on Intelligent Transportation Systems (ITSC)*. 2018.

[23] Andras Varga and Rudolf Hornig. "An Overview of the OMNeT++ Simulation Environment". In: *1st International Conference on Simulation Tools and Techniques for Communications, Networks and Systems (SIMUTools)*. 2008.

[24] Christoph Sommer, Reinhard German, and Falko Dressler. "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis". In: *IEEE Transactions on Mobile Computing* 10.1 (2011), pp. 3–15.

[25] Maxime Gueriau and Ivana Dusparic. "Quantifying the impact of connected and autonomous vehicles on traffic efficiency and safety in mixed traffic". In: *IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. 2020.

AUTHORS

Nathalie Mitton received the MSc and PhD. degrees in Computer Science from INSA Lyon in 2003 and 2006 respectively. She received her Habilitation à diriger des recherches (HDR) in 2011 from Université Lille 1. She is currently an Inria full researcher since 2006 and from 2012, she is the scientific head of the Inria FUN team which designs protocols for wireless communications in constrained networks. Her research interests focus on self-organization from PHY to routing for wireless constrained dynamic and mobile networks. She has been nominated as one of the 10 women stars in computer Science in 2020 by the IEEE Communication Society. She has published her research in more than 40 international revues and more than 120 international conferences. She coordinates the Horizon Europe SLICES-PP project, participates in different Horizon Europe projects (NEPHELE, etc) and in several program and organization committees such as Infocom (since 2019), PerCom (since 2018), DCOSS (since 2018), ICC (since 2015), Globecom (since 2017), etc. She also supervises several PhD students and engineers.

Yasir Saleem received the B.S. degree in Information Technology from the National University of Sciences and Technology (NUST), Islamabad, Pakistan, in 2012, the M.Sc. degrees in Computer Science by Research from Sunway University, Malaysia, and Lancaster University, U.K. (under a dual degree program) in 2015, and the Ph.D. degree from Institut Mines-Telecom, Telecom SudParis, Institut Polytechnique de Paris, France in 2020. He is currently a Lecturer in Computer Science at Aberystwyth University, United Kingdom since January 2022. Previously, he was a Postdoctoral Researcher in Future Ubiquitous Networks (FUN) team at Inria Lille–Nord Europe, France. He has over 30 publications in international journals and conferences. His research interests include Internet of Things, smart cities, vehicular networks, intelligent transportation systems, wireless networks and artificial intelligence. http://www.yasirsaleem.com/.

Valeria Loscri is a permanent researcher in the FUN Team at Inria Lille (France) since Oct. 2013. From Dec. 2006 to Sept. 2013, she was Research Fellow in the TITAN Lab of the University of Calabria, Italy. She received her MSc and PhD degrees in Computer Science in 2003 and 2007, respectively, from the University of Calabria and her HDR (Habilitation à diriger des recherches) in 2018 from Université de Lille (France). Her research interests focus on emerging technologies for new communication paradigms such as Visible Light Communication (VLC), mmWave, cooperation and coexistence of wireless heterogeneous devices and cyber security in wireless networks. She is involved in the activity of several European Projects (Horizon Europe MLSysOps, H2020 CyberSANE, FP7 EU project VITAL, etc.). She has been nominated to the 2021 Women Stars in Computer Networking and Communications by the IEEE Communication Society. She is in the editorial board of IEEE COMST, IEEE Transactions on Information Forensics and Security (TIFS), Elsevier ComNet, ComCom. She is Action Chair and scientific Holder of BEiNG-WISE COST Action (since 2023). She is serving as TPC members in several primary internaional conferences, such as IEEE ESORICS, IEEE CNS, IEEE INFOCOM, IEEE PerCom. Since 2019, she is Scientific International Delegate for Inria Lille.

Christophe Bureau ("Xtof") has a Ph.D in Quantum Physics from Ecole Normale Superieure in Paris. He is co-Founder and CEO of #dotdot, the first ever marketplace for the anonymous P2P resale of unused Gigabytes of mobile data allowance, based on a self-healing WiFi mesh of autonomous pocket-size nodes. For the past 25 years, Xtof has walked the entire path of creating, driving and launching sources of growth via new business opportunities: through business model game change and/or technology breakthrough, he has "turned ideas into invoices", which is a fair de inition of innovation after all! He has invented technologies now implanted in millions of humans, driving these technologies from research to product development, to regulatory approval to market launch. He has secured funding of these projects through both institutional and private equity, as well as financial return to shareholders through trade sale. In between 2 startups (#dot:dot is number 5), he has occupied diverse leadership positions in large groups in the pharma, medical device and automotive industries, his last position being Head of Innovation at Valeo. Fan of consumer centricity and design thinking and their link to willingness to pay, he has a passion for value creation for the buyer, whatever the technical field.