



Revisiting WiFi offloading in the wild for V2I applications

Furong Yang^{a,b,c}, Andrea Ferlini^d, Davide Aguiari^e, Davide Pesavento^f, Rita Tse^g,
Suman Banerjee^h, Gaogang Xie^{i,b}, Giovanni Pau^{e,j,c,*}

^a Institute of Computing Technology, Chinese Academy of Sciences, China

^b University of Chinese Academy of Sciences, China

^c Sorbonne University, France

^d University of Cambridge, UK

^e University of Bologna, Italy

^f National Institute of Standards and Technology, USA

^g Macao Polytechnic Institute, Macao SAR, China

^h University of Wisconsin – Madison, USA

ⁱ Computer Network Information Center, Chinese Academy of Sciences, China

^j University of California, Los Angeles, USA

ARTICLE INFO

Keywords:

V2I
WiFi offloading
Measurement

ABSTRACT

This paper revisits the opportunities of using WiFi offloading for Vehicle to Internet (V2I) communication, and how this has changed over the last decade. With the rollouts of provider-managed WiFi networks that are more structured and operate under authenticated regimes, WiFi offloading, or use of available (roadside) WiFi networks for V2I data communication, has different opportunities and challenges. To study the current landscape, we develop a system (*X-Fi*), which efficiently selects, associates to, authenticates with, and performs WiFi offloading for V2I communication with these networks, and a tool (*X-Perf*), which illustrates opportunities of WiFi offloading available today in these networks, with measurements and experiments across four metro areas across three continents over 22 months. Our results indicate the feasibility of achieving 1 GB/hour application goodput, an order of magnitude higher than the number provided by open WiFi networks in the past, which can take a significant load away from alternative communication paths for V2I systems. Moreover, we provide several implications on transport protocols and WiFi deployments to shed light on the use of such WiFi networks for V2I communication.

1. Introduction

Vehicle to Internet (V2I) communication has been a topic of significant research over the years. New systems that improve V2I communication [1,2] and measurement studies of how commercial and research systems fare under the prevalent network conditions have also described the evolution of such systems [3,4]. Many efforts observed that if every vehicle were to connect and communicate with Internet-hosted infrastructure, the aggregate data volumes are likely to overwhelm common cellular capacities. Therefore, a common thread of work has also been to explore *opportunistic WiFi offloading*, often using open WiFi Access Points (APs). Cabernet [3], circa 2008, is a good exemplar of such a system where various location-related data from Boston taxis were uploaded using such opportunistic encounters with such open WiFi APs. The authors conducted an extensive measurement

study that demonstrated the feasibility of such a system and for their region of interest, achieved a mean throughput of 86 Kbps.

Evolution of urban WiFi: The last decade has seen dramatic changes in the way WiFi networks are available in large metropolitan areas. In many cases, wireless providers have begun to offer WiFi service to their paid customers through APs mounted across street poles, and backed over fiber, or point-to-point wireless links. This allows users to sometimes have widespread data services across the entire city. Unlike opportunistic and open WiFi networks, these use *authentication* for access, and communication occurs over *encrypted* data channels. Further, these networks are deployed in a structured manner by the operator to maximize coverage across the city. Such structured and authenticated WiFi networks with encrypted data paths have significant implications

* Corresponding author at: University of Bologna, Italy.

E-mail addresses: yangfurong@ict.ac.cn (F. Yang), af679@cam.ac.uk (A. Ferlini), davide.aguiari2@unibo.it (D. Aguiari), davide.pesavento@nist.gov (D. Pesavento), ritatse@ipm.edu.mo (R. Tse), suman@cs.wisc.edu (S. Banerjee), xie@cnic.cn (G. Xie), giovanni.pau@unibo.it (G. Pau).

<https://doi.org/10.1016/j.comnet.2021.108634>

Received 9 August 2021; Received in revised form 24 October 2021; Accepted 16 November 2021

Available online 30 November 2021

1389-1286/© 2021 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Table 1

MNOs' cellular data plans [18]. Throughput is presented in an order of downstream and upstream, which is throttled to 128 Kbps by many providers if the data cap is reached; Italian providers stop any service if consume > 100 GB/month.

Operator	Monthly cost	Unrestricted usage	Data cap	Throughput claimed (Mbps)	Throughput sensed (Mbps)
AT&T (USA)	85 USD	×	>22 GB	at least 10/1	27.5/6
T-Mobile (USA)	75 USD	×	>33 GB	usually 5–31/3–15	25.8/8.6
Verizon (USA)	70 USD	×	>75 GB	940/880	25.3/7.9
CTM (Macao)	60 USD	×	>10 GB	112/37.5	35.5/30.94
Vodafone (IT)	40 €	✓	–	1200/150	24.3/7
TIM (IT)	50 €	×	>100 GB	2000/300	24.1/7.6
Wind (IT)	12 €	×	>100 GB	1000/75	21.6/7.5
Free (FR)	20 €	×	>100 GB	Unclaimed	14.2/5.2

in delays for access, and the ability for “handoffs” impacting the use of WiFi offloading in vehicular settings.

Evolution of connected vehicles: In parallel to the developments in the wireless world, significant changes have emerged in vehicular systems and applications. Vehicles are more and more becoming large data *prosumers*. There is the vision of fully autonomous vehicles on the horizon (Level 5 automation [5]), with major trials being led by big players (Waymo, Uber, etc.) in select markets. Intel speculates each autonomous vehicle will produce ~4 TB of data per day, mostly from the LiDAR and the cameras [6]. Although the actual deployment of autonomous vehicles at scale is still a few years out, various forms of assisted driving systems are gaining in popularity. Also, systems such as Tesla Autopilot, requiring the download of 500 MB per update [7–10], and other Internet-connected Advanced Driver Assistance Systems perform continuous data collection from vehicles on the road (~5.6 Gbps), aggregate them centrally, and share learnings back with the vehicles again. Moreover, these connected vehicles perform numerous data exchanges with cloud-hosted services, e.g., for over-the-air firmware updates, such as BMW [11], to support Connected Battery Management Systems [12], for 3D map updates [13,14], and for various AI-based driving functions.

Today's approaches for connected vehicles: At present, connected vehicles exclusively rely on cellular networks. However, this approach poses significant costs to all the stakeholders in the connected vehicle arena: (1) cellular networks are expensive for large data users; (2) mobile network operators (MNOs) suffer from such users because they are already facing the pressure of the ever-growing mobile traffic [15], and, even in their most expensive plans, reduce data rates or stop Internet access past a given data consumption threshold¹; (3) adding millions of high demanding vehicles to the network will force operators to make massive infrastructure investments [16,17]. Table 1 shows a summary of several popular MNOs' cellular data plans in Europe, the USA, and Macao, which corroborates the above statements.²

The potential of WiFi offloading: We believe that offloading part of V2I traffic to WiFi would be a critical part of the future, in which vehicles roam onto provider-managed WiFi infrastructures that are increasingly getting deployed across major metros. Those provider-managed WiFi services are marketed as a perk attached to the home-Internet contract and, therefore, included in its cost. Therefore, users can access their operator's WiFi nationwide without additional payment. Moreover, even for those users who are currently not subscribers

¹ Reducing the data rate for a specific user is a common practice among mobile network operators — in the US, high-end AT&T and T-Mobile plans have a data cap per billing cycle of 50 GB and 100 GB, respectively. Once a user reaches the data cap, their service speeds are reduced.

² As there is currently no publicly available data on the pricing plans for connected vehicles, we take the ones for smartphones as references. We note that the pricing plans for connected vehicles may be different from the ones for smartphones in the future. But, we expect that the price will still be much higher than that of provider-managed WiFi networks, since MNOs will have CAPEX and OPEX for the infrastructures to support connected cars as for the current cellular networks.

of those WiFi providers, they can subscribe to the cheapest plan of those WiFi providers, which is usually much cheaper than the common plan that provides cellular data access. For instance, the cheapest subscription offered by Free (France) is 2€ per month, which is much cheaper than the common subscription (19.99€ per month).³ The ongoing transition toward 5G [19,20] and the deployment of future WiFi standards (i.e. WiFi 6) [21] are not going to change the current cost–benefit ratio. 5G cells will trade high performance at the price of coverage. Besides, 5G architecture is designed to natively integrate multi-radio access, and WiFi is one of them [22,23]. Thus, it is reasonable assuming 5G cellular access will have a per-user cost similar to LTE, with WiFi remaining a free perk of home network providers. We note that WiFi offloading alone is suitable for non-real-time bulk transfer applications. But, with some supports from multi-path transport (e.g. Multipath TCP [24], Multipath QUIC [25,26]), WiFi offloading together with cellular networks may also contribute to real-time applications.

The problem: Although the benefits of WiFi offloading in terms of costs are appealing, today's provider-managed WiFi networks introduce new challenges. For instance, the authentication and DHCP server, deployed in core networks, may be geographically distant from the APs, which introduce extra latencies that may prevent vehicular clients from quickly getting Internet connectivity. Thus, an investigation on the feasibility of offloading connected vehicles' data to provider-managed WiFi networks is of paramount importance. Further, a quantitative analysis of the transmission capability of such networks in the V2I context is also crucial for incentivizing all players in the connected vehicle arena to explore the possibility of WiFi offloading. However, all past studies [3,27–31] focused on open WiFi networks or limited WiFi deployments around ten years ago, which almost no longer exist [32,33] and are fundamentally different from provider-managed WiFi deployments that are authentication-required, encrypted, more structured, and powered by more advanced networking techniques (e.g. MIMO has been introduced since 802.11n [34]).

Our solution: To fill the above gap, we present *an extensive measurement study of V2I WiFi offloading feasibility in-the-wild, across 4 cities and 3 continents, covering more than 5620 km, over a period of 22 months*. To conduct the measurement study, a system, *X-Fi*, that enables a moving car to connect to provider-managed WiFi networks, proven to work across multiple countries and operators, and a tool, *X-Perf*, to measure TCP performances in intermittent networks are introduced.

Main takeaway: The key observations made by this paper are: (1) the feasibility of WiFi offloading based on provider-managed networks is confirmed; (2) while today's provider-managed WiFi networks provide comparable coverage, they offer an order of magnitude higher goodput (~1 GB/h) than what was observed in previous studies based on open WiFi networks [3,27–31]; (3) the time spent for establishing Internet connectivity via roadside APs is lengthened in provider-managed networks, compared with the one in open WiFi networks [3], due to the more complicated interactions with the authentication and DHCP server via multi-hop paths; nevertheless, it only counts for a small part

³ <https://mobile.free.fr/>.

Table 2

Comparison with previous work (UO=unplanned open; PM = provider-managed; [abgn/ac] denotes 802.11 standards; C = Cities; a few numbers marked with † are indirectly calculated based on the data from the papers). On a high level, compared with open WiFi networks, current provider-managed WiFi networks provide much higher goodput and comparable coverage. All the numbers regarding connectivity are medians.

	Date	WiFi deployment	Measurement scale	Availability		Long-term goodput (MB/h)	End-to-End Internet connectivity			
				Code	Data		Establishment time (s)	Duration (s)	Connectivity hole (s)	
Soroush et al. [29]	2004~2008	UO APs [abg] (mainly)	1C (US)	✗	✓	90	✗	✗	✗	
Bychkovsky et al. [27]	2006	UO APs [b]	1C (US)	✗	✗	9.7†	12.9	13	10	
Eriksson et al. [3]	2008	UO APs [bg]	1C (US)	✓	✗	38	0.29	4	32	
Balasubramanian et al. [31]	2010	UO APs [b] (mainly)	3C (US)	✗	✗	14.4†	✗	✗	✗	
Deshpande et al. [28]	2010	PM APs [bg]	1C (US)	✗	✗	128†	✗	2.5	3.5	
Soroush et al. [30]	2011	UO APs [abg]	2C (US)	✓	✗	274	✗	20	25	
Our work	Macao	2019~2021	PM APs [n/ac]	4C (US/ASIA/EU)	✓	✓	752	1.4	9	31
	Los Angeles						1471	2.1	22	12
	Bologna						1256	0.9	10	5
	Paris						181	0.6	9	17

(10%–18%) of the whole association duration; (4) the AP selection algorithm of stock WiFi client implementations, which solely relies on the initial signal strength (ISS), is not sufficient in the V2I context, as no correlation between the ISS and performance is observed; (5) 5 GHz APs generally offer better performance than 2.4 GHz APs when it comes to V2I WiFi offloading; (6) the ability to allow clients to retain IP addresses across associations is not a must implemented by provider-managed networks by default; thus, application designers must have this caveat in mind, and, either handle IP handovers at the application layer or adopt transport protocols (e.g. QUIC [35,36]) that tolerant IP handover; (7) BBR [37] and CUBIC [38] deliver similar application goodput; moreover, for both BBR and CUBIC, application goodput can benefit from using multiple concurrent flows.

Contributions of this paper: To sum up, the contributions of this paper are: (1) the measurement system, including *X-Fi* and *X-Perf*, built for studying the current landscape of WiFi offloading based on provider-managed APs; (2) the one-of-its-kind dataset, which demonstrates the feasibility and capability of WiFi offloading via provider-managed networks; (3) the analysis and implications made from the dataset, which shed light on the use of provider-managed WiFi for connected vehicles. We have released the dataset⁴ and the code⁵ of *X-Fi* and *X-Perf* for the community to further study the vehicular WiFi offloading landscape.

The remainder of this paper is organized as follows. We describe the related work in Section 2. Next, we present the measurement system in Section 3. Then, the methodology that we used for the study is presented in Section 4. After that, Section 5 gives a summary about our dataset. Thereafter, we present the detailed analysis regarding connectivity and transport performance in Section 6 and Section 7 respectively. In Section 8, we discuss the future work and the relationship between V2I WiFi offloading via provider-managed WiFi networks and other V2I communication solutions (e.g. 5G). Finally, we conclude this paper in Section 9.

2. Related work

Using WiFi to provide moving vehicles with Internet access has been studied by many researchers in the past. Table 2 summarizes a comparison between our work and previous art. Notably, our results,

especially in goodput, are very different from the previous ones, which shows how the V2I WiFi-offloading landscape has changed.

Accessing the network: Ott et al. [39] confirmed a vehicle can connect with individual WiFi APs placed along the roadside. Gass et al. and Hadaller et al. [40,41] studied the impact of the back-haul on WiFi-based V2I connections. Besides, Hadaller et al. [41] identified limits of stock WiFi stack implementations and provided guidelines to improve them. *While they focus on characterizing the V2I connectivity from individual WiFi APs, our work focuses on deployment-ready WiFi-based in-the-wild offloading using provider-managed WiFi networks.* Mahajan et al. [42] studied link-layer characteristics of V2I WiFi communications and found that the WiFi connectivity often has intermittent poor-quality periods. Soroush et al. [29] reported link-layer V2I connectivity results based on their experience operating a metro-scale V2I WiFi testbed.

V2I WiFi performance evaluation: The authors of [27,31] measured the TCP performance of V2I connectivity using in-situ open, unencrypted, WiFi networks. They confirmed vehicles can upload hundreds of KB per WiFi association. However, their results are not representative as of today: (1) open WiFi networks have almost disappeared in urban areas [33]; (2) their measurement results are based on HostAP-MadWiFi, a companion that used low-level features that no longer exist. The authors of Cabernet [3] proposed a streamlined WiFi client to shorten the association overhead with transient APs from moving vehicles and an enhanced transport protocol to cope with intermittent connectivity. Soroush et al. [30] proposed to maintain concurrent associations with different APs on the same channel to enhance the performance and reliability of V2I WiFi connectivity. Again, both their evaluations were conducted based on open WiFi networks, thus, unable to reflect today's V2I WiFi offloading landscape based on provider-managed WiFi networks. Deshpande et al. [28] performed a head-to-head comparison between 3G and a commercially operated WiFi network in the V2I context. However, their setup is not in accord with the situation of current provider-managed WiFi networks, as WiFi authentication and IP acquisition are only needed once at the beginning of a whole drive in their experiments. Further, to reduce the effects of WiFi intermittent coverage, they limited the server-side TCP connection timeout to 1s. *In contrast, our measurement system follows the standard WiFi association and DHCP flow, and, uses the stock TCP implementations to study today's landscape of V2I WiFi offloading based on provider-managed WiFi networks as-is.*

Custom transport protocols for V2I: The authors of [3,43,44] proposed new transport protocols to hide the intermittent characteristics

⁴ https://github.com/yangfurong/X-Fi_dataset.

⁵ <https://github.com/yangfurong/X-Fi-code>.




Fig. 1. Connection procedure handled by *X-Fi/X-Perf*. Auth., assoc., req., resp., ack., and disc. are abbreviations for authentication, association, request, response, acknowledgment, and discovery respectively. MAC auth. represents the authentication with APs, while access auth. means the authentication with the central authentication server via the EAP protocol [49]. If the connection is interrupted at any stage, the system goes back to the scanning stage.

of V2I WiFi connectivity and improve the throughput of applications. More recently, Lee et al. [33] studied the usage of provider-managed WiFi networks as one of the connectivity options for MPTCP [24] to enhance the performance of low-latency applications on connected vehicles. These works require network stack upgrades and/or deployments of proxies on the Internet. *Since our goals are to understand how well today's urban WiFi networks can support vehicular applications as-is, we decided to use TCP as our transport protocol that is predominantly used by existing applications.* Hence, we do not impose any modification to the existing transport protocol.

Network assisted WiFi offloading for V2I: Several researchers [45–48] proposed better AP handoff mechanisms to enhance the V2I WiFi performance. However, all of them require modifications to the APs and, therefore, control over the infrastructure. *In our work, we could not adopt any of their approaches as we have no control over the infrastructure.*

3. Measurement system overview

X-Fi (Section 3.1) and *X-Perf* (Section 3.2) is a ready-to-deploy toolkit for WiFi offloading in-the-wild. *X-Fi* incorporates a set of optimizations to provide enhanced access to provider-managed WiFi networks from a moving vehicle. *X-Perf* is a measurement tool to study transport performance of intermittent connectivity. Fig. 1 illustrates the workflow of the toolkit.

3.1. *X-Fi* : connecting a moving vehicle to provider-managed WiFi hotspots

Motivation: Given that the encounters between vehicles and APs are commonly short, a WiFi system that can quickly set up connectivity with encountered APs is of great importance. Although Cabernet [3] proposed a vehicular WiFi client tuned for open WiFi networks via a set of heuristics, it does not support EAP authentication [49], which is prevalent in current provider-managed WiFi networks and is more suitable for vehicular uses than WEB-based authentication. On the other side, the mainstream WiFi systems for connectivity setup, such as *wpa_supplicant* and *dhcpcd* on Linux, support EAP authentication, but, they are not well-tuned for vehicular uses. Therefore, we propose *X-Fi* that consolidates the advantages from the two sides above.

X-Fi consists of two different pieces: (1) an enhanced *wpa_supplicant*, which uses a number of best-of-practice heuristics to speed up the WiFi association process; (2) an enhanced DHCP client (based on *dhcpcd*) customized for V2I.

WiFi association: The WiFi association process consists of a series of packet exchanges between the WiFi client, the AP, the EAP authentication server, and the DHCP server as shown in Fig. 1. Given that the vehicular WiFi connectivity has high loss rates in the fringe area of an AP's coverage [41], the WiFi client needs to retransmit the lost packets as fast as possible. As *wpa_supplicant* uses the timeout/retry mechanism to detect and retransmit lost packets, the proper setting of the timeouts is crucial. Therefore, we shortened the timeouts in *wpa_supplicant* to 100 ms and the max retry limit to 5 times, as suggested by Cabernet [3].

Besides the changes on timeouts, we also optimized the AP scanning strategy of *wpa_supplicant* via opportunistic active scans. *X-Fi* normally uses passive scans where the WiFi client dwells on each channel for around 80–100 ms to listen for beacons from nearby APs to get a full list of nearby APs. Then it tries to associate with the AP with the highest Signal-to-Noise Ratio (SNR). If this association fails, *X-Fi* will try to launch an active scan on the channels that have other APs with target ESSIDs,⁶ (this information was learned from the previous passive scan). The rationale behind this heuristic is that the vehicle is likely still in the range of those APs' coverage, thus, no need to waste time on scanning other channels. If this active scan does not find any AP with a target ESSID, *X-Fi* goes back to a full passive scan. Otherwise, *X-Fi* tries to associate with the AP with the highest SNR. If this association still fails, *X-Fi* can repeat the active scan and association process as long as there still are channels that have untried APs with target ESSIDs and the time since the last passive scan has not exceeded a threshold. In our experiments, we configured this threshold to 5 s, considering that the median AP coverage is 50–100 m (see Fig. 5(b)) and 30 km/h is a typical speed of cars running in urban areas.

IP address acquisition: Once the supplicant has established L2 connectivity to the access point, it is necessary for the client to obtain an IP address in order to communicate with other hosts on the Internet. This is done through the DHCP. The off-the-shelf DHCP client, i.e. *dhcpcd* on Linux, is not tuned for vehicular environments. Specifically, the inefficiency of *dhcpcd* mainly comes from two perspectives: (1) the timeouts for DHCP messages are too long for vehicular scenarios as demonstrated by Cabernet [3]; (2) the ARP probing mechanism, sending ARP probe packets to check whether the assigned address is in-use by other clients in the same LAN [50], is enabled by default, which may lengthen the IP acquisition process by several seconds. Therefore, we implemented a custom DHCP client based on *dhcpcd* that has shortened timeouts and disables the ARP probing mechanism. We note that the timeouts are set to 100 ms in our DHCP client, as it is the value recommended by Cabernet [3].

Evaluation of *X-Fi*: We compared *X-Fi* with the off-the-shelf WiFi system for connectivity setup, *wpa_supplicant/dhcpcd*, to demonstrate the efficiency of *X-Fi*. The experiments were done as follows. We equipped a car with an on-board computer (see Section 4.2 for more details about the hardware). Then, we drove the car for 10 laps following a fixed route in Bologna (the driving time is around 6 h in total), where we used *X-Fi* for connectivity setup for half of the laps and the *wpa_supplicant/dhcpcd* suite for the other half of the laps. For every WiFi association, we recorded the “time until IP acquisition” that means the time spent for the WiFi association and DHCP process in total. The results are shown in Fig. 2: *X-Fi* is significantly faster than the vanilla *wpa_supplicant/dhcpcd* suite with default configurations. The major improvement comes from disabling the ARP probing mechanism in DHCP, while the other heuristics, e.g. shortened timeouts, improve the tail latency of connectivity setup.

⁶ Like *wpa_supplicant* *X-Fi* allows users to supply a configuration file to specify the ESSIDs of the WiFi networks that it wants to connect to.




Fig. 2. *X-Fi* versus the vanilla *wpa_supplicant/dhcpd* suite. “Vanilla” stands for the standard *wpa_supplicant/dhcpd* suite with default configurations. “NoARP” represents that the time spent for the ARP probing in DHCP is excluded when we compute the total connectivity setup time for the associations using the vanilla *wpa_supplicant/dhcpd* suite.

3.2. *X-Perf* : measuring transport performance over intermittent connectivity

TCP underpins the vast majority of internet applications. Therefore, studying TCP performance over WiFi V2I scenarios is of paramount importance. The popular TCP performance measurement tool, *iPerf*, lacks intermittent networks support.⁷ Specifically, *iPerf* has two problems: (1) it cannot continue the transmission session if an IP handover caused by an AP handoff happens; (2) even if the vehicle’s IP address does not change after the interruption caused by an AP handoff, *iPerf* can take a very long time to resume the transmission because the TCP RTO timer may increase to a very large value. Therefore, to make full use of the transient transmission window of V2I intermittent connectivity to test the TCP performance of WiFi offloading without assuming that the WiFi networks support IP roaming, we designed *X-Perf*, a dedicated performance measurement tool for intermittent networks.

The working process of *X-Perf* is illustrated in Fig. 3. Once *X-Perf* is up and running, it executes *gpxlogger* and *X-Fi*, respectively to track GPS coordinates and to start the attempts to associate with WiFi hotspots. Then, *X-Perf* monitors the status of the WiFi interface operated by *X-Fi*. Once the network interface obtains an IP address, *X-Perf* starts to probe the performance of TCP over the wireless link. Specifically, it first decides the transmission configuration for the following TCP bulk transfer, which is randomly selected from a predefined configuration pool. We note that the transmission configuration includes the direction of the transfer (upload or download), the TCP congestion control algorithm (CCA), and the number of concurrent TCP flows (*flow number*). By supplying different configurations to the configuration pool, *X-Perf* allows us to study the impact of different parameters on TCP performance. Thereafter, it measures the TCP throughput by performing data bulk transfers between a remote server and the local onboard computer. During the TCP transmission, all exchanged TCP packets are recorded by *tcpdump* at the client-side. As soon as *X-Perf* finds that the WiFi interface has lost its IP address, which means that the WiFi connectivity is interrupted, *X-Perf* stops the TCP throughput measurement and archives all data collected in this measurement run. The collected data contain a wealth of information from L2 to the application layer including packet-level traces, and all information is also geo-referenced and synced with GPS. *X-Perf* repeats the above process each time *X-Fi* connects to an AP until the end of the experiment or until the vehicle is turned off.

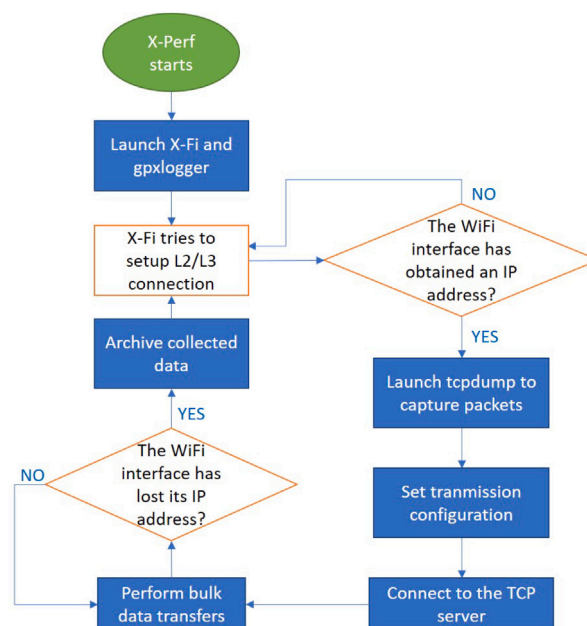


Fig. 3. The flowchart of *X-Perf*. For the readability of the flowchart, the terminal state, “the experiment is done/the vehicle is off”, is removed from the flowchart.

4. Measurement methodology

In this section, we detail our the methodology of our measurement study. First, an overview of the cities selected for our measurement activities is presented in Section 4.1. Next, we describe our measurement apparatus setup in Section 4.2. Thereafter, the metrics used for data analysis are introduced in Section 4.3.

4.1. Cities selected for measurement

To understand the state-quo of V2I WiFi offloading across different traffic patterns and networking infrastructures, we conducted our measurement activities across four different cities in Europe (Bologna and Paris), in Asia (Macao), and in North America (Los Angeles). The diversity of the cities enables us to collect measurements in radically different urban scenarios, traffic patterns, and viability. Specifically, *Bologna* is a good sample of a medium-size touristy medieval city with narrow streets, which have an unpredictable traffic pattern. *Macao* is, again, a medium-size city characterized by a number of high-rise buildings as skyscrapers and casinos. *Paris* is an extremely touristy and crowded metropolis with a well-defined city center. *Los Angeles* (LA) is a very populated megalopolis, with an alternating of residential areas, working districts, and rural regions.

In each city, we subscribed to one or more local providers offering access to urban WiFi APs. This gave us the WPA credentials/SIM cards needed to authenticate and connect while driving. Different cities, and different providers, have different authentication strategies: PEAPv0/EAP-MSCHAPv2 [51] in Bologna, LA, and Macao; EAP-SIM [52] in Paris. Besides, the access point deployment strategies that we encountered during our data collection can be grouped into two main approaches: (1) leveraging existing hardware installed in customers’ premises; (2) installing new dedicated APs. In the first approach, the providers rent the APs to their subscribers. The rented APs present two different ESSIDs when in operation. While one is exclusively used by the customers and secured with WPA/WPA2-PSK [53], the other serves as a WiFi hotspot for public users. Therefore, both the physical capacity of the 802.11 channel and the available bandwidth of the connection, result to be shared between private customers and

⁷ <https://github.com/esnet/ipperf/issues/835>.

public users. Quality of Service rules, as token-bucket algorithms, may be applied to prioritize the customers. On the other hand, the second deployment strategy consists of the installation, done by the network providers, of dedicated APs at either indoor or outdoor spots.

4.2. Measurement apparatus setup

Hardware: The computer used for the data collection is a Slimpro SP675FP Fanless Mini PC. It features an Intel x64 3rd Gen Core processor, a 3 × 3 MIMO full-size Mini PCIe 802.11n card (WLE350NX-7 A), an u-Blox EVK-7N GPS unit, 16 GBs of RAM, and 300 GBs of SSD for storage. While for the deployment in Macao we adopted three dipole antennas, in Bologna, Los Angeles, and Paris, we opted for three multi-polarized antennas. Using different antennas in Macao was not possible due to constraints imposed by the vehicle supplier; it is noteworthy that this limitation slightly affected the *X-Fi* performance in the Macao scenario.

Software: The onboard computer runs Ubuntu 16.04.1 with 4.10.0-33-generic Linux kernel. At the computer boot, *systemd* launches *X-Perf* (Section 3.2). Once up, it executes *gpxlogger* and *X-Fi* (Section 3.1), respectively to track the GPS coordinates and to start the association and connection attempts with nearby WiFi hotspots. Thereafter, *X-Perf* starts measuring the performance of WiFi offloading and collecting measurement results as explained in Section 3.2. We configured *X-Perf* to use the following configuration pool (explained in Section 3.2): CCA ∈ {BBR, CUBIC}, direction ∈ {upload, download}, flow number ∈ [1, 16]. BBR [37] and CUBIC [38] are chosen because the two are now the predominant CCAs on the Internet [54]. BBR is a recent work adopted by many Google services and outperforms classic window-based CCAs across various scenarios, especially over lossy links. CUBIC is window based and is the default Linux congestion control. This configuration pool enables us to compare different CCAs and study the impact of concurrent flows. It is noteworthy that system parameters related to Linux TCP stack (e.g. buffer size, etc.) are set by the operating system to the default values.

Permanent deployment in Macao: In Macao, we placed the Slimpro Mini PC under the back seats of a granted van. Two WiFi antennas and the GPS antenna were installed on its roof and one extra WiFi antenna under the back seats for space concerns. The van runs errands on behalf of the granting institution and transport staff. For the whole duration of the experiment, the mobility pattern of the van remained unaltered from its usual one, allowing us to collect in-the-wild samples representative of transporters' real-world mobility. The TCP server was hosted in a virtual machine (VM) based in Singapore to reduce latency.

Short-term trial in Paris: In Paris, we instrumented a vehicle with: three WiFi antennas and one GPS antenna placed on the roof of the vehicle, the Slimpro Mini PC located on the front seat and powered by the vehicle. A *Dekart* SIM card reader is used for the EAP-SIM authentication required by *Free* [55]. We deployed the TCP server on a machine located in Paris. The driving routes covered more than half of the districts in Paris downtown.

Short-term trial in LA and Bologna: The vehicle used in LA was equipped as the one in Paris, except for the SIM card reader. Here, we methodically mapped both residential and commercial areas. We deployed the TCP server used for our tests on a cloud VM in LA. In Bologna, we replicated the LA setup, but with the server in Paris.

4.3. Metrics of interests

Connectivity: Here now follows the list of metrics used to understand the characteristics of the V2I WiFi connectivity. We note that the first nine metrics are also illustrated via Fig. 4 to make them more intuitive.

- **Association setup time:** The time needed for the association setup process.




Fig. 4. The illustration of part of the metrics. In the figure, A1, A2, A3, and A4 stand for four successive associations.

- **DHCP time:** The time for the DHCP process, reflecting the overhead for IP acquisition.
- **Time until the first TCP ACK:** The time between the start of the association and the first TCP ACK received from the server, reflecting the total time spent for establishing end-to-end Internet connectivity.
- **Link-layer connectivity duration:** The time that an association lasts.
- **IP connectivity duration:** The duration that an association has an IP address.
- **TCP connectivity duration:** The time between the start and termination of a TCP session.⁸ *X-Perf* stops a TCP session right after the WiFi interface loses its IP. Hence, the end of a TCP session coincides with that of IP connectivity.
- **Link-layer connectivity hole:** The idle period between two successive associations.
- **IP connectivity hole:** The idle period without an IP address between two successive associations.
- **TCP connectivity hole:** The idle period between two successive TCP sessions.
- **Initial signal strength (ISS):** The signal strength sensed during the WiFi scanning phase. It is the only metric considered by stock WiFi implementations during the AP selection. In a stationary case, selecting the AP with the highest ISS works well. We investigate whether this works as well, for in-motion vehicles.
- **Average vehicle speed per association:** The average vehicle speed during an association, which is calculated based on GPS traces and may affect the performance of V2I WiFi offloading.
- **WiFi frequency:** We seek to find whether there are differences between 5 GHz and 2.4 GHz APs in the context of V2I WiFi offloading.
- **IP address roaming support:** By analyzing the IP address at each association, we can observe how often a vehicle changes its IP (the local one) during a drive. This affects the design of applications and transport protocols for in-motion vehicles to transfer content via provider-managed WiFi networks.

Transport performance: We now present a list of the metrics used to assess the transport performance of V2I WiFi offloading.

- **Average TCP goodput:** The aggregated average goodput of a TCP session calculated based on *pcap* traces.
- **Data volume:** The total number of bytes that have been successfully transferred during a TCP session.

5. Dataset summary

The salient features of our dataset are summarized in Table 3. The WiFi networks we connected to in Bologna and Paris (*Fastweb*, *Emilia*

⁸ A TCP session represents one or multiple TCP flows (*X-Perf* may launch multiple concurrent flows to measure transport performance).

Table 3
Dataset summary.

	Paris	Bologna	LA	Macao	Total
Driving time analyzed	11 h 18 m 08 s	12 h 48 m 29 s	32 h 38 m 06 s	311 h 53 m 8 s	368 h 37 m 51 s
Driving distance analyzed (km)	144.71	197.12	707.91	4570.89	5620.63
Avg. driving speed (km/h)	12.8	15.39	21.69	14.66	N/A
# of distinct days analyzed	4	5	7	334	350
# of assoc. attempts	4525	5559	6746	21555	38385
# of link-layer assoc.	1812	1501	1856	5412	10581
# of 2.4 GHz assoc.	1812	1218	479	1336	4845
# of 5 GHz assoc.	0	283	1377	4076	5736
# of IP acquisitions.	790	1302	1583	5042	8717
# of TCP tests	664	1179	1490	4736	8069
# of distinct associated APs	1610	917	1385	1013	4925
# of distinct 2.4 GHz APs	1610	796	391	332	3129
# of distinct 5 GHz APs	0	121	994	681	1796
# of WiFi operators	1	3	4	1	9
Total TCP download (GB)	0.7	7	33.1	97.03	137.83
Total TCP upload (GB)	1.3	8.7	13.8	132	155.8




Fig. 5. Failed attempts are excluded. (a) CCDF of the average vehicle speed per association. Mean: Bologna 18.42, LA 23.43, Macao 10.38, Paris 14.83; (b) CDF of AP coverage. Mean: Bologna 70.53, LA 180.17, Macao 58.15, Paris 49.71.

Romagna regional provider, ALMAWIFI and Free) present a higher AP density than the ones in LA and Macao (*Spectrum, eduroam, TWCPasspoint, CableWiFi* and *Companhia de Telecomunicações de Macau* (CTM)). This leads to, on average, more associations per hour: 160, 117, 57, and 17 respectively for Paris, Bologna, LA, and Macao. Remarkably, we observed a lower ratio of IP acquisitions over associations in Paris than in the other cities (43.6% in Paris, 86.7% in Bologna, 85.3% in LA, and 93.1% in Macao). This much lower DHCP success rate in Paris is caused by failures in responding to the DHCP requests, even with good link-layer connectivity (-40 dBm). We confirmed that by connecting to some of these APs which failed to respond to the DHCP requests at locations where their signal strength was decent, between -40 dBm and -50 dBm. As we have no control nor inside knowledge of the network infrastructure used in Paris, we could not go further explaining the root cause of that. Although the ratios of IP acquisitions over associations are much different across the four cities, the ratios of TCP sessions over IP acquisitions are similar. Therefore, in each of the analyzed cities, once a moving vehicle obtains an IP address it is very likely able to establish TCP connections.

The complementary CDF (CCDF) of the average driving speeds per association in every city is shown in Fig. 5(a). Given the diversity of APs density and traffic conditions across the four cities, the average driving speeds per association are also dissimilar. In Paris, Macao, as well as in the downtown area of Bologna, where the speed limits are stricter, most of the associations happened at speeds lower than 30 km/h. Instead, outside the center of Bologna, as in LA, we had the chance to connect with outdoor APs placed along the main roads, driving at higher speed (40–62.5 km/h).

Fig. 5(b) reports the CDF of AP coverage, which is the diagonal length of the smallest bounding box enclosing all GPS points collected during the associations with that AP. The APs in LA present, on average, a much wider coverage (180.17 m) than the ones in the other cities,




Fig. 6. Failed attempts are excluded. (a) CDF of association setup time. Mean: Bologna 0.34, LA 0.73, Macao 0.77, Paris 0.67; (b) CDF of DHCP time. Mean: Bologna 0.60, LA 1.38, Macao 0.61, Paris 0.28.

denoting how the providers placed several outdoor APs along roads and intersections. Our results indicate that the coverage of a single provider-managed AP is generally similar to that of a single open AP [27].

6. Connectivity analysis

This section presents the measurement results from the connectivity perspective. We first analyze the time needed for connectivity setup in Section 6.1. Next, the duration of connectivity and holes without connectivity is considered in Section 6.2. Thereafter, we analyze the impact of vehicle speed, ISS, and WiFi frequency on connectivity in Section 6.3. Finally, the IP address roaming support of each provider-managed WiFi network is analyzed in Section 6.4.

6.1. Time for connectivity setup

Association setup time: The CDFs of association setup time is reported in Fig. 6(a). In Bologna, Macao, and Paris, 90% of the association setup processes terminate in less than one second. On the other hand, on average, the association setup process in LA takes slightly longer than in the other cities. None of the distributions is heavy-tailed, with the means only lightly skewed from the medians. The data points with relatively longer association setup time are due to the re-transmissions of the packets (i.e. authentication packets) necessary for association setup.

DHCP time: The CDF of DHCP time is plotted in Fig. 6(b). In Bologna, Macao, and Paris, the 90% of DHCP processes terminate within 1 s. As the association process, also the DHCP process in LA is slightly longer than in the other cities. The differences between cities might be caused by the ways how operators deploy their DHCP servers.

A more detailed breakdown of the phases prior to gaining IP is shown in Fig. 8(a). As expected, the DHCP and EAP authentication




Fig. 7. Failed attempts are excluded. (a) CDF of the *time until the first TCP ACK*; (b) CCDF of *link-layer connectivity duration*. Mean: Bologna 16.74, LA 29.28, Macao 21.27, Paris 16.58.




Fig. 8. (a) Breakdown of the overhead before establishing the IP connectivity. All the values of each phase are averaged over all associations having IP connectivity; (b) average ratio of *TCP connectivity duration* against the time needed to set it up.

account for the major part of the overhead. The reasons for this are two-fold: (1) the DHCP and EAP authentication cause more packet exchange between clients and APs than the other two; (2) during the DHCP and EAP authentication phase, clients communicate with DHCP and authentication servers normally in core networks, which causes longer latency than communication between clients and APs.

Time until the first TCP ACK: Fig. 7(a) presents the CDF of the *time until the first TCP ACK*. On average, it takes 1.09 s in Bologna, 2.76 s in LA, 1.48 s in Macao, and 1.11 s in Paris. We note that these numbers are larger than the one (0.37 s) in open WiFi networks [3]. This is due to the fact that provider-managed WiFi networks require clients to interact with the authentication and DHCP server located in core networks, which leads to lengthened process. Nevertheless, the time to obtain TCP connectivity only counts for a very small part of the entire association (10%–18%, Fig. 8(b)), meaning that, for more than the 80% of the time that a vehicle is associated with an AP, it is possible to transfer data.

Takeaway: The time spent for establishing end-to-end Internet connectivity in provider-managed networks is larger than the one in open WiFi networks because of the more complicated and lengthened interactions with the authentication and DHCP server in core networks. Therefore, optimizations, such as adopting authentication methods for faster handoffs [56,57] and granting extended DHCP leases to clients to retain its IP address, may reduce this time. Nonetheless, on average, this overhead only counts for a small part (10%–18%) of the whole duration of an association.

6.2. Connectivity duration and holes

Link-layer connectivity duration: Fig. 7(b) shows the CCDF of the *link-layer connectivity duration*. Median values are 9.22 s in Bologna, 17.29 s in LA, 7.99 s in Macao, and 9.04 s in Paris. The distributions are heavy-tailed, with a few long-lasting associations due to traffic lights and traffic jams, typical of urban mobility.




Fig. 9. Failed attempts are excluded. (a) CCDF of *IP connectivity duration*. Mean: Bologna 17.47, LA 31.91, Macao 21.97, Paris 15.84; (b) CCDF of *TCP connectivity duration*. Mean: Bologna 18.08, LA 33.57, Macao 23.24, Paris 17.24.

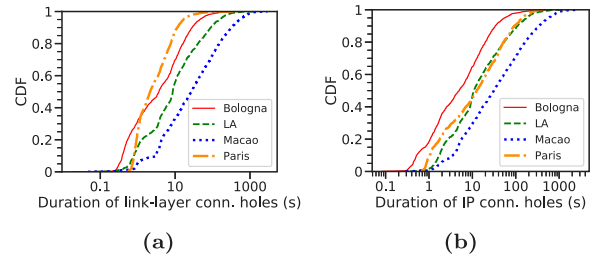


Fig. 10. (a) CDF of the duration of the *link-layer connectivity holes*. Mean: Bologna 12.23, LA 33.08, Macao 112.27, Paris 5.70; (b) CDF of the duration of the *IP connectivity holes*. Mean: Bologna 15.69, LA 40.70, Macao 120.72, Paris 34.42.

IP connectivity duration: Fig. 9(a) depicts the CCDF of the *IP connectivity duration*. As for the link-layer, also these are heavy-tailed. In Bologna, LA, and Macao, both mean and median *IP connectivity duration* are slightly longer than the ones at link-layer. This is due to the exclusion of failed IP acquisitions (Table 3), likely to have shortened *link-layer connectivity duration*.

TCP connectivity duration: The CCDF of the *TCP connectivity duration* is reported by Fig. 9(b). The mean and median values of the *TCP connectivity duration* are slightly higher than the ones of the *IP connectivity duration* due to the exclusion of associations unable to establish TCP connectivity, which tend to have shorter connectivity duration.

Link-layer connectivity holes: Due to the limited coverage of the existing urban WiFi deployments, there are gaps between successive associations. The CDF of the duration of the *link-layer connectivity holes* (Fig. 10(a)) reflects the frequency at which vehicles encounter APs, denoting their density in the region. Both in Bologna and in Paris, the AP density is capillary, with a median inter-arrival time between two consecutive associations of 3.69 s and 2.13 s, respectively. On the other hand, APs in LA and Macao are rarer, with longer median inter-arrival time (8.64 s and 27.31 s). Since in Paris we only drove through densely populated neighborhoods, the mean is not too far from the median, with no sudden nor significant changes in the AP density across different locations. For all the other cities, the mean values are skewed from their median by a factor of 4x, meaning the holes at the tail of the CDF are particularly long. This denotes how the AP distribution changes significantly across different regions in Bologna, LA, and Macao. The presence of a few extremely short holes (tens of milliseconds) is due to de-authentication packets sent by some APs to force the disassociation from them.

IP connectivity holes: As not every association succeeds in acquiring the IP address (Table 3), we expect the *IP connectivity holes* duration to be enlarged by failed IP acquisitions. For Bologna, LA, and Macao, the mean *IP connectivity holes* duration (Fig. 10(b)) are just marginally longer than those of the *link-layer connectivity holes* duration. This is




Fig. 11. (a) CDF of the TCP connectivity holes duration. Mean: Bologna 18.13, LA 43.45, Macao 124.39, Paris 42.45; (b) CDF of the ISS of associations. Mean: Bologna -67 , LA -60 , Macao -67 , Paris -59 .

because only a small fraction of associations in Bologna, LA, and Macao did not acquire the IP. Though, in Paris, the mean duration of the IP connectivity holes is greater than that at the link-layer. This is due to the DHCP issue discussed in Section 5.

TCP connectivity holes: From Fig. 11(a) we observe how the duration of the TCP connectivity holes is slightly longer if compared to that of IP. This is because of the exclusion of the fraction of associations that failed to establish a TCP connection, i.e. 15.9% in Paris, 9.4% in Bologna, 5.9% in LA, and 6% in Macao. From our dataset, it is possible to conclude how, on any commute lasting more than a handful of minutes, vehicles are likely to encounter several usable APs to actually transmit data.

Takeaway: The connectivity offered by provider-managed WiFi networks for in-motion vehicles is generally similar to that provided by open WiFi networks [3,27]: vehicles can expect to offload data for tens of seconds via an encounter with a provider-managed AP, and, are likely to have several such encounters every few minutes.

6.3. Impact of vehicle speed, ISS, and frequency

Next, we study the factors that may affect V2I WiFi connectivity. Specifically, we analyze the impact of vehicle speed, ISS, and WiFi frequency on the duration of link-layer connectivity. In the interest of space, we do not present the impacts on the duration of L3/L4 connectivity, as they are similar to the link-layer connectivity's ones.

Average vehicle speed per association: The first factor that we consider is the vehicle speed. Intuitively, higher speeds lead to shorter connectivity duration, as the vehicle leaves the area covered by APs quicker. Fig. 12 confirms the correlation between the vehicle speed and the link-layer connectivity duration.

Initial signal strength: Now, we consider ISS. In the stationary case, choosing the AP with the highest ISS works well. However, due to the continuous changes of the clients' position caused by vehicular mobility, the ISS alone is not enough to infer the network performance. In all the cities, the ISS ranges from being very weak (-80 dBm) up to an excellent signal level (greater than -50 dBm) as shown in Fig. 11(b).

To understand whether is enough to use the ISS as the only metric to infer, a priori, the network performance, we plot the correlation between the ISS and the link-layer connectivity duration (Fig. 13). From the plots, we cannot evince any obvious pattern confirming how associations with higher ISS last longer when mobility is involved. Surprisingly, some associations with a higher ISS are shorter than those with a lower ISS. This may be due to the vehicle moving from one extremity to the other of the AP coverage when it senses a weak ISS; while the vehicle could be already leaving the middle of the AP coverage when it senses a strong ISS. Hence, more sophisticated AP selection strategies, based on not only ISS, are called for improving V2I WiFi performance.




Fig. 12. Correlation between driving speed and connectivity duration: associations are partitioned into 2.5 km/h bins according to their average speed, where each bin contains associations with the average speed in that range. Red lines, green diamonds, lower box edges, upper box edges, lower whiskers, and upper whiskers are medians, means, $Q1$ (the 1st quartile), $Q3$ (the 3rd quartile), $Q1 - 1.5 \times IQR$, and $Q3 + 1.5 \times IQR$ respectively, where IQR (Interquartile Range) equals to $Q3 - Q1$. R denotes the Pearson's Correlation Coefficient. As expected, slower speeds are likely to cause longer connections. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 13. Correlation between ISS and connectivity duration: associations are partitioned into 2.5 dBm bins according to their ISS. We draw the box-whisker plot as in Fig. 12. R denotes the Pearson's Correlation Coefficient.

WiFi frequency: The last factor we take into account is the frequency of the WiFi APs. Since Free has only 2.4 GHz APs available in Paris, we omit it. Table 4 reports the correlations between the APs frequency and the connectivity duration. Interestingly, from our results, the duration of the connectivity offered by 5 GHz APs is sometimes better or at least comparable to the one offered by 2.4 GHz APs. This is likely because 2.4 GHz band has fewer independent channels than 5 GHz

Table 4
Correlation between WiFi frequency and connectivity duration.

City	WiFi frequency	Link-layer duration (s)	
		Mean	Stdev
Bologna	5 GHz	26.22	39.62
	2.4 GHz	14.54	19.56
LA	5 GHz	30.09	33.38
	2.4 GHz	26.94	30.67
Macao	5 GHz	20.22	40.33
	2.4 GHz	24.52	41.57

Table 5
IP address roaming support in different WiFi networks.

City	Wifi network	IP changed?	IP duration (s)
Bologna	ALMAWIFI	✓	687 ± 1153
	Fastweb	✗	N/A
	EmiliaWifi	✓	184 ± 594
LA	Spectrum	✓	1644 ± 2653
	CableWifi	✓	3105 ± 5097
	Eduroam	✓	106 ± 156
Macao	CTM	✗	N/A
Paris	Free	✓	216 ± 384

band, thus, leading to severer radio interference, which further causes bursty beacon losses and interrupts connectivity.

Takeaway: (1) The AP selection algorithm of stock WiFi client implementations, which solely relies on ISS, is not sufficient when it comes to vehicular mobility. Applying location-aware AP selection approaches based on GPS history [58] is a good direction, which is likely to achieve better performance as of today since the provider-managed WiFi APs are very likely more stable and predictable than open APs; (2) 5 GHz APs often offer longer connectivity than 2.4 GHz APs in the V2I context likely due to the less severe radio interference.

6.4. IP address roaming support

Lastly, we investigated how the IP address (the local IP of the vehicle) assigned by a provider-managed WiFi network changes across different APs during a continuous drive. Usually, the IP address roaming (the ability of a host to retain the same IP address when roaming across APs) is not supported when a WiFi client switches to another operator's WiFi network unless the operators have a particular agreement. Hence, we focus on the IP address roaming inside WiFi networks operated by the same provider. Specifically, we investigate: (1) Can a vehicle keep its IP address unchanged during a continuous drive? (2) If not, for how long can a vehicle keep the same IP address?

Table 5 reports the data collected for every WiFi network.⁹ used in our measurement campaign. The column IP duration reports the average time during which vehicles keep an IP address unchanged during a continuous drive. The IP duration values shown are formatted as *average ± standard deviation*

From our results, as only two networks apparently support it, we could not find any strong evidence of provider-managed WiFi networks supporting, by default, IP address roaming. Hence, either applications

⁹ As we only get 12 associations from the TWCPasspoint WiFi network, we are not able to analyze it.




Fig. 14. CCDF of the TCP average goodput per association. All associations unable to establish TCP connections are excluded. Mean values (Download/Upload): Bologna 2.97/3.69, LA 7.62/3.21, Macao 8.66/13.36, Paris 0.70/0.74.

or transport layer protocols should take care of it as, otherwise, switching AP may break the connections established by connection-oriented protocol like TCP.

Takeaway: The IP address roaming in provider-managed WiFi networks is not a must implemented by operators by default. Therefore, vehicular application designers must bear the caveat in mind if they want to leverage provider-managed WiFi networks. That said, transport protocols like QUIC [35,36], Multipath QUIC [25,26], and Multipath TCP [24], which supports IP handover, thus, hiding the IP changes from applications, are good options for vehicular applications using provider-managed WiFi networks.

7. Transport performance analysis

In this section, we first present: (1) the average goodput a vehicle can experience from an association; (2) the amount of data that can be transmitted during a single association. These results in Section 7.1 demonstrate that today's provider-managed WiFi networks can provide decent performance for V2I communication. Then, in Section 7.2, we compare the results between different CCAs, and study the impact of *flow number* on TCP goodput. Finally, the impact of vehicle speed, *ISS*, and *WiFi frequency* is analyzed in Section 7.3.

7.1. Goodput and data volume

Here, we seek to understand the capabilities of today's provider-managed WiFi networks in the V2I context from a high-level point of view. Therefore, we do not distinguish between TCP sessions with different *flow number* and CCA. It is also noteworthy that applications also use various TCP configurations in reality.

Average TCP goodput: To support real world applications, a connection with decent goodput is fundamental. In Fig. 14 we plot the CCDF of the average goodput per association, for both upload and download. The results in Macao seems very promising: for every association able to establish TCP connections, the mean goodput is 13.36 Mbps in upload and 8.66 Mbps in download; with peaks up to 103.44 Mbps and 90.48 Mbps, respectively. Peak performances are likely to be achieved when vehicles, due to traffic jams or lights, stop at places with good APs. Further, while the results from LA and Bologna are also quite encouraging, TCP does not perform as well in Paris, with 90% of the associations providing less than 2 Mbps for both upload and download. Like in the other cities, also the CCDF of Paris is heavy-tailed, biased by the vehicle stops by good APs. Notice, the maximum average goodput both in upload and download is ~14 Mbps.

Data volume: To gauge clearer insights on the network capabilities in such a scenario, we consider the actual amount of data uploaded/downloaded per association. Fig. 15 shows the CCDF of the data volume transmitted during each association. Both in LA and Macao is possible to transmit a considerable amount of data during every association (the mean data volume downloaded per association is 45.22 MB and 40.31 MB, up to 1.04 GB and 2.26 GB, respectively). On the




Fig. 15. CCDF of the transfer size per association. All associations unable to establish TCP connections are excluded. Mean transfer size (Download/Upload): Bologna 12.46/14.67, LA 45.22/19.15, Macao 40.31/55.05, Paris 2.39/3.88.

other hand, on average, the amount of data uploaded per association is 19.15 MB in LA and 55.05 MB in Macao, with a maximum of 1.01 GB and 1.86 GB. Once again, in Paris, the results are less attractive, with only 3.88 MB uploaded and 2.39 MB downloaded on average. These figures are skewed by the outliers at the tail of the distribution.

The reasons for the discrepancies between Paris and the other cities could be multi-folds. Firstly, the AP density in Paris is considerably higher than in the other cities and, therefore, we could suffer from higher radio interference in Paris. Secondly, only 2.4 GHz APs are provided by Free. In general, 5 GHz APs offer better performance, as they have more independent channels. Lastly, the APs in Paris are shared between private customers and public users, with customers who are often prioritized by the providers. Besides, in the other cities, there are also dedicated APs, which are with better hardware and friendlier to public users.

Considering the total driving hours and the downloaded and uploaded data volume, our results from Bologna, LA, and Macao confirm that vehicles are likely able to download or upload around 1 GB (1.23 GB in Bologna, 1.44 GB in LA, and 0.73 GB in Macao) per hour by using stock TCP implementations, while on-the-go. Furthermore, this figure could improve with more APs coming into operation, and with a more capillary diffusion of 802.11ax APs. Moreover, as TCP is known to be inefficient when it comes to intermittent connectivity, this could get even better by using novel, alternative, transport protocols, such as QUIC [35,36].

Takeaway: Compared with the numbers provided by open WiFi networks in the past [3,27–31], the TCP goodput (~ 1 GB/h) offered by provider-managed WiFi networks is an order of magnitude higher, which shows the great potential of offloading data from cellular networks for connected vehicles.

7.2. Impact of CCA and flow number

Next, we compare the average TCP goodput per association of different TCP configurations to investigate the impact of CCA and flow number. It is noteworthy that the following analysis is based on the data collected from Macao, as it has more data for us to draw the relationships in a statistically significant way. Fig. 16 shows the CCDF of average TCP goodput per association under different TCP configurations.

First, from Fig. 16, we observe that no obvious differences between BBR and CUBIC, despite the superiority of BBR in conventional scenarios reported by previous studies [37]. The most likely reason is that BBR underestimates the round-trip propagation time (RT_{prop}) in such a dynamic environment, which leads to an inappropriate estimation of congestion window size, thus, counterbalancing the merits of BBR. We note that this phenomenon is also reported by a TCP measurement study in high-speed rail scenarios [59].

Another observation made from Fig. 16 is that multiple TCP flows often deliver better performance than a single TCP flow. The reasons are two folds: (1) uncongestional packet losses have a severer impact




Fig. 16. CCDF of average TCP goodput per association under different TCP configurations: S denotes single flow, and M denotes multiple flows.




Fig. 17. Correlation between vehicle speed and TCP goodput: associations are partitioned into 2.5 km/h bins according to their average vehicle speed. The box-whisker plot is drawn as in Fig. 12. R denotes the Pearson's Correlation Coefficient.

on single flow than multiple flows as it is likely that only a part of flows is affected in the case of multiple flows; (2) multiple flows are more competitive in bandwidth contention than single flow.

Takeaway: (1) BBR and CUBIC perform close to each other in terms of TCP goodput when used for V2I WiFi connections; several techniques [59–61] that compensates the underestimation of RT_{prop} may be applied to boost BBR's performance; (2) vehicular applications can benefit from the use of concurrent TCP flows when offloading traffic via provider-managed WiFi networks.

7.3. Impact of vehicle speed, ISS, and frequency

Now, we consider how vehicle speed, ISS, and frequency would impact TCP goodput. As in the previous subsection, we conduct our analysis based on the data collected from Macao due to its' larger scale.

First, the correlation between average vehicle speed per association and average TCP goodput per association is drawn in Fig. 17. We notice that, unlike the correlation between vehicle speed and connectivity duration in Fig. 12, the TCP goodput is not strongly related to vehicle speed: for download and upload, the R between vehicle speed and TCP goodput is only -0.12 and -0.08 respectively. On the other hand, we also acknowledge that the maximum goodput achieved in each speed bin tends to be higher when the vehicle speed is lower. We note that our results may be limited to the fact that the vehicle speeds were relatively slow in Macao, 0~35 km/h, due to the local regulations on driving speed in populated urban areas.

Then, we draw the correlation between ISS and average TCP goodput per association in Fig. 18. Similar to the correlation between ISS and connectivity duration in Fig. 13, we do not find any evidence indicating that stronger ISS leads to higher TCP goodput. The results again call for more sophisticated AP selection algorithms that consider more factors rather than just ISS.

Lastly, we investigate how different APs with different WiFi frequencies could perform. Table 6 summarizes the average goodput achieved from 5 GHz and 2.4 GHz APs in both upload and download scenarios. We can observe that 5 GHz APs provide higher TCP goodput than 2.4 GHz APs in the V2I context, which comes from the fact that the




Fig. 18. Correlation between *ISS* and TCP goodput: associations are partitioned into 2.5 dBm bins according to their *ISS*. The box-whisker plot is drawn as in Fig. 12. *R* denotes the Pearson's Correlation Coefficient.

Table 6
Correlation between *WiFi frequency* and TCP goodput.

Trans. direction	WiFi frequency	Goodput (Mbps)	
		Mean	Stdev
Upload	5 GHz	16.25	13.21
	2.4 GHz	4.33	5.63
Download	5 GHz	10.26	10.92
	2.4 GHz	3.71	4.22

5 GHz band has more independent channels, and, therefore, more chances to use 40 MHz channel and less radio interference than the 2.4 GHz band.

Takeaway: (1) TCP goodput is not correlated with *ISS*, which corroborates that *ISS* alone is not sufficient for the AP selection in the V2I scenario; (2) 5 GHz APs provide notably higher goodput than 2.4 GHz APs. Thus, for the operators who are willing to support V2I WiFi offloading, they may prioritize 5 GHz APs in their future WiFi deployment.

8. Discussion

While this work focuses on TCP due to its' predominance, we are aware that cutting-edge transport protocols may outperform TCP. QUIC [35,36] is already deployed by many companies [35,62,63] thanks to its' benefits in terms of 0-RTT connection establishment and better loss recovery mechanisms. Further, QUIC's connection migration mechanism enables connections to continue across different APs, even if IP roaming is not supported. We leave a study of how QUIC performs in our setup as future work.

As of today, given the WiFi promising bandwidth reported in our results, we can conclude V2I WiFi connectivity can benefit non-real-time bulk transfer applications. Yet, it is still not clear how applications requiring continuous Internet connectivity can benefit from intermittent, rapidly changing, V2I WiFi links. A possible direction is to develop applications on top of multi-path transport protocols [24–26], combining cellular and WiFi links. This would allow applications to dynamically adapt to network changes, leveraging cellular links for key urgent information (e.g. malfunctioning, etc.), whilst offloading over WiFi, saving the cost of cellular networks, otherwise. However, extra care is paramount while designing the multi-path scheduler [64], to avoid incurring additional Head-of-Line (HoL) blocking caused by injecting data into the ever-changing WiFi link, as it would impair application-perceived performance [33,65–67].

Although the full-fledged 5G Standalone (SA) connectivity [68] is poised to enhance V2I communication performance, recent measurements [69] have shown that current deployments of 5G Non-standalone (NSA) services are still far from being satisfactory. Ultimately, we

believe that in the transition to, or even in the world of the full-fledged 5G SA deployments, V2I WiFi offloading via prevalent provider-managed WiFi networks can complement 5G due to its' low-cost and performance.

Currently, the key industries of connected vehicles are also actively working on DSRC and C-V2X to provide network access for connected vehicles [70]. However, the applications that DSRC and C-V2X target are mostly real-time ones, such as sharing sensor data with other vehicles and roadside facilities in real-time to improve road safety. The non-real-time bulk transfer applications, e.g. infotainment, FOTA updates, and telemetry reports, still mainly rely on cellular networks, which put a lot of pressure on cellular infrastructures and add costs to users and car vendors. From this point of view, WiFi offloading is appealing because of the low price and decent bandwidth of WiFi networks. On the other side, WiFi networks can also complement DSRC and C-V2X for better application performance. For instance, WiFi offloading can be used for the control traffic in DSRC, in order to improve throughput [17].

9. Conclusion

In this paper, we revisit the opportunities of using WiFi offloading for V2I communication, and how this has changed over the last decade. Through an extensive measurement study across 4 cities located in 3 continents over 22 months, we confirm the feasibility of using provider-managed WiFi networks for V2I WiFi offloading, and, quantitatively analyze the transmission capability offered by such networks. Our data show that current provider-managed WiFi networks are much superior to open WiFi networks in the past for V2I WiFi offloading, especially when it comes to application goodput. Moreover, several implications on transport protocols and WiFi deployments are made from our measurement study. We hope that our work will shed light on the use of provider-managed WiFi networks for V2I WiFi offloading in the future.

CRedit authorship contribution statement

Furong Yang: Writing – original draft, Software, Methodology, Conceptualization, Data curation, Investigation. **Andrea Ferlini:** Investigation, Software, Validation, Writing – review & editing. **Davide Aguiari:** Visualization, Investigation, Writing – review & editing. **Davide Pesavento:** Software, Writing – review & editing. **Rita Tse:** Resources, Writing – review & editing. **Suman Banerjee:** Writing – review & editing. **Gaogang Xie:** Writing – review & editing. **Giovanni Pau:** Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We are grateful to the review panel for their insightful comments and feedback, which have greatly improved the quality of this paper. We appreciate Luca De Mori and Émilie Mespoulhes for their contribution to the experiments in Paris. We also thank Zhenyu Li for his constructive suggestions during the revision stage of this paper. This work is supported in part by the National Key R&D Program of China (Grant No. 2019YFB1802800), the National Natural Science Foundation of China (Grant No. 61725206), and the CAS-Austria Joint Project (171111KYSB20200001), the Macao Polytechnic Institute project (RP/ESCA-01/2020), the US National Science Foundation grants CNS-1719336, CNS-1647152, CNS-1629833, and CNS-2003129 and an award from the US Department of Commerce with award number 70NANB21H043.

References

- [1] P. Rodriguez, R. Chakravorty, J. Chesterfield, I. Pratt, S. Banerjee, Mar: A commuter router infrastructure for the mobile internet, in: Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services, 2004, pp. 217–230.
- [2] J. Hare, L. Hartung, S. Banerjee, Beyond deployments and testbeds: experiences with public usage on vehicular WiFi hotspots, in: Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services, 2012, pp. 393–406.
- [3] 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 Press, San Francisco, California, USA, 2008, p. 199, <http://dx.doi.org/10.1145/1409944.1409968>, URL <http://portal.acm.org/citation.cfm?doid=1409944.1409968>.
- [4] C.E. Andrade, S.D. Byers, V. Gopalakrishnan, E. Halepovic, D.J. Poole, L.K. Tran, C.T. Volinsky, Connected cars in cellular network: a measurement study, in: Proceedings of the 2017 Internet Measurement Conference, IMC'17, ACM Press, London, United Kingdom, 2017, pp. 235–241, <http://dx.doi.org/10.1145/3131365.3131403>, URL <http://dl.acm.org/citation.cfm?doid=3131365.3131403>.
- [5] U.S. DOT - NHTSA, Automated driving systems, 2017, URL https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/13069a-ads2.0_090617_v9a_tag.pdf.
- [6] K. Winter, For self-driving cars, there's big meaning behind one big number: 4 Terabytes, 2019, URL <https://newsroom.intel.de/news-releases/self-driving-cars-theres-big-meaning-behind-one-big-number-4-terabytes/>.
- [7] Tesla autopilot, 2020, https://en.wikipedia.org/wiki/Tesla_Autopilot (Accessed 16 August 2020).
- [8] K.J. Ryan, Tesla explains how a.I. Is making its self-driving cars smarter, 2019, URL <https://www.inc.com/kevin-j-ryan/how-tesla-is-using-ai-to-make-self-driving-cars-smarter.html>.
- [9] B. Marr, The amazing ways tesla is using artificial intelligence and big data, 2019, URL <https://bernardmarr.com/default.asp?contentID=1251>.
- [10] R. Waters, Why tesla is taking a different approach to self-driving cars, 2019, URL <https://www.ft.com/content/b423fe54-6cfa-11e9-80c7-60ee536e681d>.
- [11] L. Yvkoff, BMW rolls-out its intelligent personal assistant feature via over-the-air update, 2019, URL <https://www.forbes.com/sites/lianeyvkoff/2019/05/30/bmw-rolls-out-its-intelligent-personal-assistant-feature-via-over-the-air-update/#5030b92f5894>.
- [12] K.L. Man, T. Ting, T. Krilavicius, K. Wan, C. Chen, J. Chang, S. Poon, Towards a hybrid approach to SoC estimation for a smart battery management system (BMS) and battery supported cyber-physical systems (CPS), in: 2012 2nd Baltic Congress on Future Internet Communications, IEEE, 2012, pp. 113–116.
- [13] F. Ahmad, H. Qiu, R. Eells, F. Bai, R. Govindan, CarMap: Fast 3D feature map updates for automobiles, 2020, pp. 1063–1081, URL <https://www.usenix.org/conference/nsdi20/presentation/ahmad>.
- [14] Q. Liu, T. Han, J.L. Xie, B. Kim, LiveMap: Real-time dynamic map in automotive edge computing, in: IEEE INFOCOM 2021 - IEEE Conference on Computer Communications, IEEE, Vancouver, BC, Canada, 2021, pp. 1–10, <http://dx.doi.org/10.1109/INFOCOM42981.2021.9488872>, URL <https://ieeexplore.ieee.org/document/9488872/>.
- [15] H. Zhou, X. Chen, S. He, J. Chen, J. Wu, DRAIM: A novel delay-constraint and reverse auction-based incentive mechanism for WiFi offloading, IEEE J. Sel. Areas Commun. 38 (4) (2020) 711–722, <http://dx.doi.org/10.1109/JSAC.2020.2971871>, Conference Name: IEEE Journal on Selected Areas in Communications.
- [16] N. Lu, N. Cheng, N. Zhang, X. Shen, J.W. Mark, Connected vehicles: Solutions and challenges, IEEE Internet Things J. 1 (4) (2014) 289–299, <http://dx.doi.org/10.1109/JIOT.2014.2327587>, URL <http://ieeexplore.ieee.org/document/6823640/>.
- [17] H. Qin, B. Cao, W. Chen, X. Xiao, Y. Peng, WiFi-assisted control traffic offloading in vehicular environments, IEEE Trans. Veh. Technol. 69 (4) (2020) 4542–4547, <http://dx.doi.org/10.1109/TVT.2020.2975605>, Conference Name: IEEE Transactions on Vehicular Technology.
- [18] OpenSignal - mobile network experience report October 2019, 2020, <https://www.opensignal.com/reports/2019/10/italy/mobile-network-experiencem> (Accessed 06 February 2020).
- [19] M. Shafi, A.F. Molisch, P.J. Smith, T. Haustein, P. Zhu, P. De Silva, F. Tufvesson, A. Benjebbour, G. Wunder, 5G: A tutorial overview of standards, trials, challenges, deployment, and practice, IEEE J. Sel. Areas Commun. 35 (6) (2017) 1201–1221.
- [20] A. Festag, Standards for vehicular communication—from IEEE 802.11 p to 5G, E I Elektrotech. Inf. tech. 132 (7) (2015) 409–416.
- [21] D.X. Yang, Y. Guo, O. Aboul-Magd, 802.11 Ax: The coming new WLAN system with more than 4x MAC throughput enhancement, in: 2017 IEEE 86th Vehicular Technology Conference, VTC-Fall, IEEE, 2017, pp. 1–5.
- [22] M. Ayyash, H. Elgala, A. Kheishah, V. Jungnickel, T. Little, S. Shao, M. Rahaim, D. Schulz, J. Hilt, R. Freund, Coexistence of WiFi and LiFi toward 5G: concepts, opportunities, and challenges, IEEE Commun. Mag. 54 (2) (2016) 64–71.
- [23] Y. Mehmood, N. Haider, M. Imran, A. Timm-Giel, M. Guizani, M2M communications in 5G: state-of-the-art architecture, recent advances, and research challenges, IEEE Commun. Mag. 55 (9) (2017) 194–201.
- [24] A. Ford, C. Raiciu, M.J. Handley, O. Bonaventure, TCP Extensions for Multipath Operation with Multiple Addresses, Request for Comments, (6824) RFC Editor, 2013, <http://dx.doi.org/10.17487/RFC6824>, URL <https://rfc-editor.org/rfc/rfc6824.txt>.
- [25] Q. De Coninck, O. Bonaventure, Multipath QUIC: Design and evaluation, in: Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies, in: CoNEXT '17, Association for Computing Machinery, New York, NY, USA, 2017, pp. 160–166, <http://dx.doi.org/10.1145/3143361.3143370>.
- [26] Y. Liu, Y. Ma, Q.D. Coninck, O. Bonaventure, C. Huitema, M. Kühlewind, Multipath extension for quic, (draft-lmbdnhk-quic-multipath-00) Internet Engineering Task Force, 2021, p. 26, <https://datatracker.ietf.org/doc/html/draft-lmbdnhk-quic-multipath-00>, Work in Progress.
- [27] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, S. Madden, A measurement study of vehicular internet access using in situ Wi-Fi networks, in: Proceedings of the 12th Annual International Conference on Mobile Computing and Networking, MobiCom '06, ACM Press, Los Angeles, CA, USA, 2006, p. 50, <http://dx.doi.org/10.1145/1161089.1161097>, URL <http://portal.acm.org/citation.cfm?doid=1161089.1161097>.
- [28] P. Deshpande, X. Hou, S.R. Das, Performance comparison of 3G and metro-scale wifi for vehicular network access, in: Proceedings of the 10th Annual Conference on Internet Measurement, IMC '10, ACM Press, Melbourne, Australia, 2010, p. 301, <http://dx.doi.org/10.1145/1879141.1879180>, URL <http://portal.acm.org/citation.cfm?doid=1879141.1879180>.
- [29] H. Soroush, N. Banerjee, M. Corner, B. Levine, B. Lynn, A retrospective look at the UMass DOME mobile testbed, ACM SIGMOBILE Mob. Comput. Commun. Rev. 15 (4) (2012) 2–15.
- [30] H. Soroush, P. Gilbert, N. Banerjee, B.N. Levine, Concurrent Wi-Fi for mobile users: Analysis and measurements, in: CoNEXT, 2011, p. 12.
- [31] A. Balasubramanian, R. Mahajan, A. Venkataramani, Augmenting mobile 3G using WiFi, in: Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services, MobiSys '10, ACM Press, San Francisco, California, USA, 2010, p. 209, <http://dx.doi.org/10.1145/1814433.1814456>, URL <http://portal.acm.org/citation.cfm?doid=1814433.1814456>.
- [32] D. Pesavento, G. Grassi, G. Pau, P. Bahl, S. Fdida, Demo: Car-Fi: Opportunistic V2I by Exploiting Dual-Access Wi-Fi Networks, 2015, pp. 173–175, <http://dx.doi.org/10.1145/2789168.2789171>.
- [33] H. Lee, J. Flinn, B. Tonshal, RAVEN: Improving interactive latency for the connected car, in: Proceedings of the 24th Annual International Conference on Mobile Computing and Networking, MobiCom '18, ACM Press, New Delhi, India, 2018, pp. 557–572, <http://dx.doi.org/10.1145/3241539.3241571>, URL <http://dl.acm.org/citation.cfm?doid=3241539.3241571>.
- [34] R. Van Nee, V.K. Jones, G. Awatar, A. Van Zelst, J. Gardner, G. Steele, The 802.11n MIMO-OFDM standard for wireless LAN and beyond, Wirel. Pers. Commun. 37 (3–4) (2006) 445–453, <http://dx.doi.org/10.1007/s11277-006-9073-2>, URL <http://link.springer.com/10.1007/s11277-006-9073-2>.
- [35] A. Langley, A. Riddoch, A. Wilk, A. Vicente, C. Krasic, D. Zhang, F. Yang, F. Kouranov, I. Swett, J. Iyengar, J. Bailey, J. Dorfman, J. Roskind, J. Kulik, P. Westin, R. Tennen, R. Shade, R. Hamilton, V. Vasiliyev, W.-T. Chang, Z. Shi, The QUIC transport protocol: Design and internet-scale deployment, in: Proceedings of the Conference of the ACM Special Interest Group on Data Communication, in: SIGCOMM '17, Association for Computing Machinery, New York, NY, USA, 2017, pp. 183–196, <http://dx.doi.org/10.1145/3098822.3098842>, URL <https://doi.org/10.1145/3098822.3098842>.
- [36] IETF quic working group. Quic working group website, 2021, <https://quicwg.org/> (Accessed 12 January 2021).
- [37] N. Cardwell, Y. Cheng, C.S. Gunn, S.H. Yeganeh, V. Jacobson, BBR: Congestion-based congestion control, 2016.
- [38] S. Ha, I. Rhee, L. Xu, CUBIC: a new TCP-friendly high-speed TCP variant, Oper. Syst. Rev. 42 (5) (2008) 64–74.
- [39] J. Ott, D. Kutscher, Drive-thru internet: IEEE 802.11b for "automobile" users, in: IEEE INFOCOM 2004, Vol. 1, IEEE, Hong Kong, PR China, 2004, pp. 362–373, <http://dx.doi.org/10.1109/INFCOM.2004.1354509>, URL <http://ieeexplore.ieee.org/document/1354509/>.
- [40] R. Gass, J. Scott, C. Diot, Measurements of in-motion 802.11 networking, in: Seventh IEEE Workshop on Mobile Computing Systems & Applications, WMCSA'06, IEEE, Orcas Island, WA, USA, 2006, pp. 69–74, <http://dx.doi.org/10.1109/WMCSPA.2006.14>, URL <http://ieeexplore.ieee.org/document/1691716/>.
- [41] D. Hadaller, S. Keshav, T. Brecht, S. Agarwal, Vehicular opportunistic communication under the microscope, in: Proceedings of the 5th International Conference on Mobile Systems, Applications and Services, MobiSys '07, ACM Press, San Juan, Puerto Rico, 2007, p. 206, <http://dx.doi.org/10.1145/1247660.1247685>, URL <http://portal.acm.org/citation.cfm?doid=1247660.1247685>.
- [42] 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 Press, San Diego, California, USA, 2007, p. 321, <http://dx.doi.org/10.1145/1298306.1298351>, URL <http://portal.acm.org/citation.cfm?doid=1298306.1298351>.

- [43] J. Ott, D. Kutscher, A disconnection-tolerant transport for drive-thru internet environments, in: Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies, Vol. 3, IEEE, Miami, FL, USA, 2005, pp. 1849–1862, <http://dx.doi.org/10.1109/INFCOM.2005.1498464>, URL <http://ieeexplore.ieee.org/document/1498464/>.
- [44] X. Hou, P. Deshpande, S.R. Das, Moving bits from 3G to metro-scale WiFi for vehicular network access: An integrated transport layer solution, in: 2011 19th IEEE International Conference on Network Protocols, 2011, pp. 353–362.
- [45] Z. Song, L. Shangguan, K. Jamieson, Wi-Fi goes to town: Rapid picocell switching for wireless transit networks, in: Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '17, ACM Press, Los Angeles, CA, USA, 2017, pp. 322–334, <http://dx.doi.org/10.1145/3098822.3098846>, URL <http://dl.acm.org/citation.cfm?doi=3098822.3098846>.
- [46] A. Giannoulis, M. Fiore, E.W. Knightly, Supporting vehicular mobility in urban multi-hop wireless networks, in: Proceeding of the 6th International Conference on Mobile Systems, Applications, and Services, MobiSys '08, ACM Press, Breckenridge, CO, USA, 2008, p. 54, <http://dx.doi.org/10.1145/1378600.1378608>, URL <http://portal.acm.org/citation.cfm?doi=1378600.1378608>.
- [47] A. Balasubramanian, R. Mahajan, A. Venkataramani, B.N. Levine, J. Zahorjan, Interactive WiFi Connectivity For Moving Vehicles, SIGCOMM, 2008, p. 12.
- [48] 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 Press, Wroclaw, Poland, 2009, p. 263, <http://dx.doi.org/10.1145/1555816.1555843>, URL <http://portal.acm.org/citation.cfm?doi=1555816.1555843>.
- [49] D. Stanley, J.R. Walker, D.B.D.A. Ph.D., Extensible Authentication Protocol (EAP) Method Requirements for Wireless LANs, in: Request for Comments, RFC Editor, 2005, <http://dx.doi.org/10.17487/RFC4017>, URL <https://rfc-editor.org/rfc/rfc4017.txt> RFC 4017.
- [50] R. Droms, Dynamic Host Configuration Protocol, in: Request for Comments, RFC Editor, 1997, <http://dx.doi.org/10.17487/RFC2131>, URL <https://rfc-editor.org/rfc/rfc2131.txt> RFC 2131.
- [51] protocol/EAP PEAP URL <https://wiki.freeradius.org/protocol/EAP-PEAP>.
- [52] J.A. Salowey, H. Haverinen, Extensible authentication protocol method for global system for mobile communications (GSM) subscriber identity modules (EAP-SIM), RFC Editor, 2006, <http://dx.doi.org/10.17487/RFC4186>, URL <https://rfc-editor.org/rfc/rfc4186.txt>.
- [53] Wi-fi protected access, 2021, Wikipedia URL https://en.wikipedia.org/w/index.php?title=Wi-Fi_Protected_Access&oldid=1031381932 Page Version ID: 1031381932.
- [54] A. Mishra, X. Sun, A. Jain, S. Pande, R. Joshi, B. Leong, The great internet TCP congestion control census, Proc. ACM Meas. Anal. Comput. Syst. 3 (3) (2019) 1–24.
- [55] Free France, 2020, <https://www.free.fr/freebox> (Accessed 02 February 2020).
- [56] S. Zrelli, N. Okabe, Y. Shinoda, EAP-kerberos: A low latency EAP authentication method for faster Handoffs in wireless access networks, IEICE Trans. Inf. Syst. E95-D (2) (2012) 490–502, <http://dx.doi.org/10.1587/transinf.E95.D.490>, URL http://www.jstage.jst.go.jp/article/transinf/E95.D/2/E95.D.2_490/article.
- [57] Low latency of re-authentication during handover - re-authentication using a signed token in heterogeneous wireless access networks, in: Proceedings of the 10th International Conference on Signal Processing and Multimedia Applications and 10th International Conference on Wireless Information Networks and Systems, SCITEPRESS - Science and Technology Publications, Reykjavik, Iceland, 2013, pp. 248–254, <http://dx.doi.org/10.5220/0004608602480254>, URL <http://www.scitepress.org/DigitalLibrary/Link.aspx?doi=10.5220/0004608602480254>.
- [58] A.J. Nicholson, B.D. Noble, BreadCrumbs: forecasting mobile connectivity, in: Proceedings of the 14th ACM International Conference on Mobile Computing and Networking, MobiCom '08, ACM Press, San Francisco, California, USA, 2008, p. 46, <http://dx.doi.org/10.1145/1409944.1409952>, URL <http://portal.acm.org/citation.cfm?doi=1409944.1409952>.
- [59] J. Wang, Y. Zheng, Y. Ni, C. Xu, F. Qian, W. Li, W. Jiang, Y. Cheng, Z. Cheng, Y. Li, X. Xie, Y. Sun, Z. Wang, An active-passive measurement study of TCP performance over LTE on high-speed rails, in: The 25th Annual International Conference on Mobile Computing and Networking, ACM, Los Cabos Mexico, 2019, pp. 1–16, <http://dx.doi.org/10.1145/3300061.3300123>, URL <https://dl.acm.org/doi/10.1145/3300061.3300123>.
- [60] F. Yang, Q. Wu, Z. Li, Y. Liu, G. Pau, G. Xie, BBRV2+: Towards balancing aggressiveness and fairness with delay-based bandwidth probing, 2021, [arXiv: 2107.03057](https://arxiv.org/abs/2107.03057) [Cs], URL <https://arxiv.org/abs/2107.03057>.
- [61] R. Kumar, A. Koutsafis, F. Fund, G. Naik, P. Liu, Y. Liu, S. Panwar, TCP BBR for ultra-low latency networking: Challenges, analysis, and solutions, in: 2019 IFIP Networking Conference, IFIP Networking, 2019, pp. 1–9, <http://dx.doi.org/10.23919/IFIPNetworking.2019.8816856>, ISSN: 1861-2288.
- [62] Facebook mvfst client and server implementation of IETF quic protocol in C++ by facebook, 2021, <https://github.com/facebookincubator/mvfst> (Accessed 12 January 2021).
- [63] CloudFlare quiche - IETF QUIC transport protocol and HTTP/3 implementation by CloudFlare, 2020, <https://github.com/cloudflare/quiche> (Accessed 12 January 2021).
- [64] O. Bonaventure, M. Piraux, Q.D. Coninck, M. Baerts, C. Paasch, M. Amend, Multipath schedulers, (draft-bonaventure-icrg-schedulers-02) Internet Engineering Task Force, 2021, <https://datatracker.ietf.org/doc/html/draft-bonaventure-icrg-schedulers-02>, Work in Progress.
- [65] S. Ferlin, O. Alay, O. Mehani, R. Boreli, BLEST: Blocking estimation-based MPTCP scheduler for heterogeneous networks, in: 2016 IFIP Networking Conference (IFIP Networking) and Workshops, 2016, pp. 431–439, <http://dx.doi.org/10.1109/IFIPNetworking.2016.7497206>.
- [66] C. Raiciu, C. Paasch, S. Barre, A. Ford, M. Honda, F. Duchene, O. Bonaventure, M. Handley, How hard can it be? Designing and implementing a deployable multipath TCP, in: 9th USENIX Symposium on Networked Systems Design and Implementation, NSDI 12, USENIX Association, San Jose, CA, 2012, pp. 399–412, URL <https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/raiciu>.
- [67] Y.E. Guo, A. Nikraves, Z.M. Mao, F. Qian, S. Sen, Accelerating multipath transport through balanced subflow completion, in: Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking, in: MobiCom '17, Association for Computing Machinery, New York, NY, USA, 2017, pp. 141–153, <http://dx.doi.org/10.1145/3117811.3117829>, URL <https://doi.org/10.1145/3117811.3117829>.
- [68] S. Abdel Hakeem, 5G-V2x: Standardization, architecture, use cases, network-slicing, and edge-computing, Wirel. Netw. 26 (2020) <http://dx.doi.org/10.1007/s11276-020-02419-8>.
- [69] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, H. Ma, Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption, in: Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication, in: SIGCOMM '20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 479–494, <http://dx.doi.org/10.1145/3387514.3405882>, URL <https://doi.org/10.1145/3387514.3405882>.
- [70] L. Zhao, X. Li, B. Gu, Z. Zhou, S. Mumtaz, V. Frascolla, H. Gacanin, M.I. Ashraf, J. Rodriguez, M. Yang, S. Al-Rubaye, Vehicular communications: Standardization and open issues, IEEE Commun. Stand. Mag. 2 (4) (2018) 74–80, <http://dx.doi.org/10.1109/MCOMSTD.2018.1800027>, Conference Name: IEEE Communications Standards Magazine.



Furong Yang is currently a Ph.D. candidate at ICT/CAS. He is also in a joint doctoral program between the University of Chinese Academy of Sciences and Sorbonne University since 2018. His main research interests lie in TCP congestion control, multipath transport protocol, and Internet measurement.



Andrea Ferlini is a Ph.D. Candidate, under the supervision of Prof Cecilia Mascolo, in the Systems Research Group (SGR) of the University of Cambridge and a graduate student at Churchill College. His research interests range between wearables, applied deep learning on constrained devices, privacy, and mobile health. Before starting his Ph.D. at the University of Cambridge, he was a research engineer at LIP6 (Sorbonne Université, Campus Pierre et Marie Curie). Precisely, he was working with Prof Giovanni Pau on connected vehicles and wireless technologies.



Davide Aguiari is a research fellow at CIRI-ICT in Bologna where he develops a prediction and monitoring battery system for connected electric vehicles. He received the Computer Science master degree at Alma Mater Studiorum — University of Bologna in 2017, the Computer Science Ph.D. degree from Sorbonne Université and the Electronic and Computer Engineering Ph.D. from University of Cagliari, both in 2021 thanks to a joint Ph.D. programme. His current research interests include Internet of Things, connected Electric Vehicles and Big Data management.



Davide Pesavento is currently a researcher at NIST's Advanced Network Technologies Division. Prior to that, he received his M.S and B.S. degree in Computer Science from the University of Padova in Italy. His research interests focus on Name Data Networking and Vehicular networking.



Rita Tse received the Math/Computer Science degree from UCLA and the Ph.D. degree from Hull University in 2004. She has been the Coordinator of the Computer Studies/Computing Program at the Macao Polytechnic Institute from 2006 to 2016. She is now the Director of School of Public Administration. Her current research interests include Ubiquitous Computing, Urban Sensing, Social Network-based Sensing, Casino Gaming Systems, and Course Design and Development.



Suman Banerjee (Fellow IEEE/ACM) received the undergraduate degree from IIT Kanpur in computer science and engineering and was a gold medalist in his graduating class, and the MS and Ph.D. degrees from the University of Maryland. His Ph.D. dissertation was the university's nomination for the ACM Doctoral Dissertation Award. He is a professor in computer sciences with UW-Madison where he is the founding director of the WiNGS laboratory which broadly focuses on research in wireless and mobile networking systems. While at Wisconsin, he received the CAREER award from the US National Science Foundation and the



Gaogang Xie (Member, IEEE) received the B.S. degree in physics and the M.S. and Ph.D. degrees in computer science from Hunan University in 1996, 1999, and 2002, respectively. He is currently a Professor with the Computer Network Information Center (CNIC), Chinese Academy of Sciences (CAS), and the University of Chinese Academy of Sciences (UCAS), and the Vice President of CNIC. His research interests include Internet architecture, packet processing and forwarding, and Internet measurement.



Giovanni Pau is a Full Professor in the Department of Computer Science and Engineering (DISI) at the University of Bologna; He also serves as an Adjunct Professor at UCLA. He holds the Italian Laurea in Computer Science and a Ph.D. in Computer Engineering awarded by the University of Bologna in 1998 and 2002 respectively. Before joining the University of Bologna, Dr. Pau served as ATOS/Renault smart mobility chair of excellence Professor at the Sorbonne Université, Paris France. Dr. Pau's core research interests are in Network Systems with a focus on Vehicular Networks and pervasive mobile sensor systems.