IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 9 December 2022, accepted 5 January 2023, date of publication 11 January 2023, date of current version 19 January 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3236100

RESEARCH ARTICLE

Millimeter-Wave Feasibility in 5G Backhaul: A Cross-Layer Analysis of Blockage Impact

TÂNIA FERREIRA[1], ALEXANDRE FIGUEIREDO[1], DUARTE RAPOSO[1], MIGUEL LUÍS[1,2], PEDRO RITO[1], AND SUSANA SARGENTO[1,3]

[1]Instituto de Telecomunicações, 3810-193 Aveiro, Portugal
[2]Instituto Superior de Engenharia de Lisboa (ISEL), Instituto Politécnico de Lisboa, 1959-007 Lisbon, Portugal
[3]Departamento de Electrónica, Telecomunicações e Informática (DETI), University of Aveiro, 3810-193 Aveiro, Portugal

Corresponding author: Tânia Ferreira (tania.s.ferreira@av.it.pt)

This work was supported in part by the European Regional Development Fund (FEDER), through the Regional Operational Programme of Centre (CENTRO 2020) of the Portugal 2020 framework; and in part by the National Public Funds through FCT I.P. (OE) under the MIT Portugal Program, Project SNOB-5G (CENTRO-01-0247-FEDER-045929), under Grant 045929.

ABSTRACT The wireless backhaul has become a key enabler for 5G technology by presenting a cost-effective and scalable alternative to the typical fiber backhaul. WiGig protocols, such as IEEE 802.11ad and later IEEE 802.11ay, have been considered for backhaul connectivity of 5G mobile networks, thanks to the availability of high bandwidths capable of achieving fiber-like data rates. However, this band suffers from high propagation loss that can only be compensated using highly directional antennas, making mmWave links more susceptible to blockage and errors. Thus, to effectively evaluate the viability of WiGig-based technologies in wireless backhaul scenarios, it is crucial to characterize the impact of obstruction across the different network layers. This article presents an extensive measurement campaign and cross-layer analysis of physical (PHY), medium access control (MAC), and transport layers metrics measured for outdoor WiGig-based hardware submitted to short-term and long-term blockage. This study found that maintaining constant and higher modulation and coding schemes (MCSs) in long-term blocked channels may induce packet errors as high as 100%, round-trip-time (RTTs) that can be in the order of a few seconds, and packet losses as high as 90%. Even dynamically adjusting the MCS, the performance can be highly degraded. This effect was exacerbated in short-term links, as they suffered from more extreme MCS changes upon sudden obstructions. Temporary line of sight (LOS) obstruction was shown to cause maximum delays of half a second and a PER of around 20%; in more extreme cases, it has even led to temporary link failures.

INDEX TERMS Milimeter-wave communications, backhaul network, WiGig, IEEE 802.11ad, milimeter-wave reliability, blockage, cross-layer analysis.

I. INTRODUCTION

In recent years, a large amount of research emerged regarding mmWave technologies and their unique applications. In particular, the IEEE 802.11ad standard (WiGig) was proposed for 5G small cell wireless backhauling [1], [2], [3], [4] due to the support for data rates as high as 2.5 Gbps (in the Single Carrier mode). This multi-gigabit wireless standard operates in the unlicensed 7 GHz band located between 57 – 64 GHz. It features fast session transfer for seamless data fallback and data increase between 60 GHz and 2.4/5 GHz PHYs, and MAC and PHY layer enhancements which include support for directional antennas, beamforming, and spatial reuse techniques. The standard novelty resides in the new Directional Multi-Gigabit (DMG) PHY layer that, based on multiple-antennas beamforming techniques, is capable of creating highly directional beams that concentrate the power in a specific direction, which ultimately compensates the increased attenuation in the 60 GHz band [5]. Therefore, a reliable communication link can only be established the beamforming training, which points the stations' transmitting and receiving patterns to each other. In the next evolution,

The associate editor coordinating the review of this manuscript and approving it for publication was Xiaolong Li.

the IEEE 802.11ay, a Task Group [6] was formed to propose the next generation wireless multi-gigabit standard. Built on top of its predecessor IEEE 802.11ad, it defines new PHY and MAC layers that enable WiFi devices to achieve theoretical peak data rates up to 100 Gbps, - by employing MIMO techniques -, while also ensuring backward compatibility and coexistence with 802.11ad DMG STAs operating in the same band [7]. Still, at the time of writing, few IEEE 802.11ay-compliant COTS devices are available on the market [8].

While the highly directional links employed in WiGig technologies address the path loss issue associated with higher frequencies, it also makes the 60 GHz communication more vulnerable to blockage and errors. MmWave signals suffer less diffraction than microwave signals, which once again makes them much more vulnerable to blockage (20 – 30 dB in temporary or long-term link outage [9]). Also, relying on non-line-of-sight (NLOS) transmission is not always an option (received signal decreases 10 – 20 dB using first-order and second-order reflected signals [10]). Therefore, a sudden obstruction of the line of sight path is enough to cause a significant decrease in the achievable throughput and lead to delays above acceptable limits for critical applications. In the worst cases, link obstruction may even cause link failure. Thus, to explore the viability of WiGig-based technologies in wireless backhaul scenarios, it is crucial to understand the real impact caused by different types of obstruction in the major link Key Performance Indicators (KPIs).

This work aims to characterize the impact caused by long and short-term obstruction based on multi-layer data acquired in a millimeter-wave outdoor deployment. To the best of our knowledge, this is the first study on the real effects of obstruction in mmWave links. The remaining of this paper is organized as follows. Section II presents a comprehensive study of previous experimental works made at the 60 GHz band. Section III details the deployed backhaul infrastructure and describes the multi-layer data acquisition process. Section IV presents a detailed analysis of short-term and long-term obstruction's impact on the metrics at multiple layers. At last, section V summarizes our findings and discusses future research directions.

II. RELATED WORK

Up to now, experimental studies on the 60 GHz band used wireless docking stations with radio chipsets that Wilocity [11] typically manufactures. Moreover, they mainly focused on indoor measurements campaigns [12], [13], [14], [15], [16], [17].

In 2015, [12] investigated the feasibility of full-scale indoor 60 GHz networks by conducting several measurements using 802.11ad-compliant COTS devices in an indoor office environment. The authors conducted an analysis using the received signal strength indicator (RSSI), PHY data rate, and TCP throughput to characterize the impact of the transmitter-receiver distance, antenna height, location, and orientation on the system's performance. They found that RSSI is a weak indicator for PHY data and TCP throughput and that PHY data rate is not always a good indicator of higher-layer performance. Furthermore, they showed that static human blockage had a more significant impact on the TCP throughput than transient blockage.

In 2017, [17] proposed a novel approach of leveraging WiFi channel fingerprints for localizing WiGig coverage area along with reducing its beamforming training (BT) complexity over conventional exhaustive search BT. For that purpose, they linked WiFi fingerprints with Wigig information (RSS readings and the best beam sector IDs) collected in a non-obstructed meeting room environment with Wigig-compliant APs.

Later in 2018, Kacou et. al. [13] presented an experiment to compute the channel's root mean square (RMS) delay spread and associated path loss under LOS and NLOS conditions. Their measurements showed that their 802.11ad equipment has a 65 – 70 dB of propagation loss under normal operation, and a 25 dB of additional path loss, in the case of obstruction by a load-bearing wall or by a wooden cabinet close. Additionally, they found that the TCP throughput decreases from 1.6 Gbps to 2.5 Gbps in LOS conditions to 0 to 1.6 Gbps, when faced with static obstruction. Besides, it was observed that four people blocking a LOS path still achieved a throughput greater than 1 Gbps in indoor office environments.

In 2020, authors of [15] reported the ability to differentiate with 96% accuracy temporary from permanent blockage caused by humans in an indoor environment by analyzing the signal quality and TCP throughput. To do so, they deployed an experimental testbed with an INTEL WiGig sink chipset W13100 as the access point (AP) and the INTEL tri-band wireless-AC 17265 wireless cards as the wireless station (STA). However, similar to the previously listed works, the blockage characterization is limited to a small set of metrics, including the TCP throughput, RSSI, and PHY data rate. The work in [14] accesses lower layer metrics (MCS, data rate, and CRC failures), and presents a preliminary study of their behavior for blocked and non-blocked indoor mobility scenarios that use a COTS 802.11ad device. More specifically, they describe the impact that transmit-receiver distance has at multiple layers, and define the influence of transient blockage at different speeds on the signal-to-noise ratio (SNR) and the beam adaptation process. Their results show that transient blockage may induce long-term suboptimal beam alignment on a 60 GHz link, causing a significant performance decrease. Later in 2019, [16] extended this work by providing a more in-depth analysis of static and transient blockage impact. This analysis shows how different types of blockage affect the TCP throughput, the number of control packets, the MCS, and the beam pattern changes.

The works reported in [14] and [16] present huge advances, but some limitations are still present from a higher-layer perspective, as they only focus on TCP throughput. The performance of critical mmWave indicators -delay and packet loss- are excluded from their analysis. Furthermore, the measurement campaign was conducted indoors with equipment

whose characteristics differ entirely from those intended for outdoors.

The few outdoor studies available evaluate the feasibility of indoor' IEEE 802.11ad solutions on outdoor environments submitted to static and transient blockage. In fact, to the best of our knowledge, only [18] and most recently [19] address outdoor scenarios. Zhu et al. presented in [18] the results of measurements (range, attenuation and interference) using COTS Wilocity radios and a horn antenna emulating a 10 × 10 antenna array to evaluate the feasibility of 60 GHz picocells in outdoor scenarios. However, this work still has two major limitations. First, the results of simulating a large array using a horn antenna can only be considered valid in an environment with no interference. Second, the cross-layer performance analysis only includes RSSI and TCP throughput. Most recently, Tran et al. [19] conducted an outdoor experiment campaign in a customized mmWave Meshed Backhaul that employs WiGig to demonstrate the potential of Mobile Edge Computing (MEC) technology in this type of network. The tests performed evaluated the UDP throughput and end-to-end delay achieved in non-obstructed single-hop and multi-hop backhaul links with a constant MCS of 9. So far, this is the only work that proposes a physical deployment for a mmWave backhaul.

In contrast, this work presents the most extensive PHY, MAC and transport layer measurement campaign. Furthermore, it will join the analysis of metrics from the transport (TCP throughput, number of TCP retransmissions, packet loss rate, RTT) and PHY/MAC layers (RCPI, SNR, Packet Error Rate, Beam Index TX, Beam Index RX, and L2 data rate) to characterize static and transient blockage on 60 GHz links. More specifically, this work proposes a cross-layer analysis of the impact caused by long-term and short-term blockage on an experimental WiGig-based outdoor network intended for 5G backhauling.

Table 1 provides a brief summary of the aspects discussed for the current experimental work using the IEEE 802.11ad standard. Specifically, we compare these works to our own in terms of equipment, test environment, metrics and layers analyzed, types of blockage investigated and point out their limitations.

III. EXPERIMENTAL SETUP

This section describes the outdoor deployment infrastructure and the data collection process.

A. NETWORK OVERVIEW

The deployed mmWave backhaul network comprises three CCS Metnet 60 GHz nodes [20], similar to the ones presented in Figure 1. These nodes were designed for deployment on urban street furniture, such as lamp posts. However, in this work, a smaller-scale network was deployed on the rooftop of our premises' building, - at the height of about 1.5 meters and with a 50 meter distance between the nodes-, to ensure that the

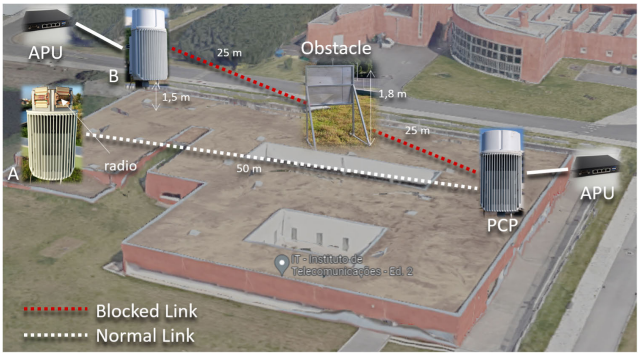

FIGURE 1. Nodes and topology of the deployed outdoor testbed, presenting the long-term blockage scenario.

tests were performed under a closed and controlled outdoor environment. Figure 1 shows the testbed and equipment used throughout this work. In this infrastructure, two accelerated processing units (APUs) were placed at each node end to inject traffic into the mmWave connections.

The network nodes leverage the standardized IEEE 802.11ad (WiGig) technology to form a wireless 5G meshed backhaul capable of accommodating data rates up to 4.6 Gbps in the Single Carrier mode. Each node has four radio modules covering a 90° range, overlapping to cover a 300° horizontal field of view. Each radio employs a 19 dBi (37dBm EIRP) beamforming steerable antenna (a 16 × 2 element array), which allows coping with the propagation losses by establishing highly directional links between stations (STAs). This beamforming capability allows the equipment to divide the 300° horizontal field into a set of 64 discrete sectors (with a 5° horizontal beamwidth) that can be used to send and receive the signal toward a specific direction.

The system also takes one step further by engineering a multi-point to multi-point topology. Each node can cover up to 500 meters depending on the location, availability, and capacity. The deployed network adopts an architecture where multiple small cell base stations (SBSs) can be connected to the core network through fiber connections at Point of Presence (PoP) locations. These fiber-enabled nodes participate in an election process to select one network coordinator, known as the Personal Basic Service Set (PBSS) coordinator or PCP/AP. The remaining nodes, A and B, are remote and thus can only gain access to the core network through node PCP.

The network also features an Element Management System (EMS) based on HTTPS and SNMP protocols, which provides a graphical overview of the nodes' current state and the mmWave links. In addition, the EMS GUI displays the real-time MAC layer metrics (e.g., RCPI, SNR, MCS, PER) collected for each mesh link. A more detailed description of the operation of the deployed testbed, as well as some other additional real-life measurements, can be found in our previous published work [21].

TABLE 1. Comparison of existing experimental measurement work for WiGig-based networks.

Work	Date	DUTs	Environment	Metrics	Scenarios Investigated	Limitations
[12]	2015	Dell Latitude E420 laptop with wil6210 (STA) Dell Wireless Dock D5000 (AP)	Indoor	TCP throughput	Normal Transient human blockage Static human blockage	Lacks extensive PHY and MAC layer cross-analysis
[17]	2017	Costumized WiGig AP and STA prototypes Commercial Wi-Fi APs	Indoor	RSS Best Beam Sector IDs	Normal	
[13]	2018	Qualcomm 802.11ad mini PCIe cards	Indoor	TCP throughput	Normal Static obstruction from common construction materials Transient human blockage	
[15]	2020	INTEL tri-band wireless-AC 17265 wireless cards (STA) INTEL WiGig sink chipset W13100 (AP)	Indoor	TCP throughput RSSI	Normal Transient human blockage	
[14]	2016	D5000 wireless docking station (AP) Dell E7440 laptop (STA)	Indoor	AGC attenuation MCS, data rate, SNR, CRC failure rate and beam adaptation	Normal Transient blockage at different speeds	
[16]	2019	D5000 wireless docking station (AP) Dell E7440 laptop (STA)	Indoor	TCP throughput No. of sent control packets, MCS Beam pattern	Normal Static Blockage Transient blockage at different speeds	
[18]	2014	D5000 wireless docking station (AP) Dell 6430U laptop (STA)	Outdoor	RSS TCP throughput	Normal Static blockage from common construction materials Transient human blockage	Lacks extensive PHY and MAC layer cross-analysis Uses a 60 GHz horn antenna to simulate a 10x10 phased-array
[19]	2019	Custom setup based on Panasonic radio transceivers	Outdoor	Latency	Normal	Lacks extensive PHY and MAC layer cross-analysis Does not study the impact of transient and human blockage
This work	2022	CCS Metnet 60G nodes	Outdoor	RSSI, SNR PER, MCS Beam Index TX Beam Index RX L2 data rate RTT, PL TCP throughput TCP retransmissions	Normal Static blockage from metal object Transient human obstruction	Does not include UDP metrics

B. MULTIPLE-LAYER PROBING

The physical layer metrics, accessible through the EMS, were collected and stored inside a local Mongo database collection using a python script that sends HTTP requests to the monitoring interface API. This database collection is composed of entries that display the instantaneous KPI values for a specific path, presented in Table 2.

The network metrics were collected simultaneously with the physical ones using a custom *iperf3*-based python probe. The values assumed for the different TCP session-related parameters, such as timeout and window sizes, are the iperf3 defaults. The probe developed comprises two main components: an application server and the client. For the client part, a python script was developed, which instantiates an *iperf3* client whose job is to collect TCP metrics (configured to use half of the Ethernet capacity, 500 Mbps). The metrics collected are the received and transmitted data rate, the number of retransmissions and the lost packets. Additionally, the *pingparser* library was used to parse the output of the ping request to obtain additional network metrics, such as the RTT-based metrics, the number of transmitted, received and lost packets, and the duplicated packets related metrics (calculated by averaging the values measured during ten packets). The TCP-based and Ping metrics were also loaded to the same Mongo database collection that stored the PHY and MAC layer metrics.

By default, the client uploads to the server installed at node PCP. However, the client located at the node end could also be configured to operate in reverse mode, making it download from the server instead. Depending on the configured mode, the *iperf3* traffic flows more in one direction than the other, i.e., either the uplink or the downlink.

TABLE 2. Physical level metrics collected for dataset generation, and data acquisition scenarios.

Physical-layer metrics	
SNR (dB)	(Signal to Noise Ratio), used to select the best tx and rx beams/sectors during beamforming training
PER (%)	(Packet Error Rate), represents the ratio between the number of packets with errors after FEC and the sum of tx packets
MCS TX and MCS RX	(Modulation and Coding Scheme), used during transmission in one of the directions (uplink/downlink)
Beam Index TX (Min/Max/Mean)	The maximum, minimum and average beam ID used for tx, respectively, in a one second window
Beam Index RX (Min/Max/Mean)	The maximum, minimum and average beam ID used for rx, respectively, in a one second window
Mbps TX (Max) and Mbps TX	The peak and mean data rate in Mbps
Data acquisition scenarios	
Normal	LOS between each pair of tested STAs, for the entire test duration (15 minutes)
Short-term blockage	Temporary LOS blockage throughout the entire experiment duration (5 minutes). Data was obtained by having 1 and 2 people at a time, moving freely during 5 minutes, across the path formed by two STAs
Long-term blockage	Continuous LOS blockage for the entire experiment (15 minutes), with a metallic object between each STA pair, as shown in Figure 1.

C. PROBING VERIFICATION

To characterize the impact of obstruction on real mmWave backhaul links, the physical and network metrics were compared to distinguish the network behavior under distinct blockage events. For that purpose, the multi-layer metrics were collected for each possible path for the two pairs of STAs and for each MCS mode (constant [1, 3, 5, 7 and 9] and dynamic/automatic) under three scenarios. These scenarios were defined according to the type and duration of the blockage event crossing the LOS path, presented in Table 2.

Note that the term path refers to the physical link between a specific pair of STAs. For instance, a link between node PCP and B would ideally have eight paths to test using the four radios on each side if correct signal reception was always guaranteed. Also note that, due to the nature of the short-term blockage scenario data acquisition, it was unrealistic to collect data for all possible paths with all MCS modes for the same 15 minutes mark used in the other scenarios. So instead, tests were only performed for a smaller set of paths and only using the automatic MCS for a five minute-duration.

IV. MULTI-LAYER DATA ANALYSIS

The multi-layer metrics collected for each experiment (i.e., for each MCS and path combination) for the three scenarios were analyzed to characterize the impact of long-term and short-term blockage. The results from the analysis are presented in the following two sections. The first section presents the behavior of the physical and MAC-level KPIs under short-term and long-term blockage, while the second section details the impact of blockage in the transport and network-layer metrics.

A. MAC AND PHY LAYERS

By jointly analyzing the metrics collected at the physical layer for different paths, several conclusions can be drawn. The next section provides insights into the most relevant physical and MAC-level KPIs: SNR, PER, and MCS. Table 3 displays the statistics of the low-level metrics measured for the worst-performing link between stations PCP and A, under normal operation, long-term blockage (identified as ''LT''), and short-term blockage (identified as ''ST''), which served as the basis for the analysis.

1) SNR

This section discusses the results obtained concerning the SNR. Some considerations are presented:

- A link suffering from either long/short-term blockage has a lower average SNR when compared with the non-blocked scenario;
- The mean SNR of the uplink/downlink of a given path may not be similar, even in a non-blocked scenario, since they use different beams for transmitting/receiving to avoid interference. Thus, each direction has its communication channel, which is likely to be subject to different conditions;
- The SNR of two path directions suffering from long-term blockage may not experience the same average drop, even if both uplink/downlink have similar values under normal operation. Depending on the location and rotation at which the obstacle is placed, one communication channel direction may be affected more than the other;
- The average SNR drop is not the same for all long-term blocked links since different paths (formed by different radios) cover different regions of space. Thus, placing an obstacle in the LOS connection formed by two radios facing each other is more likely to affect the signal quality than in paths formed by radios on different sides;
- Under short-term blockage, the uplink and downlink formed between a pair of STAs followed a similar

TABLE 3. Statistics of the measured physical and MAC-level KPIs in the worst-performing link.

MCS Mode	Scenario	SNR				PER (%)				MCS	
		Mean	std dev	Min	Max	Mean	Std dev	Min	Max	Range	Mean
1	Normal	18,9	1,6	15,5	22,8	0	0	0	0		
	LT blockage	4,2	2,0	0,3	8,8	1,3	5,1	0	46,7		
3	Normal	18,8	1,6	16,0	22,8	0	0	0	0		
	LT blockage	6,8	1,9	1,00	9,3	10,2	24,7	0	100		
5	Normal	19,1	2,1	14,3	22,8	0	0	0	0		
	LT blockage	6,5	1,5	3,5	9,8	49,1	43,2	0	100		
7	Normal	19,2	1,0	15,8	22,0	0,3	0	0	1,8		
	LT blockage	5,5	1,1	3,0	9,00	66,8	33,4	0	100		
Auto	Normal	18,6	1,4	16,3	21,8	1,6	1,4	0	19,0	8 – 9	9
	LT blockage	9,8	2,0	6,3	13,8	6,4	9,3	0	64,7	1 – 6	3
	ST blockage	15,4	4,5	-2,3	21,0	2,4	4,8	0	88,2	1 – 9	8

behavior, apart from the maximum value reached that is close to what was registered in LOS conditions. In both directions, the LOS blockage occurred at the same instants, although with different intensities;

- In some cases, the SNR drops were more pronounced when the link was short-term blocked than when it was long-term blocked, as there is a higher chance of achieving complete beam obstruction.

Figures 2(a) and 2(b), obtained for paths PCP/A/2/2[1] and PCP/A/1/1, respectively, present examples in which all the conditions described above are met under the three scenarios: 1) the uplink and downlink's average SNR are different in all three scenarios; 2) both directions present a similar behavior for the same scenario since the instants at which transitions are observed are the same; and 3) in a short-term blocked link, the SNR alternates between a maximum value close to the non-obstructed scenario's average, and a minimum value that may be lower than the mean registered in the long-term blocked scenario. The minimum and maximum SNR are obtained in instants where the LOS path is obstructed or not, respectively. This behavior was also verified for the SNR of the remaining two paths (i.e., PCP/A/1/2 and PCP/A/2/1).

Because data is acquired with a 1-second periodicity, determining the exact value of the duration of the SNR drop caused by short-term human obstruction is not possible. Nevertheless, since each metric collected represents the average measured for that time window, we get an upper limit of 1 second. Therefore, it is observed that each LOS obstruction is reflected, in most cases, in the generation of a single point.

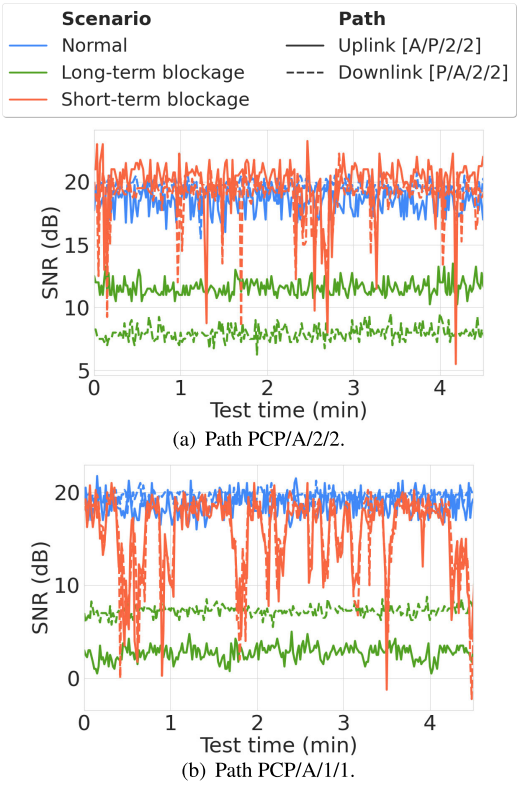

(a) Path PCP/A/2/2.

(b) Path PCP/A/1/1.

FIGURE 2. SNR of two mmWave paths under normal, long-term blocked and short-term blocked scenarios.

2) PACKET ERROR RATE

This section discusses the results obtained with respect to the PER. Some considerations are depicted as follows:

- The mean PER is higher for a long-term blocked link than for a link in normal operation when the MCS is constant, as there is consistently lower SNR caused by a lower RSSI;
- Even if a specific link has an SNR that is equal to or higher than the established threshold for a given MCS, achieving negligible PER is only possible if the RSSI is also superior to the minimum defined in the standard;
- Non-obstructed links operating with constant MCS mode present *negligible* PER (mostly under 1%);
- Non-obstructed links operating in the automatic mode present a more significant variability than in the constant mode (usually up to 10%) due to slight MCS adjustments;

[1]The term PCP/A/x/y identifies the link between node PCP and node A established by using radios y and x, respectively.

- Generally, a PER peak is detected, regardless of the scenario, at instants where a significant signal quality decrease occurred, caused by human obstruction or natural factors (e.g., strong wind);
- Setting a higher MCS in a long-term blocked link results in increased PER as the channel's margin decreases;
- The PER of links operating in the automatic MCS mode present more significant fluctuations when compared to the values recorded for the same scenario with the constant MCS. These fluctuations are due to the node's ability to vary the modulation and the redundancy introduced accordingly to the measured SNR and PER;
- As an extension of the previous statement, channels that face short-term obstructions present a larger PER variability when compared with either long-term blocked and non-obstructed channels. The PER remains smaller under long-term blockage because the channel fluctuations are slower than in short-term fluctuations;
- Under the short-term blockage, the PER is more likely to achieve its maximum value at times the LOS path is unexpectedly blocked and its minimum value at instants where there is LOS;
- Even in non-obstructed and long-term blocked scenarios, the PER may sporadically reach peak values comparable to those obtained with short-term blockage, due to natural causes;
- Rapid SNR variations do not always necessarily lead to higher error rates.

Figure 3 shows how the PER of the uplink A/PCP/2/1 is affected by the MCS increase under long-term blockage. This specific link had an average SNR of 7.5 dB, which in theory is enough to support modulations up to MCS 7 [22] as its threshold is set to 5.3 dB. However, this path only received on average −75 dBm (before antenna gain), which does not comply with the −68 dBm of MCS 7 nor the −69 dBm requirement of MCS 5. Thus, the uplink cannot maintain a reliable connection with either one of these schemes. Figure 4(a) confirms this conclusion, since it shows that long-term blockage has caused PERs to be within the 30 – 100 range for the MCS 7, being null in the non-obstructed environment.

On the other hand, in links operating with an automatic MCS mode, the PER fluctuations observed are smaller than those obtained with the constant MCSs. Auto mode allows the modulation scheme used by a link to be adjusted as soon as there is a significant drop in signal quality and/or when high PERs are experienced. This capability causes instants that follow PER peaks to exhibit lower PERs soon after the maximum, as shown in Figure 4(b) for link A/PCP/2/2. This image shows evidence that a link suffering from long-term blockage is more likely to present larger PER variability than when it is short-term blocked or operating in LOS conditions. Still, this relationship is not always true. Higher PERs depend on different aspects of the positioning of the transient obstacle, such as its rotation and distance from the transmitter or receiver. Therefore, it is not possible to estimate precisely

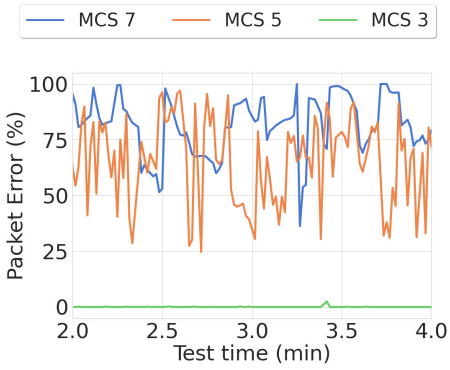

FIGURE 3. The influence of the MCS in the PER for link A/PCP/2/1 under long-term blockage.

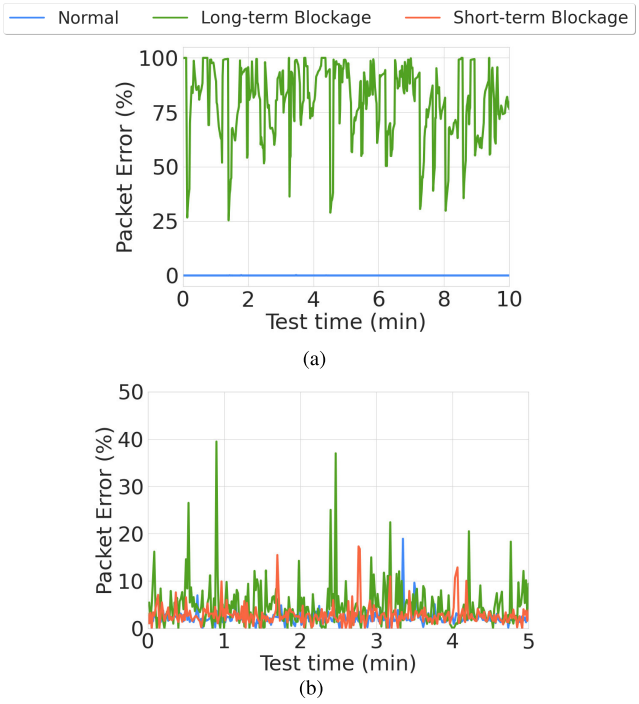

FIGURE 4. PER obtained for A/PCP/2/1 under normal and long-term blockage scenarios with a constant MCS of 7 (a). PER obtained for A/PCP/2/2 under normal, long-term and short-term blockage scenarios with automatic MCS (b).

how significant the impact of the two types of obstructions will be without knowing the specific details of the blocker's position.

As shown by the PER statistics of Table 3, long-term and short-term blocked scenarios can cause PERs as high as 64% and 88%, respectively. Furthermore, the experiments conducted have shown that a received power below the −70 dBm mark leads to PERs between 90 and 100% in low constant MCS. Note that, with −70 dBm, only an MCS equal to or below three can be supported with a small PER.

Finally, the relationship between signal quality drops and the error rate is further explored. Figure 5(a) demonstrates a situation in which the PER varied at a rate roughly equal to that of the SNR. In this figure, is clear that the error measured

for the uplink A/PCP/2/1 during short-term blockage is lower at instants where the signal quality is higher, and is higher when the signal quality drops. However, as shown in the experimental results of uplink A/PCP/2/2 in figure 5(b), SNR drops do not always translate to PER maximums. Therefore, achieving or not a maximum PER value depends more directly on whether the amount of redundancy introduced by the MCS at the time when the obstruction occurs is sufficient to compensate for the additional error introduced by natural or blockage-induced channel variations rather than the obstruction value itself.

3) MODULATION AND CODING SCHEME

This section discusses the results obtained when changing the modulation and coding scheme. Some considerations are depicted as follows:

- Increasing the MCS on a link enables it to transfer data at higher bit rates, if the received signal strength and SNR meet the outlined threshold values;
- In the automatic mode, the link MCS is dynamically adjusted according to both SNR and PER. Therefore, a decrease in the modulation supported may be triggered by either detecting a significant SNR loss, a high PER or both;
- As soon as the line of sight is no longer obstructed, the MCS resets to its pre-loss value.

The CCS mesh nodes used in the experimentation adjust the MCS according to both the SNR and PER measured at a given instant. Figure 6 shows the MCS adjustment made by the system to establish and maintain a link between node A and PCP (using radio 2 at both ends) that is submitted to short-term blockage. In this figure, it is clear that time windows with fewer MCS changes are the ones that present smaller SNR fluctuations around the mean value of 20 dB and a controlled small PER. In these windows, the MCS is set up to the maximum possible value, 9 for that connection, only changing occasionally to a smaller MCS, 8, due to sporadic small decreases in the signal quality.

Because dynamic MCS adjustment is performed on an instantaneous SNR and PER basis, it was difficult to pinpoint the MCS decrease to a single metric. Instead, a careful analysis had to be performed considering both metrics to understand the instantaneous MCS changes properly. Under this work context, the MCS changes were either due to one of three different conditions. First, the MCS decrease can be triggered due to a significant SNR drop alongside a high PER, as shown in instant t_6 of Figure 6, where the MCS decreases to 8. Second, the MCS can be changed solely due to the presence of high error rates, as shown in instants t_2 and t_4. A maximum PER of 18% and 16% was registered at t_2 and t_4, respectively, while the SNR value remained practically unaltered. Third, the MCS change can be triggered in some cases due to a signal quality decrease, which is the case of instant t_3. In t_3, the SNR has decreased from 20 to 9 dB with a 5% error rate.

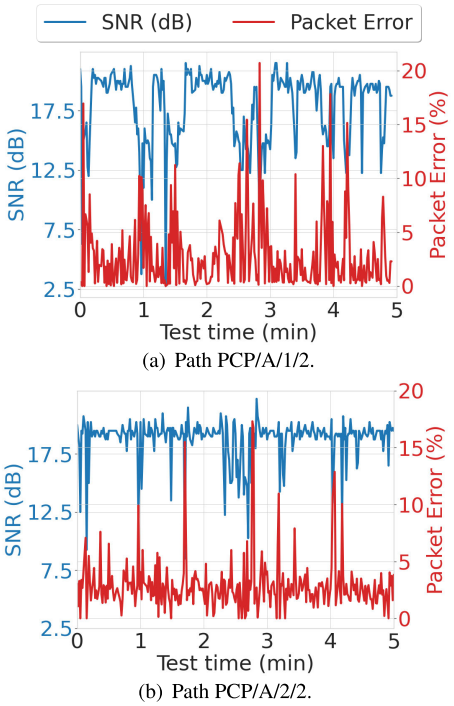

(a) Path PCP/A/1/2.

(b) Path PCP/A/2/2.

FIGURE 5. PER and SNR observed for two mmWave links operating in automatic MCS mode, under short-term blockage.

Still, it is only possible to delineate a general behavior of the dynamic adjustment algorithm. As shown in instant t_1 and t_5, not every decrease in signal quality and/or increase of the PER results in a MCS adjustment.

B. NETWORK AND TRANSPORT LAYERS

Once the impact of the main KPIs at the physical layer is understood, these findings are used to detect general patterns perceptible to the human eye at the transport layer, such as the behavior of packet delay and TCP performance in highly variable mmWave networks. The following section provides insights relating to the most relevant TCP and network-level KPIs: Round Trip Time (RTT), packet loss, TCP throughput, and TCP retransmissions. Table 4 presents the statistics of the TCP and network metrics measured for the worst-performing link between stations PCP and A, which served as the basis for the analysis.

1) RTT

To obtain the RTT metric, the *pingparser* library pings the client host for a given amount of time. In each time window, the 10 ICMP packets transmitted through the network were used to compute an average, maximum and minimum round-trip-time for each possible communication path and every scenario. The following discussion presents the findings from our analysis of the average and maximum RTT:

- A non-blocked link has an average RTT that decreases as the MCS index increases (for the constant mode). Even so, some small fluctuations around the mean are

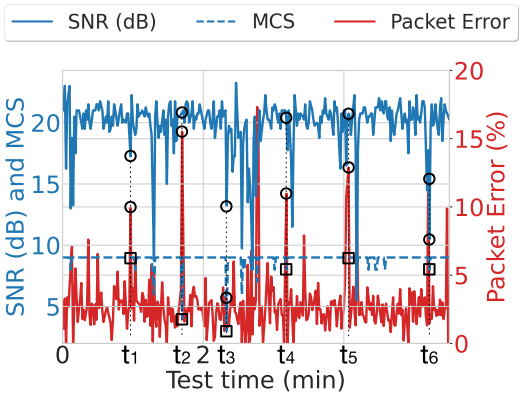

FIGURE 6. The influence of the SNR and PER in the dynamic MCS adjustment mechanism for link A/PCP/2/2 under short-term blockage.

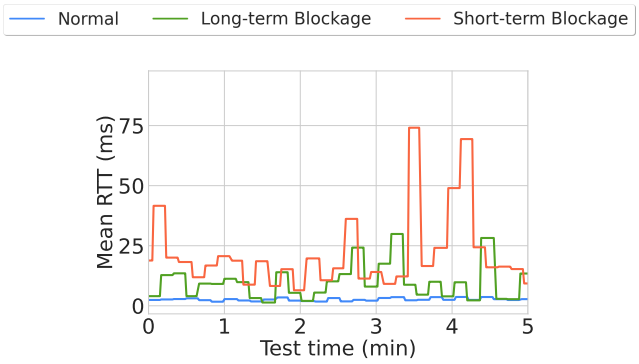

FIGURE 7. RTT in link PCP/A/2/2 under normal, long-term and short-term blocked scenarios for the automatic MCSs.

observed due to the varying nature of the mmWave channel;

- In a non-blocked link, the average RTT value measured with MCS 1 is ten times higher than any other MCS tested since the throughput traversing the network is at least 2.2 times lower than the one observed in other tested modulations;
- Links operating with higher MCSs in both modes, under a non-obstructed environment, have a maximum average RTT that does not surpass the 20-millisecond mark and a standard deviation that is always lower than three milliseconds;
- The average RTT of a long-term blocked link resembles the behavior observed for links operating in the non-blocked state when the PER is low;
- Long-term blockage and short-term blockage are observed to cause larger variability in the mean and maximum RTT value when compared with the normal scenario;
- A heavily blocked link can have a high RTT which can be in the order of one or a few seconds;
- Long-term and short-term blocked links using the automatic MCS mode face a significant increase in both the average and maximum RTT when compared with the values registered in LOS conditions;
- Links suffering from obstruction with high PERs present higher average RTTs for MCS with the lowest code rates as they are less likely to compensate for the error introduced by the channel;
- A less robust MCS will lead to an increase in the number of retransmissions;
- The mean RTT of all paths formed between a pair of stations have similar average values for the same MCS.

Figure 8 shows the RTT measured for the link A/PCP/2/2 in the non-blocked and long-term blocked scenarios. Through the analysis of Figure 8(a), it is possible to observe that, under normal conditions, the mean RTT is higher for lower MCSs. This is expected since lower MCSs have lower associated data rates [22]. On the other hand, under the long-term blockage, the link experiences higher delays in less robust MCS (i.e., MCS with lower code rates). For instance, in Figure 8(b), a maximum RTT of 150 ms is experienced with MCS 5 (which uses a 13/16 code rate), while at MCS 7 (code rate of 5/8) that value drops to 85 ms. MCS 1 (the most robust scheme) follows the same behavior observed in its non-blocked state since the PER is close to zero. This figure also shows that long-term obstruction causes high fluctuations of the mean RTT, even in the constant MCS mode. As shown in Figure 7, the RTT behavior is even more unpredictable in the case where the MCS is automatically adjusted for short-term blocked links as they are more prone to rapid changes in the MCS than long-term blocked links where the obstacle remains static during the entire test. This high RTT variability makes it difficult to estimate with certainty the extent of the obstruction's impact.

2) PACKET LOSS

In our specific context, two types of retransmission mechanisms can be identified that act at two different layers. First, retransmissions can be made at the MAC layer of IEEE 802.11ad. Second, TCP-based retransmissions can also be performed at the transport layer. The following discussion presents our findings from the analysis of the packet losses at the transport layer:

- A non-obstructed mmWave link using both automatic and constant MCS usually presents a low average packet loss (less than 4%), with only some sporadic losses occurring due to timeouts;
- TCP is not adapted to work with the new PHY and MAC layer mechanisms of IEEE 802.11ad;
- The higher the MCS, the more likely the link is to suffer from packet losses. In these conditions, a higher number of TCP retransmissions will be triggered to overcome the channel erasures;
- Short-term blocking a link may lead to packet losses up to 9% even under dynamic MCS adjustment.

As shown in Table 4, a non-obstructed link usually presents a very low packet loss for all supported MCSs. Still, some sporadic losses were observed in normal scenarios (e.g., the

TABLE 4. Statistics of the transport and network-layer metrics for the worst-performing link.

MCS Mode	Scenario	RTT Mean (ms)				RTT Max (ms)				Packet Loss				Throughput			
		Mean	Std dev	Min	Max	Mean	Std dev	Min	Max	Mean	Std dev	Min	Max	Mean	Std dev	Min	Max
1	Normal	41,7	7,9	18,7	58,8	79,3	4,6	58,4	88,7	5,0	2,2	1,5	**11,5**	**66,4**	19,8	17,3	101,6
	LT blockage	42,9	134,5	14,5	1325,0	112,5	478,6	36,3	4677,1	1,8	0,3	1.06	2,9	238,5	5,1	225,4	250,1
3	Normal	10,0	2,9	3,5	17,2	24,3	7,1	6,5	40,0	0,9	0,5	0,0	4,0	467,8	18,8	423,5	499,1
	LT blockage	167,9	230,3	0,7	819,3	402,4	491,6	1,2	1846,6	12,6	7,5	0,0	21,8	64,6	130,5	4,2	444,8
5	Normal	3,1	0,7	1,5	5,8	7,7	3,1	3,6	21,8	0,3	0,4	0.0	2,4	495,8	11,1	387,8	499,9
	LT blockage	208,2	138,3	5,7	554,1	684,9	323,7	13,5	1360,3	49,2	36,5	2,9	**100**	1,1	2,1	0	14,2
7	Normal	2,5	0,60	1,5	4,5	6,6	2,8	2,3	20,1	0,3	0,5	0.0	2,1	**496,2**	8,3	441,0	499,9
	LT blockage	3,8	9,0	0,5	87,3	16,5	47,3	0,6	401,3	97,5	9,1	37,5	**100**	**0**	0,0	0	0,1
Auto	Normal	2,7	1,3	2,7	12,3	8,7	12,8	2,3	110,2	0,5	0,5	0,0	**2,3**	495,2	7,3	452,6	499,9
	LT blockage	5,5	15,3	0,5	134,0	32,9	105,7	0,6	897,8	0,2	0,4	0,0	2.59	328,8	60,0	213,1	488
	LT blockage	25,3	6,5	21,2	93,2	35,5	66,9	2,8	529,8	1,1	0,9	0,0	**8,7**	273,1	82,2	75,1	365,4

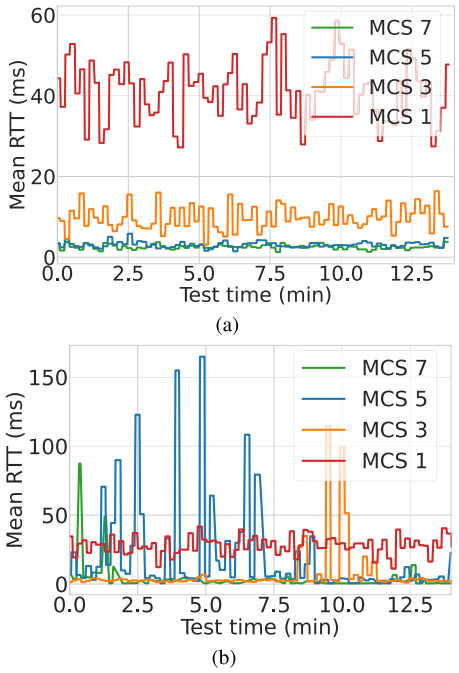

FIGURE 8. Average RTT observed for the uplink A/PCP/2/2 under normal operation (a) and long-term blockage (b).

11.5% losses measured in the normal scenario with MCS 1). These losses are likely caused by the timeouts when the channel is busy performing the 802.11ad MAC mechanisms (e.g., beamforming training), as the signal quality remains practically constant. As reported by the authors of [23], TCP does not change its timeout limit to account for these new operations and thus is forced to discard the unprocessed packets. Additionally, Table 4 shows that long-term blockage can cause average packets losses of about 98% when a high MCS is kept constant. Furthermore, it is demonstrated that short-term blockage can cause maximum packet losses up to 9% in the automatic MCS, contrasting with the 3% packet losses of non-obstructed and long-term blocked links. Figures 9 and 10 depict these aspects. As shown by figure 9, on a highly long-term obstructed link, the higher the MCS used, the higher the losses. Conversely, as shown by figure 10, when the MCS is dynamically adjusted according to the channel conditions, the losses incurred are lower than those registered with high and constant MCSs.

3) TCP THROUGHPUT

The following discussion presents the findings from our analysis of the TCP throughput:

- Increasing the MCS in a non-obstructed link leads to an increase in the maximum achievable throughput;
- Long-term blockage causes a high PER in higher index MCSs, which in turn decreases the achievable data rate;
- In the worst-case scenario, a heavily blocked link achieves rates as low as 1 Mbps when using constant MCSs. This means that the network's delivery capacity is so low that TCP reduced its congestion window from 500 to 1 Mbps, only allowing traffic generated by the retransmissions to circulate through the network;
- A short-term blocked link performing dynamic MCS adjustments has larger data rate fluctuations, as a sudden SNR decrease can make the modulation go from MCS 9 to MCS 1 within a second.

Table 4 shows the results obtained for the worst-performing link for each MCS mode. The following conclusions that validate the previous statements can be taken: 1) increasing the modulation scheme increases the achievable data rate, as we went from transmitting at a 67 Mbps rate in MCS 1 to 496 Mbps with MCS 7; 2) long-term blockage caused the average TCP throughput to drop significantly in higher MCSs. For instance, the throughput reduced from 496 Mbps in the non-blocked scenario to almost 0 Mbps for the MCS 7. This observation is further confirmed by the results of Figure 11(b) for link PCP/A/2/2, as both a constant MCS of 5 and 7 were not able to transmit more than a few Mbps under long-term blockage; and 3) the data rate of a link suffering from short-term blockage varies very quickly in the automatic MCS mode, as shown in Figure 12. This behavior is a result of the sudden changes in the MCSs that are made in an attempt to control the PER. Note that there was an irregularity observed

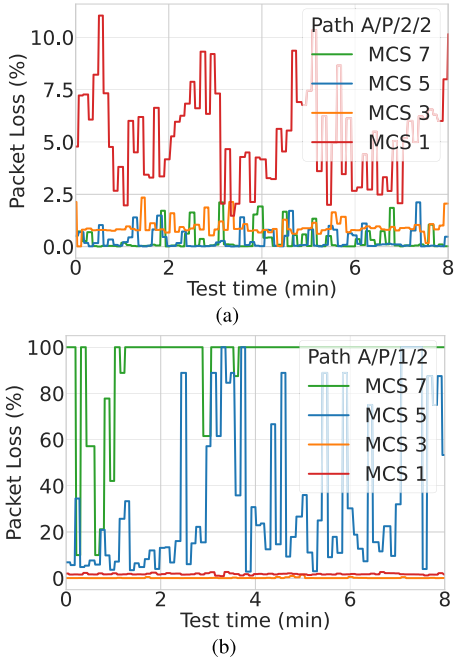

FIGURE 9. TCP packet losses for path PCP/A/2/2 in the normal scenario (a), and in the long-term blocked scenario (b) for the constant MCS mode.

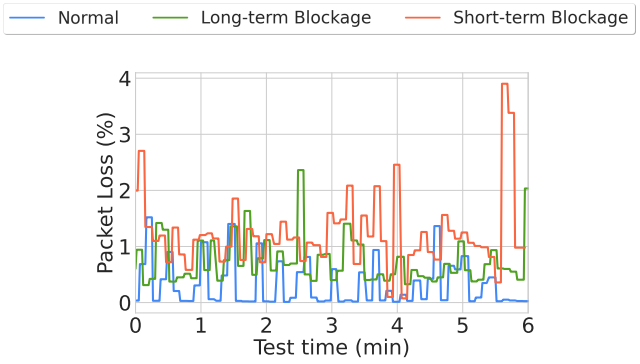

FIGURE 10. TCP packet losses for path PCP/A/2/2 under normal, long-term blocked and short-term blocked scenarios with the automatic MCS mode.

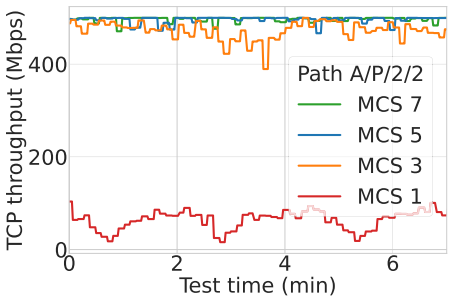

(a) PCP/A/2/2 in the non-blocked scenario.

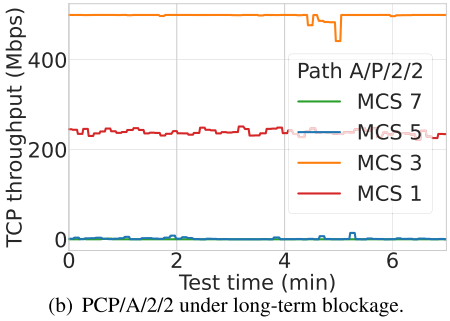

(b) PCP/A/2/2 under long-term blockage.

FIGURE 11. TCP throughput of the PCP/A/2/2 link operating under normal (a) and long-term blockage (b) conditions with constant MCSs.

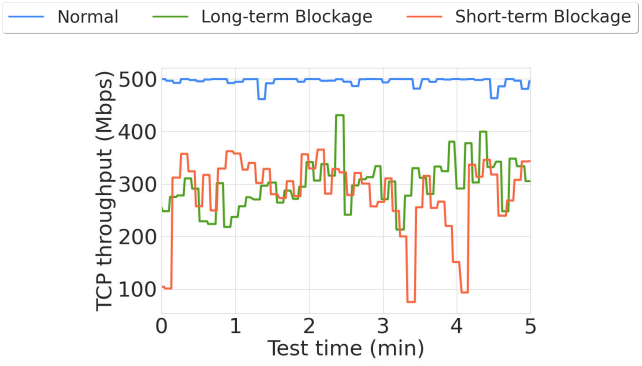

FIGURE 12. TCP throughput of the PCP/A/2/2 link operating under normal, long-term and short-term blockage conditions for the auto MCS.

for all paths tested. More specifically, the data rate increased with the introduction of the obstacle in the MCS of 1.

4) RETRANSMISSIONS

To properly compare the impact of each scenario in the number of TCP retransmissions, the TCP retransmits metric was normalized using the average TCP received rate. This enabled the removal of the effect of higher modulations on the transmitted bit rate and, consequently, in the total retransmissions. The following statements summarize the conclusions drawn from our analysis of the non-obstructed, long-term blocked and short-term blocked scenarios:

- The total number of TCP retransmissions is higher in non-blocked links (which support high MCS) than in blocked links;
- The normalized number of retransmissions of links under long-term blockage is much higher than in non-blocked environments for a constant MCS, which implies that, on average, more packet retransmissions are needed in order to deliver a given packet to its destination;
- In a long-term blocked link with high PER, setting higher MCSs increases the number of TCP retransmissions and thus, decreases the maximum achievable TCP throughput;
- Heavily blocked links in a long-term blocked environment have a limited delivery capacity when used with TCP as they mainly carry traffic from pending retransmissions instead of useful new information;
- Changing the MCS to a higher index MCS with lower code rates in heavily blocked scenarios increases the number of normalized TCP retransmissions;

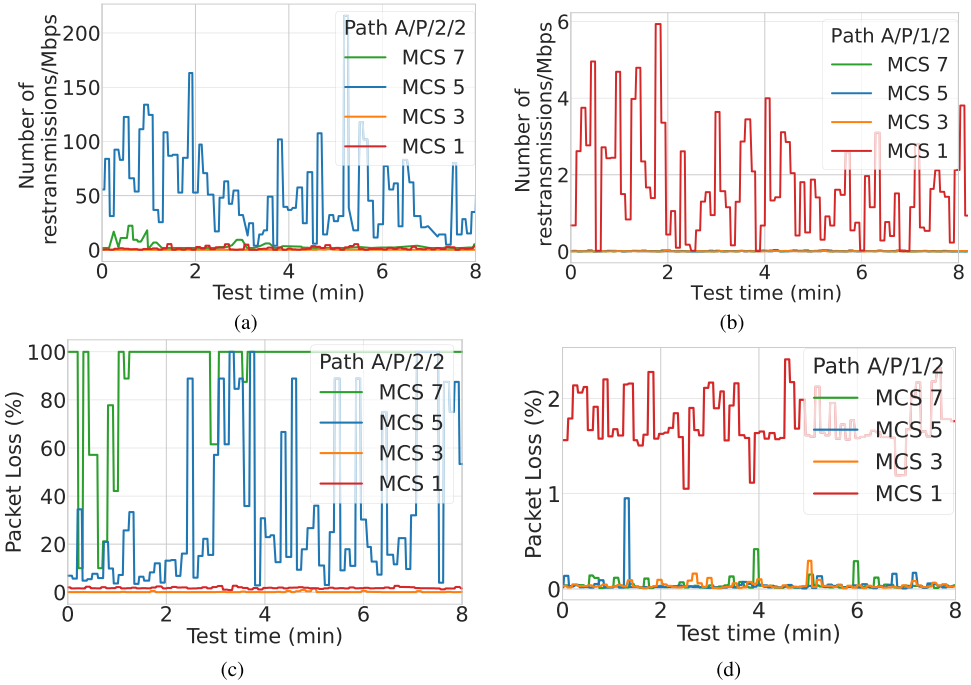

FIGURE 13. Number of TCP retransmissions per Mbps measured during long-term blockage in the uplink A/PCP/2/2 (a) and A/PCP/1/2 (b) for five different MCS levels. Packet losses measured for the uplink A/PCP/2/2 (c) and A/PCP/1/2 (d) for five constant MCS levels.

- Using MCS with higher coding rates makes it less likely that the TCP performs a high number of retransmissions in a link submitted to short-term blockage;
- The behavior of the number of TCP retransmissions in scenarios where the MCS is adjustable, especially in obstructed environments, is rather challenging to predict as the instants at which the increases in the number of retransmissions occur are not directly related to the instantaneous signal quality.

Figure 14 shows the impact that long-term blockage has on the total and the number of retransmissions per Mbps. At first glance, it may appear that a non-blocked link is worse than the long-term blocked one, but that is not true. Removing the effect of the maximum achievable throughput enables a correct analysis of the blockage impact. The truth is that a long-term blocked link, in which the MCS is constant, is unable to adjust its MCS whenever a high PER is detected. Thus, maintaining the MCS leads to a consistently high PER, which in turn forces TCP to trigger a high number of retransmissions to deliver the packets without error.

In principle, the higher the average packet losses, the higher the number of retransmissions needed to deliver a packet successfully. However, the collected metric is not referring to the average number of retransmitted packets per transmitted packet but rather to the mean number of retransmissions required to deliver the overall packets. On average, non-blocked links carry a much higher bit rate than a long-term blocked link, which means that a larger number of packets must be delivered to the receiving STA. Thus, as shown in Figure 14, under normal operation, a higher number of retransmissions every second is necessary to deliver the information, even though the number of retransmissions required per packet is lower. The latter observation provides evidence that CCS nodes most likely use an in-order delivery strategy for flow control purposes. Similar to that described in [23] for a TP-Link AP based on IEEE 802.11ad, the link layer cannot aggregate within the same block, frames from both retransmission and transmission queues. Therefore, the only type of traffic transiting in a link with pending retransmissions is the traffic generated from those retransmissions. Thus, new information can only circulate after the pending packets are successfully delivered. In our experiments, the transmitted rate was observed to be no more than 60 Mbps for heavily blocked links, suggesting that the network mainly carries traffic from pending retransmissions instead of useful new information.

Furthermore, the experimental results show that, even under normal conditions, TCP packet retransmissions still exist. According to [23], this is because the timeout value of frames that are waiting on the received queue is not adjusted to account for the time channel that was occupied performing frame transmissions of the IEEE 802.11ad mechanisms, such as beamforming and beam tracking.

Figure 13 shows the relationship between the normalized number of retransmissions and the packet losses for two links suffering from different levels of long-term obstruction. The A/PCP/2/2 link has lower SNR than the A/PCP/1/2 link in all constant MCS tested, resulting in a more heavily blocked

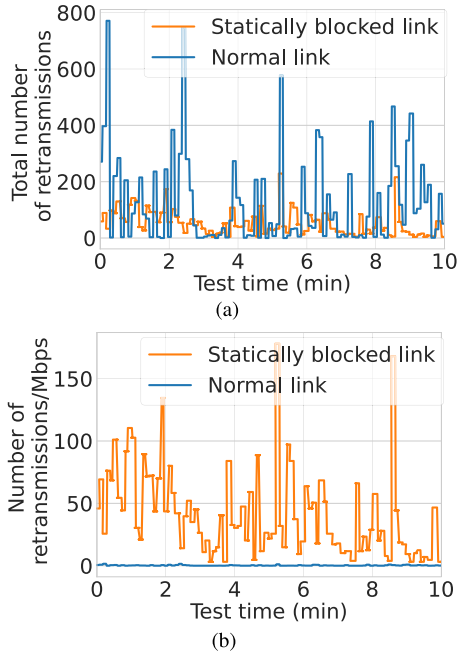

FIGURE 14. The total number of TCP retransmissions (a) and number of retransmission per Mbps (b), measured for PCP/A/2/2 under long-term blockage when using a constant MCS of 5.

channel. In a highly obstructed link like A/PCP/2/2, using higher MCSs such as 5 and 7 causes an average packet loss of 49% and 97%, respectively, and maximums that can reach the 100% mark (see table 4). The high losses trigger more TCP retransmissions. On the other hand, in a slightly blocked link similar to A/PCP/1/2, the number of retransmissions is kept low for all tested MCSs. Note that, in the A/PCP/1/2 link case, the maximum retransmissions do not surpass the six mark, whereas in the heavily obstructed link, the value is 33 times higher. This is because A/PCP/1/2 was observed to comply with the RSSI and SNR requirements of MCSs up to 7, which allowed it to operate with negligible PER for all MCS tested.

Furthermore, describing the behavior of the TCP retransmissions on a link using an adaptive MCS scheme is a complex task, as an increase or decrease in the instantaneous value of this metric can be traced to multiple sources. First, due to the MCS adjustment, it is expected that most of the time, a significant SNR decrease and/or high PER will cause the MCS to be changed to a more robust scheme. Nonetheless, suppose that the code rate of the FEC mechanism can correct the errors introduced by obstruction. In that case, there are not many losses in the channel, and thereby, the number of retransmissions should remain controlled. Yet, an increase in the TCP retransmissions could be observed due to pending retransmissions from previous unsuccessful instances. Thus, one cannot conclude with certainty whether an increase or decrease in the number of retransmissions will be observed when the MCS is dynamically adjusted.

V. CONCLUSION

In this work, a backhaul using IEEE 802.11ad was successfully deployed in an outdoor scenario, which allowed collecting physical, network, and transport layer metrics of mmWave links suffering from long-term and short-term blockage. To the best of our knowledge, this is the largest collection campaign in WiGig-based networks, which includes an extensive cross-layer analysis of the behavior of long-term and short-term blocked 60 GHz links. The analysis conducted showed that maintaining higher modulations in long-term blocked channels might only be possible at the expense of high PERs (between 90%-100%), a maximum delay that can be in the order of one or a few seconds, high losses at the transport layer (between 37%-100%), and a high number of TCP retransmissions which reduces the throughput rates to a few Mbps. Even though dynamically adjusting the MCS leads to more controlled PERs and losses at the transport layer, it was still shown to reach maximum PERs of 40% even if the SNR remains constant in a long-term blocked link. Moreover, long-term blockage was shown to cause maximum delays of 0.9 seconds in some links and data rates that vary by a few hundred Mbps.

Furthermore, it has been shown that short-term blocked links are more likely to reach maximum PERs when there are sudden decreases in signal quality and are, therefore, more likely to adapt their MCSs to a more robust scheme. Still, the MCS returns to its pre-loss value as soon as the LOS condition is re-established. Hence, short-term obstruction not only decreases the maximum throughput at layer 2 but also causes packet losses up to 8% that will lead to a significant decrease of the TCP congestion window, further contributing to the decrease in the data rate. Moreover, the temporary LOS blockage led to maximum delays of half a second in the worst-case scenario. These aspects combined create a highly unreliable backhaul network, which comprises the constant data rate and maximum delay requirements of both TCP-based and UDP-based services.

Our future work will leverage the insights provided by this work to propose mechanisms at multiple layers that improve network reliability and resilience of outdoor IEEE 802.11ad-based networks. In particular, we aim to explore the potential of network coding techniques at the transport layer and of deep-learning-based forecasting algorithms to mitigate the negative impact of LOS obstruction.

REFERENCES

[1] K. Sakaguchi, E. M. Mohamed, H. Kusano, M. Mizukami, S. Miyamoto, R. E. Rezagah, K. Takinami, K. Takahashi, N. Shirakata, H. Peng, T. Yamamoto, and S. Nanba, "Millimeter-wave wireless LAN and its extension toward 5G heterogeneous networks," *IEICE Trans. Commun.*, vol. E98.B, no. 10, pp. 1932–1948, Jun. 2015.

[2] Z. Pi, J. Choi, and R. W. Heath Jr., "Millimeter-wave gigabit broadband evolution toward 5G: Fixed access and backhaul," *IEEE Commun. Mag.*, vol. 54, no. 4, pp. 138–144, Apr. 2016.

[3] P. Legg and R. McConnell, "Meshed backhauling of small cells using IEEE802.11ad at 60 GHz," in *Proc. Eur. Conf. Netw. Commun. (EuCNC)*, Jun. 2018, pp. 393–397.

[4] O. El Ayach, S. Rajagopal, S. Abu-Surra, Z. Pi, and R. W. Heath Jr., "Spatially sparse precoding in millimeter wave MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 13, no. 3, pp. 1499–1513, Jan. 2014.

[5] T. Rappaport, Y. Xing, G. R. MacCartney, A. F. Molisch, E. Mellios, and J. Zhang, "Overview of millimeter wave communications for fifth-generation (5G) wireless networks-with a focus on propagation models," *IEEE Trans. Antennas Propag.*, vol. 65, no. 12, pp. 6213–6230, Dec. 2017.

[6] *IEEE Standard for Information Technology—Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks—Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 2: Enhanced Throughput for Operation in License-exempt Bands above 45 GHz*, IEEE Standard 802.11ay-2021 (Amendment to IEEE Standard 802.11-2020 as amendment by IEEE Standard 802.11ax-2021), 2021, pp. 1–768.

[7] Y. Ghasempour, C. R. C. M. da Silva, C. Cordeiro, and E. W. Knightly, "IEEE 802.11ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi," *IEEE Commun. Mag.*, vol. 55, no. 12, pp. 186–192, Dec. 2017.

[8] *Next Generation Applications With Peraso's 60 GHz Wireless Technology*. Accessed: Jan. 13, 2023. [Online]. Available: https://perasoinc.com/60ghz-products/

[9] S. Singh, F. Ziliotto, U. Madhow, E. M. Belding, and M. Rodwell, "Blockage and directivity in 60 GHz wireless personal area networks: From cross-layer model to multihop MAC design," *IEEE J. Sel. Areas Commun.*, vol. 27, no. 8, pp. 1400–1413, Oct. 2009.

[10] A. Maltsev, R. Maslennikov, A. Sevastyanov, A. Khoryaev, and A. Lomayev, "Experimental investigations of 60 GHz WLAN systems in office environment," *IEEE J. Sel. Areas Commun.*, vol. 27, no. 8, pp. 1488–1499, Oct. 2009.

[11] (Jul. 2014). *QCA6335: 802.11ad MAC Baseband Transceiver*. [Online]. Available: https://www.qualcomm.com/products/technology/wi-fi/qca6335

[12] S. K. Saha, V. V. Vira, A. Garg, and D. Koutsonikolas, "60 GHz multi-gigabit indoor WLANs: Dream or reality?" 2015, *arXiv:1509.04274*.

[13] M. Kacou, V. Guillet, G. El Zein, and G. Zaharia, "Coverage and throughput analysis at 60 GHz for indoor WLAN with indirect paths," in *Proc. IEEE 29th Annu. Int. Symp. Pers., Indoor Mobile Radio Commun. (PIMRC)*, Sep. 2018, pp. 1–5.

[14] A. Loch, G. Bielsa, and J. Widmer, "Practical lower layer 60 GHz measurements using commercial off-the-shelf hardware," in *Proc. 10th ACM Int. Workshop Wireless Netw. Testbeds, Exp. Eval., Characterization*, Oct. 2016, pp. 9–16.

[15] K. C. Joshi, R. Hersyandika, and R. V. Prasad, "Association, blockage, and handoffs in IEEE 802.11ad-based 60-GHz picocells—A closer look," *IEEE Syst. J.*, vol. 14, no. 2, pp. 2144–2153, Jun. 2020.

[16] G. Bielsa, "Analysis and performance improvement of consumer-grade millimeter wave wireless networks," Ph.D. dissertation, Universidad Carlos III de Madrid, Getafe, Spain, Jul. 2019. [Online]. Available: http://hdl.handle.net/20.500.12761/745

[17] E. M. Mohamed, K. Sakaguchi, and S. Sampei, "Experimental work on WiGig coverage area management and beamforming training using Wi-Fi fingerprint," in *Proc. 14th IEEE Annu. Consum. Commun. Netw. Conf. (CCNC)*, Jan. 2017, pp. 772–777.

[18] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, "Demystifying 60 GHz outdoor picocells," in *Proc. 20th Annu. Int. Conf. Mobile Comput. Netw.*, Sep. 2014, pp. 5–16, doi: 10.1145/2639108.2639121.

[19] G. K. Tran, M. Nakamura, H. Nishiuchi, K. Sakaguchi, R. Santos, and K. Koslowski, "Outdoor experiment of mmWave meshed backhaul for realtime edge content delivery," in *Proc. IEEE Wireless Commun. Netw. Conf. Workshop (WCNCW)*, Apr. 2019, pp. 1–6.

[20] (Feb. 2020). *MetNet 60G Unlicensed mmWave Mesh Datasheet*. [Online]. Available: https://www.ccsl.com/v1/uploads/files/Metnet-60G-Mesh-datasheet.pdf

[21] A. Figueiredo, T. Ferreira, D. Raposo, M. Luís, P. Rito, and S. Sargento, "On the real experimentation and simulation models for millimeter-wave," *IEEE Access*, vol. 10, pp. 51191–51208, 2022.

[22] A. Richardson and J. Brady. *Using Self-Organizing mmWave to Deliver 5G Services*. [Online]. Available: https://www.bcba.ca/application/files/5515/5812/9724/CCS_Metnet__Broadnet.pdf

[23] M. Dahhani, G. Jakllari, and A.-L. Beylot, "Association and reliability in 802.11ad networks: An experimental study," in *Proc. IEEE 44th Conf. Local Comput. Netw. (LCN)*, Oct. 2019, pp. 398–405.

TÂNIA FERREIRA received the M.Sc. degree in electronics and telecommunication engineering from the Universidade de Aveiro, in 2021. She joined the Instituto de Telecomunicações (IT). In 2020, she received a research grant at IT to work toward a multi-layer approach to enhance reliability in mmWave backhauls, using machine learning tools and techniques, such as network coding. In 2022, she became a full-time Researcher at IT. Her research interests include multi-gigabit wireless systems (mmWave), mesh networks, and network reliability.

ALEXANDRE FIGUEIREDO received the M.Sc. degree in electronics and telecommunications engineering from the University of Aveiro, Portugal, in 2021. Since 2020, he has been a Researcher with the Instituto de Telecomunicações (IT), Aveiro, Portugal, exploring wireless communications, regarding hardware and medium access control reservation in network simulation and real experimentation. His current research interests include multi-gigabit wireless systems (mmWave), and low-power WAN (LPWAN) technologies, like LoRA and IEEE802.11ad.

DUARTE RAPOSO received the M.Sc. and Ph.D. degrees in information science and technology from the University of Coimbra, in 2012 and 2020, respectively. In 2012, he joined the Centre for Informatics and Systems of the University of Coimbra (CISUC), Portugal, as a Researcher, working in industrial wireless sensor networks in the topics of network management and security, exploring state-of-the-art anomaly detection techniques. In 2012, he also joined Eneida.IO developing industrial wireless solutions in the fields of oil and gas, mining, and energy. Since 2020, he has been with the Instituto de Telecomunicações (IT), Aveiro, Portugal, as a Research Assistant in network architectures and protocols. Until now, he is the author/coauthor of more than 34 scientific works in international conferences and journals and two national patents. His current research interests include wireless communications, deterministic networks, cellular communications, software-defined networks, and edge computing and orchestration.

MIGUEL LUÍS received the M.Sc. and Ph.D. degrees in electrical and computer engineering from the Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, Portugal, in 2009 and 2015, respectively. He is currently an Adjunct Professor with the Instituto Superior de Engenharia de Lisboa (ISEL) and a Researcher with the Instituto de Telecomunicações. He has been involved in several national and European research projects targeting new communications for mobile networks. He is the coordinator of Multihomed Software Defined Vehicular Networks (MH-SDVanet), a national funded research project. He contributes to several other research projects, such as Aveiro STEAM City (EU-UIA Program), SNOB-5G (FCT-MIT Program), and IMMINENCE (Celtic-P2020). He has published more than 80 scientific works, including three book chapters and 39 publications in peer-reviewed international journals. His research interests include medium access control for wireless systems, routing and dissemination mechanisms for mobile networks and management, and orchestration and softwarization of future networks.

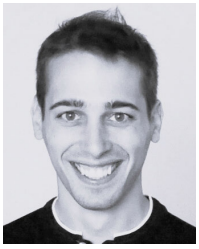

PEDRO RITO received the M.Sc. degree in electronics and telecommunications engineering from the University of Aveiro, Portugal, in 2011, and the Ph.D. degree in electrical engineering from Technische Universität Berlin, Germany, in 2019. In 2012, he joined IHP GmbH, Germany, as a Researcher. In 2018, he joined Cisco Optical GmbH, Nürnberg, Germany, as a Hardware Engineer, investigating and developing state-of-the-art high-efficient and high-bandwidth electro-optical interconnects for datacenters, metro and long-haul communications. Since 2020, he has been with the Instituto de Telecomunicações (IT), Aveiro, Portugal, as an Assistant Researcher at the Network Architectures and Protocols Group. During his research activity, he has published more than 40 publications and holds one U.S./EU patent. His current research interests include radio access networks, software-defined networking, vehicular communications, edge-computing, and smart cities.

SUSANA SARGENTO was a Visiting Ph.D. Student at Rice University, from 2000 to 2001, and a Guest Faculty at Carnegie Mellon University, in 2008. She is currently a Full Professor with the University of Aveiro and a Senior Researcher with the Institute of Telecommunications, where she is leading the Network Architectures and Protocols Group. She has been leading research projects with telecom operators and OEMs. She has also been involved in several FP7 projects (4WARD, Euro-NF, C-Cast, WIP, Daidalos, and C-Mobile), EU Coordinated Support Action 2012-316296 "FUTURE-CITIES," EU Horizon 2020 5GinFire, EU Steam City, CMU-Portugal projects (S2MovingCity, DRIVE-IN with Carnegie Melon University), and MIT-Portugal Snob5G Project. She has organized several international conferences and workshops, such as ACM MobiCom and IEEE Globecom. She has co-founded a vehicular networking company, Veniam (www.veniam.com). She is also the co-coordinator of the national initiative of digital competences in the research axis INCoDe.2030, belongs to the evaluation committee of the Fundo200M (www.200m.pt), and one of the scientific directors of CMU-Portugal Programme. She regularly acts as an expert for European research programmes. She has published more than 400 scientific papers. Her main research interests include self-organized networks, in ad-hoc and vehicular network mechanisms and protocols, such as routing, mobility, security and delay-tolerant mechanisms, resource management, and content distribution networks. She is the Winner of the 2016 EU Prize for Women Innovators and the Femina 2020 Prize in science. She has also been a Reviewer of conferences and journals, such as *IEEE Networks* and *IEEE Communications Magazine*.