Redundancy and Diversity in Wireless Networks to Support Mobile Industrial Applications in Industry 4.0

M.C. Lucas-Estañ, *Member, IEEE*, B. Coll-Perales, *Member, IEEE*, and J. Gozalvez, *Senior Member*, *IEEE*

***Abstract*—Factories are evolving into fully digitalized and networked structures for more adaptive and agile production ecosystems in the context of the Industry 4.0. Wireless communications will be a technical pillar of this evolution as it improves the reconfigurability of factories and the integration of mobile robots and objects. The integration of industrial wireless networks into the Industry 4.0 requires solutions capable to support highly reliable and deterministic low latency communications. This is particularly challenging for mobile industrial applications with constantly changing link quality conditions. This study experimentally evaluates for the first time the capacity of diversity and redundancy to improve the reliability and latency of wireless networks for mobile industrial applications. To this aim, a prototype is built in a collaborative robotics experimental facility. The prototype wirelessly connects a dual-arm robot and a mobile robot that collects and supplies components to the dual-arm robot. The prototype implements redundancy and diversity (using multipath TCP) for the wireless connections between both robots. The conducted trials show that both techniques improve the reliability of mobile industrial wireless communications even under the presence of interference. However, redundancy achieves lower latency levels and represents then the most attractive solution to support mobile industrial applications.**

***Index Terms*— Industry 4.0, redundancy, diversity, industrial wireless networks, reliability, latency, mobile industrial.**

I. Introduction

INDUSTRY 4.0 (or Factories of the Future, FoF) targets the digitalization of manufacturing for building more advanced, adaptive and zero-defect production systems [1]. This vision requires industrial communication networks that can sustain data-intensive services and transmit data with deterministic latency and high reliability. Latency and reliability communication requirements have been identified for Industry 4.0 use cases by international organizations such as the 5G Alliance for Connected Industries and Automation (5G-ACIA) [2] and European Factories of the Future Association (EFFRA) [3]. These use cases may utilize fixed and/or wireless connections. Wireless connectivity is expected to play an increasingly relevant role in Industry 4.0 given the need for more modular and reconfigurable factories. Wireless communications are also necessary to support mobile industrial applications and connect mobile objects (e.g. robots or vehicles), moving parts of the factory, and even workers. Mobile robots will be highly relevant in the factories of the future as they will transport goods or materials within the factory and support processes [1]. Connecting mobile robots requires reliable and deterministic low latency wireless connections [4]. For example, 3GPP establishes that standard mobile robot applications require latency levels between 40 and 500 ms [5]. Guaranteeing reliable and deterministic low latency connections in harsh industrial environments can be a challenge, in particular for mobile industrial use cases [6].

Most wireless devices include two or more radio interfaces (of the same or different technologies). This offers the possibility to establish multiple wireless links that can be exploited to improve the reliability and quality of wireless connections. Diversity and redundancy are usually proposed to achieve these objectives [7]. Diversity exploits the availability of multiple wireless links between transmitter and receiver, and dynamically selects the most suitable link to transmit each packet. The selection can be driven by the availability, reliability, throughput or latency of each link. Redundancy simultaneously transmits copies of each packet over multiple wireless links or network paths[1]. Diversity improves the spectrum efficiency compared to redundancy but requires algorithms to select the most suitable link. Its performance could also be compromised if none of the links provide a sufficiently good quality by themselves. Recent studies have demonstrated that diversity and redundancy can improve the reliability of industrial wireless communications ([8], [9]). These studies have focused to date on wireless communications between fixed nodes in a factory. This paper extends the current state of the art by experimentally demonstrating for the first time the impact of diversity and redundancy on industrial wireless communications for mobile applications. To this aim, we have implemented a prototype in a collaborative robotics experimental facility. The prototype wirelessly connects a dual-arm robot and a mobile robot. The mobile robot collects and supplies components to the dual-arm robot. To this aim, both robots must be able to continuously communicate and exchange data (e.g. the location of the component to be supplied and the location and trajectory of the mobile robot). The prototype implements redundant and diverse wireless connections between the dual-arm and mobile robots. For diversity, this study considers the use of multipath

Manuscript received June 20, 2019; revised Nov. 4, 2019; accepted Feb. 23, 2020. This work was supported in part by the European Commission through the FoF-RIA Project AUTOWARE (No. 723909), and by the Ministerio de Ciencia e Innovación, AEI and FEDER funds through the project TEC2017-88612-R.

M.C. Lucas-Estañ, J. Gozalvez, and B. Coll-Perales are with the UWICORE Laboratory, Universidad Miguel Hernandez de Elche (UMH), 03202 Elche, Spain (e-mail: {m.lucas, bcoll, j.gozalvez}@umh.es).

[1] Redundancy is a particular implementation of diversity. For clarity, diversity is utilized in this paper to refer to configurations that exploit multiple links but do not transmit duplicated packets simultaneously over two links.

TCP. The conducted trials demonstrate that redundancy and diversity improve the reliability of mobile industrial wireless communications even under the presence of interference. However, redundancy achieves lower latency levels and represents then the most attractive solution to support mobile industrial applications.

II. Related Work

Recent studies have shown that redundant wireless links can improve the reliability and latency of industrial wireless communications [10]. The Parallel Redundancy Protocol (PRP) [11] for real-time Ethernet (RTE) proposes transmitting two copies of the same frame over two different networks or network paths. This approach reduces the transmission latency and the likelihood that a packet is not delivered to the destination. This was confirmed in [10] where authors analyzed the performance of PRP over WiFi by means of simulations. Experimental studies were also presented in [8] and [9]. In [9], authors evaluate the reliability and latency achieved when streaming mission-critical phasor data is duplicated and transmitted over two independent IEEE 802.11b links. The transmission takes place between two fixed transmitters and one receiver that integrates two wireless interfaces. The transmitters and receiver are 200 m away, and are located on the rooftop of three buildings. There is line of sight between the sending and receiving devices. [9] experimentally demonstrates that redundancy can improve the reliability and latency of wireless connections between fixed devices. Redundant WiFi-based industrial wireless links are also analyzed in [8] where authors experimentally evaluate the effect of different IEEE 802.11 management schemes on the effectiveness of redundancy. Authors show that it is important to adequately configure the energy-saving IEEE 802.11 Delivery Traffic Indication Map mechanism to guarantee the effectiveness of redundant industrial wireless links.

Redundancy implies the simultaneous transmission of duplicated packets over two wireless interfaces. An alternative that improves the spectrum efficiency is diversity. Diversity does not transmit the same packet over two wireless links. Instead, it exploits the availability of multiple wireless links to select the link that guarantees the highest reliability, network coverage or availability, or provides the best Quality of Service (QoS) at each point in time. The dynamic selection and configuration of wireless links improves the reliability of wireless communications while efficiently utilizing the spectrum [12]. Diversity is commonly utilized in wireless networks and links can implement the same or different communication technologies [7]. An example is Multipath TCP (MPTCP), a protocol that allows the transmission and reception of data using multiple network interfaces [13]. MPTCP selects for each transmission the interface or network path that provides the lowest RTT (Round-Trip Time) between transmitter and receiver. MPTCP exploits the availability of multiple links between transmitter and receiver while efficiently utilizing the network resources and ultimately increasing the network bandwidth [14]. MPTCP and other diversity protocols require checking the link quality to detect any possible degradation and modify the selected network path. Detecting the communication conditions and reacting to possible changes might not be immediate [15], and this could affect the capacity to support deterministic low latency industrial applications [12]. An interesting study is presented in [16] where authors integrate redundancy in the MPTCP protocol to reduce the latency and jitter while achieving high reliability levels. The study experimentally measured the application level RTT between a client in Germany and an Amazon server in USA. The client is implemented in a notebook with IEEE 802.11 and LTE interfaces. [16] shows that the proposed solution can effectively reduce the average RTT when compared with only using one interface (IEEE 802.11 or LTE). However, the comparison with the traditional MPTCP was only performed through simulations in [16].

To the authors' knowledge, no study has previously analyzed experimentally how diversity and redundancy in wireless networks can improve the support of mobile industrial applications. Mobility implies the presence of Line of Sight (LOS) and Non Line of Sight (NLOS) conditions that result in significant variations of the channel quality. Exploiting the availability of multiple wireless links between transmitter and receiver could combat such variations to ensure high reliability and deterministic low latency levels. This paper progresses the current state of the art by experimentally evaluating for the first time the impact of diversity and redundancy on the reliability and latency of mobile industrial wireless communications. To this aim, a prototype has been built in a collaborative robotics experimental facility that requires the continuous connectivity between a fixed dual-arm robot and a mobile robot that collects and supplies components to the dual-arm robot.

III. Collaborative Robotics Experimental Facility

A wireless prototype has been implemented and integrated into an experimental facility for collaborative robotics at TEKNIKER, an industry-oriented research center in the north of Spain. The experimental facility (Fig. 1) is a standalone workcell for collaborative assembly in an industrial shopfloor. It includes a dual-arm robot, a mobile robot, a tool changer, interaction devices, and multiple sensors for safety and interaction. The mobile robot is autonomous and must supply components to the dual-arm robot. The dual-arm robot wirelessly sends a request for a component and its location to the mobile robot. The mobile robot computes the trajectory to the location of the component, and uses a Velodyne LiDAR sensor to autonomously reach the location, collect the requested component and bring it back to the dual-arm robot. The dual-arm and mobile robots continuously communicate. The dual-arm robot periodically asks the mobile robot for its position and the mobile robot reports it back. Thanks to this interaction, the dual-arm robot can control the complete process and detect any potential problems. This could include, for example, incorrect trajectories. In this case, the dual-arm robot could wirelessly provide corrections to the mobile robot. The complete process stops if the communication between the robots fails. The communication fails if messages are not received at the receiver side before a pre-established deadline. The deadline considered for the delivery of a message is equal to 500 ms. Reliable wireless connections are therefore necessary to ensure the correct and safe operation of the collaborative assembly process. The communication between the dual-arm and mobile robots corresponds to the third use case related to mobile robots identified in [5]. This use case includes periodic communication for standard mobile robot operation and traffic management[2]. The 3GPP establishes in [5] that standard mobile robot applications (use case 3) require latency levels between 40 and 500 ms which is in line with the maximum latency (500 ms) considered in this work.

IV. Prototype

The implemented wireless prototype supports communication between the dual-arm robot and the mobile robot at the collaborative robotics experimental facility. The prototype is illustrated in Fig. 2. It is composed of two communication nodes that are integrated each at one robot. Both communication nodes have been implemented on laptops and barebone PCs operating under Linux using the Ubuntu distribution. The Communication Node 1 or CN1 is integrated in the dual-arm robot and is connected to the robot's controller using a wired connection (Ethernet). The communication node 2 or CN2 is integrated in the mobile robot and is connected using Ethernet to the mobile robot's server.

The prototype exploits redundancy and diversity to increase the reliability of the wireless connection between the robots.

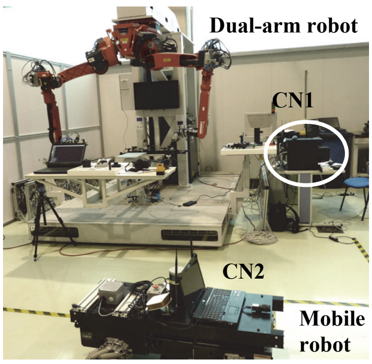

Fig. 1. Workcell for collaborative assembly.

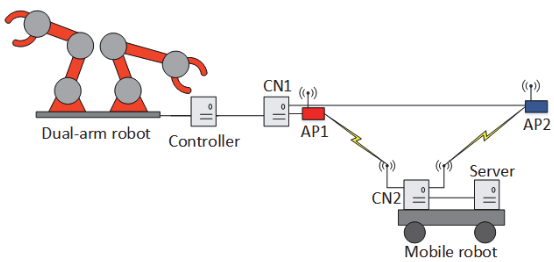

Fig. 2. Implemented prototype.

[2] Four use cases are identified in [5] related to mobile robots. Use case 1 considers periodic communication for the support of precise cooperative robotic motion control, machine control, cooperative driving. Use case 2 includes periodic communication for video-operated remote control. Use case 4 considers real-time streaming data transmission (video data) from a mobile robot to the guidance control system. Our implementation corresponds then to the third use case for standard mobile robot operation and traffic management.

To this aim, the prototype establishes two independent wireless links between the dual-arm and mobile robots. The links currently utilize WiFi (IEEE 802.11g)[3] although the prototype can integrate other wireless technologies. The redundancy and diversity solutions are implemented at the application and TCP layers respectively. They are then independent of the lower layers of the protocol stack and are hence compatible with other PHY/MAC layer solutions to improve the reliability of the wireless links (e.g. MIMO, HARQ retransmissions). CN1 is connected using Ethernet to two WiFi access points deployed in the scenario (AP1 and AP2). CN2 at the mobile robot is equipped with two wireless interfaces so that two independent wireless connections can be established with AP1 and AP2. CN2 uses its built-in wireless interface to connect to AP2. CN2 has been equipped with an additional external Wireless ExpressCard interface to communicate with AP1.

CN1 and CN2 integrate packet sniffers developed by the authors. These sniffers are used to monitor the transmitted and received packets[4]. In particular, CN1 integrates a packet sniffer to monitor the IEEE 802.3 Ethernet frames exchanged with the dual-arm robot. This sniffer can extract from the TCP header information such as the source or destination ports, the TCP sequence number, the acknowledgement (ACK) sequence number and the timestamp. CN2 integrates a packet sniffer to capture the wireless IEEE 802.11 packets. The sniffer extracts information from the radiotap header such as the IEEE 802.11 channel frequency, the packet and headers' size, the type of packet (management, data, control), the data rate (which is related to the Modulation and Coding Scheme or MCS used to transmit the packet [17]), the RSSI (Received Signal Strength Indicator) and the timestamp. We also locate two sniffer nodes next to AP1 and AP2 to capture and analyze the IEEE 802.11 packets transmitted and received by the APs. These external sniffer nodes were necessary since the operating system of AP1 and AP2 (model TP-Link TL-WA901ND) is not open.

V. Diversity and Redundancy-based Industrial Wireless Communications

A. *Multipath TCP for diversity*

Diversity has been implemented using the MPTCP protocol [13]. This protocol allows establishing multiple TCP connections between a transmitter and a receiver using multiple disjoint paths. Each TCP connection is managed as a regular TCP connection. Our prototype integrates the 'Linux Kernel MultiPath TCP project' (version 2.5.2) implementation [18]. This implementation includes two scheduling schemes to

[3] Typical ESS (Extended Service Set) IEEE 802.11 networks would see their latency performance degrade when the mobile robot switches the serving AP due to the delay introduced in this process. Our solutions can maintain two active links and better combat any degradation.

[4] The implemented sniffers use the open source library *libpcap.h* to extract information from the headers of the packets. *libpcap.h* is commonly utilized in analyzer tools such as Wireshark and Kismet. Using our own sniffers gave us flexibility to access in runtime relevant parameters and metrics to monitor the quality of the communications during the trials. It also allowed us to capture data packets and customize their logging for a most efficient post-processing.

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TII.2020.2979759

decide what TCP connection is utilized to transmit data packets. The default scheduler measures the RTT over each of the TCP connections, and sends data packets using the TCP connection with the lowest RTT. If the TCP congestion window for this TCP connection is full, the scheduler sends the data packets using the TCP connection with the second lowest RTT. The second scheduler follows a round-robin approach, i.e. data packets are equally distributed to each available TCP connection. Our trials have been conducted using the default scheduler since it can better exploit the benefits of path diversity.

Packets generated by the dual-arm robot are first transmitted from the robot's controller to CN1 using a TPC socket. CN1 establishes a socket with CN2 to exchange packets. Only one socket is created for the two wireless interfaces. MPTCP establishes two different TCP connections (or sub-flows) through AP1 and AP2. This is transparent to the upper layers of the protocol stack that operate considering that a single socket is established between CN1 and CN2. The scheduler selects then the TCP connection that experiences the lowest RTT to transmit the packet. CN2 forwards to the mobile robot's server any packet received from CN1. To this aim, CN2 uses a TCP socket created between CN2 and the server. A similar process is followed at CN2 to transmit packets from the mobile robot's server to the dual-arm robot.

B. *Redundant industrial wireless communications*

MPTCP chooses the wireless link with the lowest RTT at each point in time. In principle, this link provides the best link quality to exchange packets. However, the link quality might change and MPTCP needs to detect such changes so that the scheduler can react and select the other available TCP connection when appropriate. This can affect the latency. An alternative is establishing redundant wireless connections between CN1 and CN2. In this case, packets are duplicated and transmitted through both wireless links (i.e. using AP1 and AP2) simultaneously. Redundancy is implemented at the application layer, i.e. CN1 and CN2 manage and process duplicated and exchanged packets at the application level. Similarly to the MPTCP solution, a TCP socket is established between the dual-arm robot's controller and CN1 to exchange packets, and between CN2 and the mobile robot's server. In the redundancy-based solution, CN1 establishes two TCP sockets with CN2: one through AP1 and the other through AP2. When CN1 receives a packet from the dual-arm robot, it adds a header with a unique sequence number. Then, CN1 duplicates the packet and forwards the two copies to CN2; each copy is forwarded using one of the two APs. CN2 forwards the first received copy of the packet to the mobile robot's server, and discards the second one. The second copy is identified as it has the same sequence number as the first copy. A similar process is followed at CN2 to transmit packets from the mobile robot to the dual-arm robot.

VI. Scenarios For The Trials

The trials have been conducted at TEKNIKER. The trial area includes two open and spacious rooms that are separated by a 30 cm-thick and heavy concrete wall (Fig. 3). The rooms are communicated by an industrial vertical 5 m wide lift door. Workers and forklifts freely move around the two rooms that include work cells and high-volume machinery tools such as wind turbines, forming press, robots, and refrigerated cold chambers. There is also a cabin for wind generation tests in the trial area (the grey box represented in Fig. 3.b).

A dedicated wireless network with two APs (Fig. 3.b) has been deployed in the trial area. The location of the two APs was chosen to guarantee wireless coverage in the two rooms. The APs are configured to transmit at 2.4 GHz using the IEEE 802.11g standard (i.e. 20 MHz channel). The transmission power of both APs is limited to 20 dBm following European regulations for ISM transmissions in this frequency band. AP1 and AP2 create two different (and private) wireless networks that operate in the non-overlapping channels 1 (AP2) and 11 (AP1). These two networks coexist (without any priority in channel access) with the 2.4 GHz wireless networks available at the TEKNIKER premises. Fig. 4.a shows that there are at least 12 permanent WiFi networks operating in the trial area.

Fig. 4.b shows the wireless networks in the trial area when an interfering node is activated. One of the objectives of this study is to test the performance of redundant and diverse industrial wireless communications under adverse communication conditions. To this aim, we included in the trial area an interference node (IN, Fig. 3.b). The IN is at a fixed location and is switched on or off depending on the specific test. It has been implemented using an USRP2 (Universal Software Radio Peripheral version 2) board, and the GNU Radio Companion software development toolkit

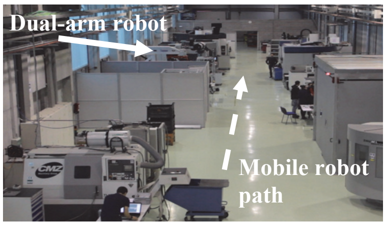

a) One of the two rooms.

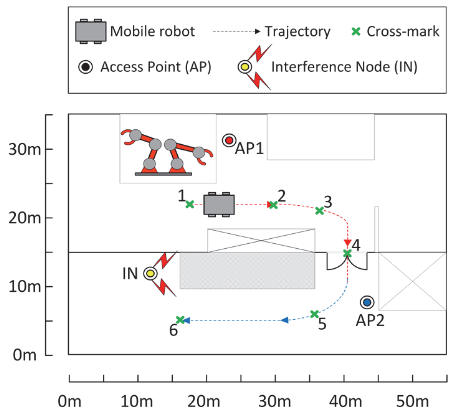

b) Simplified view of the shopfloor plan showing the path followed by the mobile robot and the position of the dual-arm robot and the APs. Empty boxes represent work cells. Crosslined-boxes represent high-volume machinery. The grey box is a cabin for wind generation tests.

Fig. 3. Industrial shopfloor at TEKNIKER.

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TII.2020.2979759

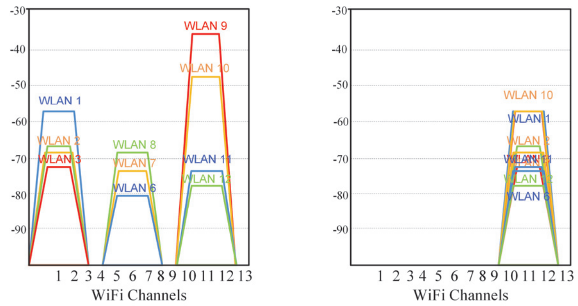

a) Interferer switched off. b) Interferer switched on.

Fig. 4. WiFi networks operating at 2.4 GHz in the facility. The figure has been obtained using Android's WiFiAnalyzer application.

running on a PC. The IN generates a jamming signal transmitted at 18dBm and centered at 2.412 GHz (i.e. IEEE 802.11g channel 1 where AP2 operates). Fig. 5 shows the generated interference signal. This signal affects all WiFi networks operating in channel 1. However, while AP2 is fixed at this channel, the WiFi networks permanently deployed at the premises dynamically select the channel based on the detected interference. As a result, when IN is switched on, all permanent networks change channels and select channel 11.

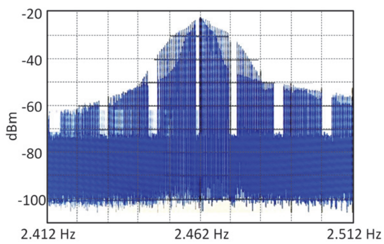

Fig. 5. Generated interference signal[5].

During the trials, the dual-arm robot requests a component to the mobile robot and indicates the location of this component. The path the mobile robot has to follow to collect the requested component is represented in Fig. 3.b. At the start of a trial, the mobile robot is located in cross-mark #1 in the floor map and the requested component is in cross-mark #6 in the other room. The mobile robot passes through cross-marks #2 to #5 until it reaches cross-mark #6. The separation distance between cross-marks #1 and #4 is approximately 23 m, and between cross-mark #4 and #6 is 25 m. When the mobile robot reaches cross-mark #6, it turns around and comes back to cross-mark #1 to bring the component to the dual-arm robot. The trial finishes when the mobile robot reaches cross-mark #1. During the execution of the trials, the dual-arm robot periodically requests the mobile robot to send its location. The location requests are transmitted at 100 Hz (i.e. one request every 10 ms) using a 40-byte application packet. When the mobile robot receives a request, it replies with a 29-byte application packet that includes its location. Packets generated by the robots are transmitted through CN1 and CN2 that implement the redundant and diverse wireless solutions. CN1 and CN2 get synchronized at the beginning of each trial using NTP (Network Time Protocol). The different obstructing elements in the trial area rapidly attenuate the transmitted signal when the mobile robot moves along the path. This is particularly the case when the transmitter and receiver are in different rooms. In this case, the concrete wall and the cabin for wind generation tests strongly attenuate the wireless signal.

VII. Performance Evaluation

The experimental performance achieved with the MPTCP and redundant solutions presented in Section V are compared in this Section against the performance obtained when a single AP is used, i.e. the mobile robot connects either to AP1 or AP2. Several trials have been conducted for each solution and scenario configuration. Same trends have been observed in all trials carried out for the same solution and configuration. For better clarity, this Section shows the results for a selected trial per solution and scenario configuration.

A. *Reliability and latency*

Fig. 6 and Fig. 7 show the results of trials where the performance obtained with MPTCP and redundant communications respectively is compared against that obtained when the mobile robot connects to a single AP. Both figures correspond to scenarios without interference (IN switched off). All plots are depicted as a function of the time (in seconds) required by the mobile robot to go from cross-mark #1 to #6 and come back to the initial position (Fig. 3.b). Fig. 6 and Fig. 7 represent the performance at the mobile robot although the wireless links between both robots are bidirectional. Fig. 6.a and Fig. 7.a represent the RSSI of the packets received at CN2 (attached to the mobile robot) from AP1 and AP2. Both figures show on the top the cross-marks (represented with #number) that the mobile robot traverses at each point in time. It is important noting that the redundant configuration always maintains active connections with both APs so the RSSI values in Fig. 7.a are continuous (unless there is a link outage). On the other hand, MPTCP only maintains an active connection with the AP selected by the default scheduler. This explains why Fig. 6.a does not report continuous RSSI values with each AP. This actually applies to all MPTCP results in Fig. 6. Fig. 6.b and Fig. 7.b represent the average data rate of the packets received from each AP; this average is computed every 0.5 s. Fig. 6.c and Fig. 7.c represent the average PER (Packet Error Ratio) experienced at computed every 0.5 s. Fig. 6.d and Fig. 7.d represent the average PER[6] when implementing MPTCP and redundancy respectively. In this case, the PER is computed considering the packets received through either AP1 or AP2 for MPTCP, and through AP1 and AP2 for the redundant configuration. Fig. 6.e and Fig. 7.e depict the transmission latency between CN1 and CN2 for the packets received through AP1 under the MPTCP and redundant configurations, and Fig. 6.f and Fig. 7.f depict the transmission latency between CN1 and CN2 for the packets received through AP2 for the MPTCP and redundant configurations respectively. The latency is measured as the time elapsed between the time instant CN1 (dual-arm robot) sends a packet and the time instant the packet is received at

[5] This signal has been measured with a different output power to that used in the trials due to the limited capabilities of the utilized spectrum analyzer.

[6] The figures also show the average and variance of the end-to-end PER.

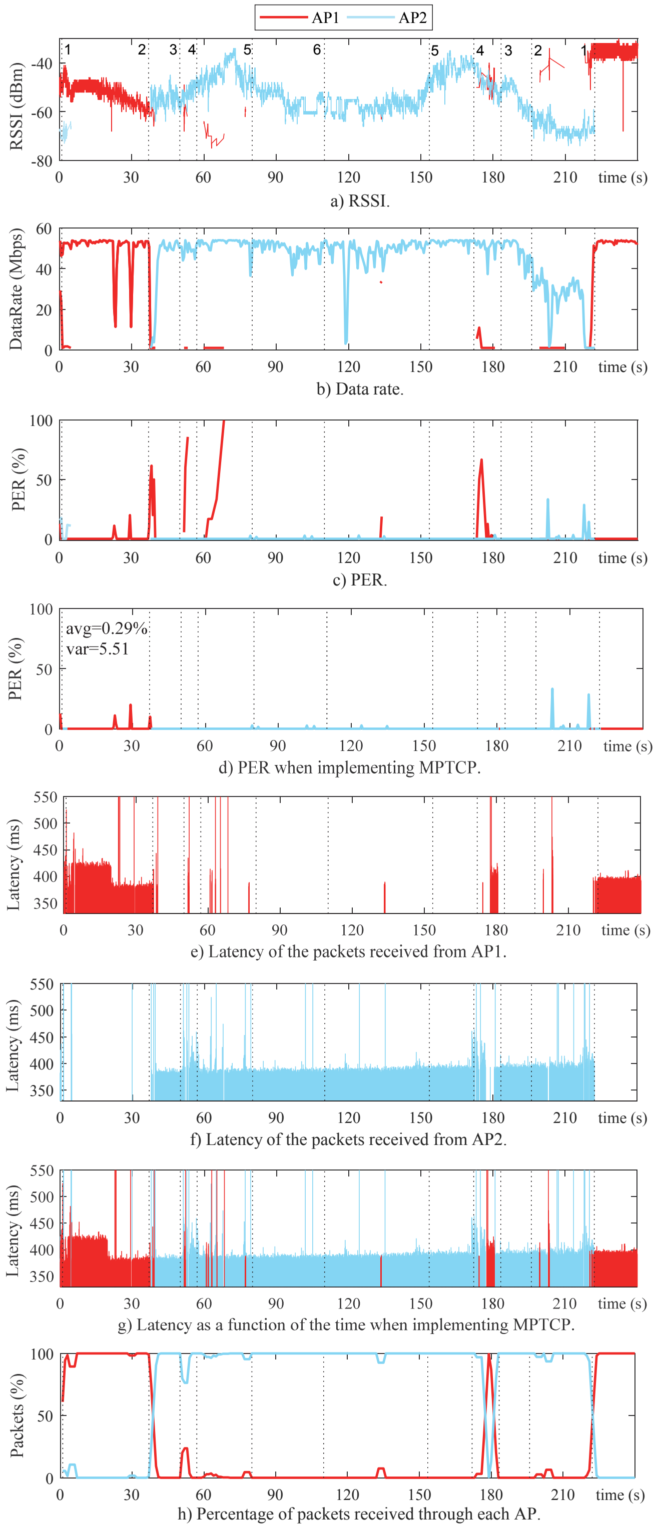

Fig. 6. Comparison of MPTCP with connections to a single AP.

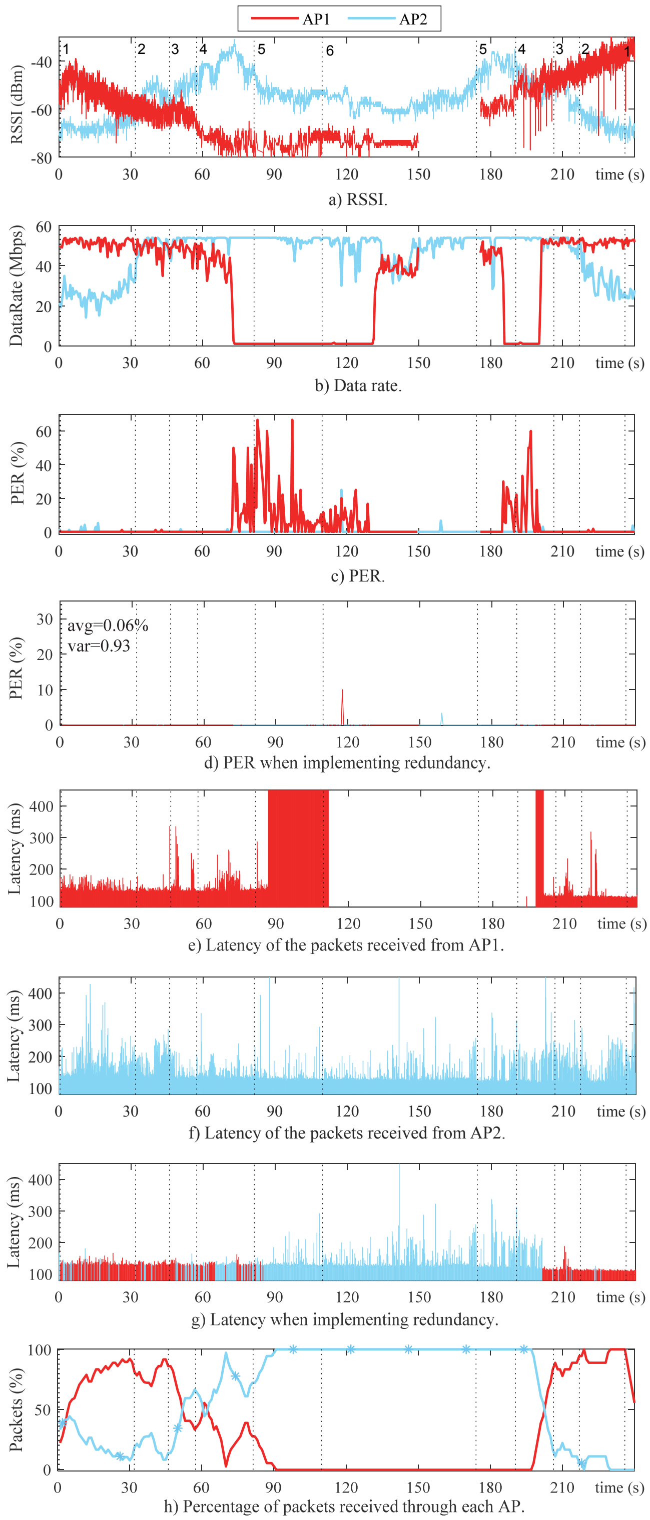

Fig. 7. Comparison of redundant wireless communications with connections to a single AP.

CN2 (mobile robot). The latency is measured at the application level, and it takes into account possible MAC and TCP retransmissions. Fig. 6.g and Fig. 7.g represent the latency between CN1 and CN2 when utilizing MPTCP and redundancy respectively. In Fig. 6.g, each packet is transmitted only through one AP (the default MPTCP scheduler selects for each transmission the AP with the lowest RTT). Different colors are used to identify the AP that transmitted each packet. In Fig. 7.g, each packet can be received through either the wireless links with AP1 or AP2, and the latency is computed considering the first copy of a packet received through either of the wireless links. The colors in Fig. 7.g identify the AP that transmitted the first copy of each received packet (i.e. the one with the lowest latency).

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TII.2020.2979759

Fig. 6.h and Fig. 7.h depict the percentage of packets that are received at the mobile robot through each AP as a function of time for the MPTCP and redundant configuration respectively.

We first analyze the performance achieved when only one AP is used for connecting the dual-arm and mobile robots. Fig. 7.a shows that when the mobile robot moves between cross-marks #4 and #6 (Fig. 3.b) the quality of the wireless link with AP1 degrades, and the experienced RSSI reaches values below -70 dBm[7]. This is due to challenging propagation conditions resulting from the presence of obstructing elements between transmitter and receiver. AP1 decreases the data rate between cross-marks #4 and #6 (Fig. 7.b). In particular, it uses a lower data rate (corresponding to a more robust MCS) to better combat poor link quality conditions. However, this cannot compensate the low RSSI and the increase in the PER (Fig. 7.c); the PER can exceed 50% between cross-marks #4 and #6. The mobile robot even loses temporally the connection (link outage) with AP1 between t=85s and t=190s, i.e. when the mobile robot is between cross-marks #5 and #6. Thirty six percent of the packets sent through AP1 during the link outage do not reach the destination[8]. Fig. 7.a also shows that when the mobile robot moves between cross-marks #2 and #6, AP2 provides better RSSI than AP1. However, when the mobile robot moves between cross-marks #1 and #2, the RSSI for packets received from AP2 decreases and lower data rates (Fig. 7.b) must be used to maintain low PER values (Fig. 7.c); peak PER values up to 5% are still observed.

Fig. 6 and Fig. 7 show that MPTCP and redundancy improve the reliability of industrial wireless communications since the end-to-end connection is not compromised when either of these two solutions is utilized. This is for example observed when comparing the end-to-end PER achieved with MPTCP (Fig. 6.d) and redundancy (Fig. 7.d) with the PER experienced through AP1 and AP2 (Fig. 6.c and Fig. 7.c). Fig. 6.d and Fig. 7.d show that the mobile robot receives most of the packets through AP1 (red color) when it is in the proximity of the dual-arm robot (between cross-marks #1 and #3 approximately). During the link outage with AP1, the robots can still communicate through AP2 (blue color packets in Fig. 6.d and Fig. 7.d between cross-markers #3 and #6). Diversity and redundancy improve the reliability of the communication between the robots. However, differences are observed between both solutions. MPTCP exploits diversity and sends packets through the AP that provides the lowest RTT. This adaptive process results in that MPTCP can guarantee overall a better wireless connection between both robots than when only one AP (AP1 or AP2) is used. This is actually observed in Fig. 8 that represents a pie plot of the data rate used to transmit the WiFi packets. The figure shows the results when packets are only transmitted through one AP (AP1 or AP2), and when redundancy and MPTCP are used. Transmitters (APs and CN2) adapt the data rate used to transmit packets based on the quality of the wireless links. The transmitters use higher data rates when the link quality conditions are good. The transmitters switch to lower data rates related to more robust MCSs when the link quality conditions degrade. Fig. 8 shows that MPTCP transmit 84% of the packets using MCSs with data rates equal or higher than 48 Mbps. This percentage decreases to 64% and 74% when transmissions only utilize AP1 and AP2 respectively. This result reflects the fact that MPTCP is capable to operate overall with better link quality conditions compared to scenarios in which transmissions rely on a single AP.

MPTCP must constantly estimate the link providing the lowest RTT to select the transmitting AP. This reactive approach can result in small performance degradations as the transmitting AP is only changed when degradations are observed. On the other hand, the redundant configuration duplicates packets and sends the copies simultaneously through both APs. The receiver selects on a per-packet basis the copy that is first correctly received, and the redundant configuration always uses the link with best quality. This guarantees the highest data rates of all evaluated configurations. In fact, the use of redundancy resulted in that 93% of the packets were transmitted using the highest data rates (48 and 54 Mbps) as shown in Fig. 8. This is a significant gain compared to the MPTCP configuration and the use of a single AP. The trends observed with the selected data rates are also observed for the experienced end-to-end PER. Redundancy guarantees the lowest end-to-end PER since the PER is the minimum at each point in time of the PER experienced in the links with AP1 and AP2. This explains why the redundant configuration (Fig. 7.d) achieves a lower end-to-end PER than MPTCP (Fig. 6.d). The redundant configuration achieved a mean PER value of 0.06% with a variance of 0.93 while MPTCP resulted in a mean PER equal to 0.29% and a variance equal to 5.51. These results also show that the redundant configuration minimizes the variability of the PER resulting from channel variations.

Fig. 6.g and Fig. 7.g show the end-to-end latency achieved with MPTCP and redundancy respectively. Fig. 6.e and Fig. 6.f, and Fig. 7.e and Fig. 7.f plot the latency achieved when transmitting packets over AP1 and AP2 when implementing MPTCP and redundancy respectively. The comparison with Fig. 6.g and Fig. 7.g shows that redundancy significantly reduces the latency of industrial wireless communications. In particular, the redundant configuration significantly reduces the percentage of packets that experience higher latency

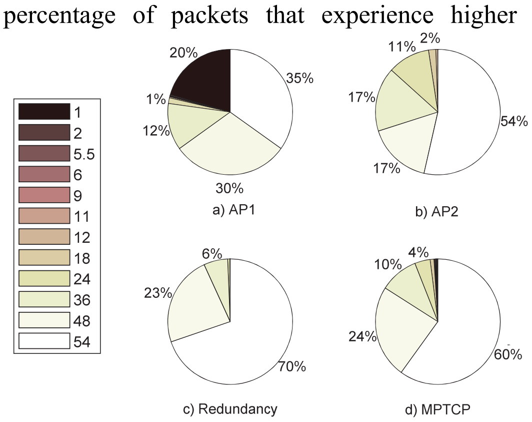

Fig. 8. Pie plot of the data rate (in Mbps) used for the transmission of the IEEE 802.11g packets.

[7] The sensitivity threshold of the CN2's wireless interfaces is -84 dBm.

[8] Packets are transmitted using TCP. Packets are stored in the transmitter's buffer until they are correctly received at the destination. If a wireless link is in outage, packets generated during the outage period are stored in the buffer to be transmitted when the link is re-established. Packets can be lost due to the overflow of the transmitter's buffer.

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TII.2020.2979759

values compared to when packets are only transmitted through AP1 or AP2. For example, 46% and 20% of the packets experienced a delay higher than 133 ms when transmitted through AP1 and AP2 respectively. This percentage is reduced to 5% with redundancy. Even during the link outage of AP1 (from t=85 s to t=190 s) the redundant configuration can maintain latency levels generally below 200 ms while it is significantly higher for MPTCP and of course for the connection through AP1. Fig. 7.g shows that redundancy can guarantee latency levels below 150 ms for most of the mobile robot's route. These results are in line with the 3GPP requirements for standard mobile robot applications that identify necessary latency levels between 40 and 500 ms (the application evaluated in this work corresponds to the third use case defined in [5] for mobile robots). With redundancy, CN2 receives at the application level the copy of each transmitted packet that is first correctly received. This results in that the redundant configuration can select on a per-packet basis the link providing the best reliability and the lowest end-to-end latency as illustrated in Fig. 7.g. This is actually visible in Fig. 6.h and Fig. 7.h that depict the percentage of packets that are received at the mobile robot through each AP. The comparison of Fig. 6.h and Fig. 7.h with Fig. 7.a and Fig. 7.b shows that the redundant configuration better adapts the reception to the quality of the links with the two APs. For example, AP2 provides higher RSSI values between t=55 s and t=195 s approximately that corresponds to the time the mobile robot is between cross-marks #4 and #6. Fig. 6.h and Fig. 7.h show that both configurations receive most of the packets through AP2 in this time period. However, the redundant configuration can also better and more quickly adapt to the fast variations of the channel quality than MPTCP. This is observed in Fig. 6.h that shows that MPTCP generally transmits all packets through one AP or the other while the redundant configuration (Fig. 7.h) better adjusts the selection of the receiving AP to match the channel quality variations.

The redundant configuration reduces the end-to-end latency compared to MPTCP (the median is reduced by 67%). In fact, MPTCP increases the end-to-end latency compared to when packets are only transmitted through AP1 or AP2. MPTCP results in a median end-to-end latency of 385 ms compared to a median of 313 ms and 143 ms when transmitting using AP1 and AP2 respectively. This is illustrated in Fig. 9 that represents the box plot of the end-to-end latency measured when only one AP is considered (AP1 or AP2), and when implemented redundancy and MPTCP[9]. In Fig. 9, the red line within the box represents the median, the edges of the box the 25th and 75th percentiles, and the whiskers represent the 5th and 95th percentiles. MPTCP includes some internal processes (e.g. at the scheduler level) that result in additional processing delays. In addition, MPTCP sends signaling packets to estimate the quality of the links. These packets increase the load of the channel and contribute towards a poorer latency

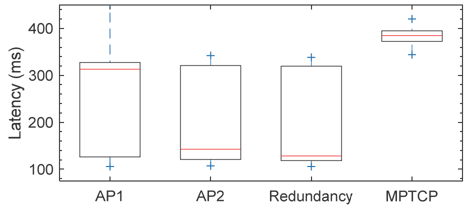

Fig. 9. Box plot of the end-to-end latency.

performance. Fig. 9 also shows that redundancy reduces the median of the experienced latency to 128.4 ms, i.e. it reduces the latency by 59.0% and 10.1% when compared to the use of AP1 and AP2 respectively[10].

The presented results show that MPTCP increases the reliability of industrial wireless communications but at the cost of degrading the latency. On the other hand, redundancy requires more bandwidth, but increases the reliability and reduces the latency (redundancy eliminates the AP selection and AP switching processes and so the delays due to these processes). Reliability and latency are the most important requirements for Industry 4.0. Redundancy is then considered the most appropriate solution to support mobile industrial applications, and the next section evaluates its capacity to combat interference.

B. *Robustness against interference*

Interference is introduced using the IN node depicted in Fig. 3.b. IN generates interference on the channel used by AP2. During the trial, the mobile robot follows the same trajectory as in previous trials. The interference is activated when the mobile robot is between cross-marks #5 and #6 (i.e. between t=69 s and t=110 s). Fig. 10 shows the performance achieved with the redundant configuration with interference. Similar trends to those explained in the previous section are observed when the interferer is switched off. Fig. 10.a shows that the RSSI at AP2 is still good (and higher than the RSSI at AP1) when IN is turned on. However, the PER experienced in the link with AP2 (Fig. 10.c) increases from 0.27% when the IN is not activated to 3.38% while IN is active (i.e. between t=69 s and t=110 s). The PER even reaches peak PER values up to 50% when the mobile robot is in the proximity of the IN. The interference also significantly degrades the latency of the packets received through AP2 (Fig. 10.f): some packets experienced end-to-end latency values higher than 1.5s. However, redundancy can maintain a connection with low PER (Fig. 10.d) and end-to-end latency (Fig. 10.g) even when the mobile robot is in proximity of the active interferer. Redundancy results in that only 0.5% of the transmitted packets experienced an end-to-end latency higher than 500 ms. This is in contrast to 20.4% of the packets with an end-to-end latency higher than 500 ms when only transmitting through AP2. Fig. 11 compares the CDF (Cumulative Distribution Function) of the end-to-end latency experienced under the presence of interference when packets are transmitted using redundancy or only through AP1 or AP2. The figure shows that redundancy can guarantee that 92.9% of the packets are transmitted in less than 350 ms. On the other hand, relying on

[9] The implemented solutions at CN1 and CN2 run at the user space of the Linux operating system, and therefore are managed together with other applications (processes or services) and kernel processes. This influences the end-to-end latency and lower latency values could be achieved with dedicated hardware and functions. Our prototype is though adequate to experimentally demonstrate the benefits that diversity and redundancy can provide to achieve Industry 4.0 QoS requirements.

[10] The internal processing of the packets from the application layer to the MAC layer entails a delay that represents on average 25-30% of the total measured end-to-end latency.

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/TII.2020.2979759

AP1 or AP2 exclusively can only guarantee that 86.9% and 40.8% of the packets respectively are transmitted in less than 350 ms. Redundancy achieves lower latency levels because it always utilizes the packet that correctly arrives first through AP1 or AP2. In this case, the redundant configuration uses the packets received through AP1 when IN interferes AP2. This is illustrated in Fig. 10.h that shows that most of the packets are

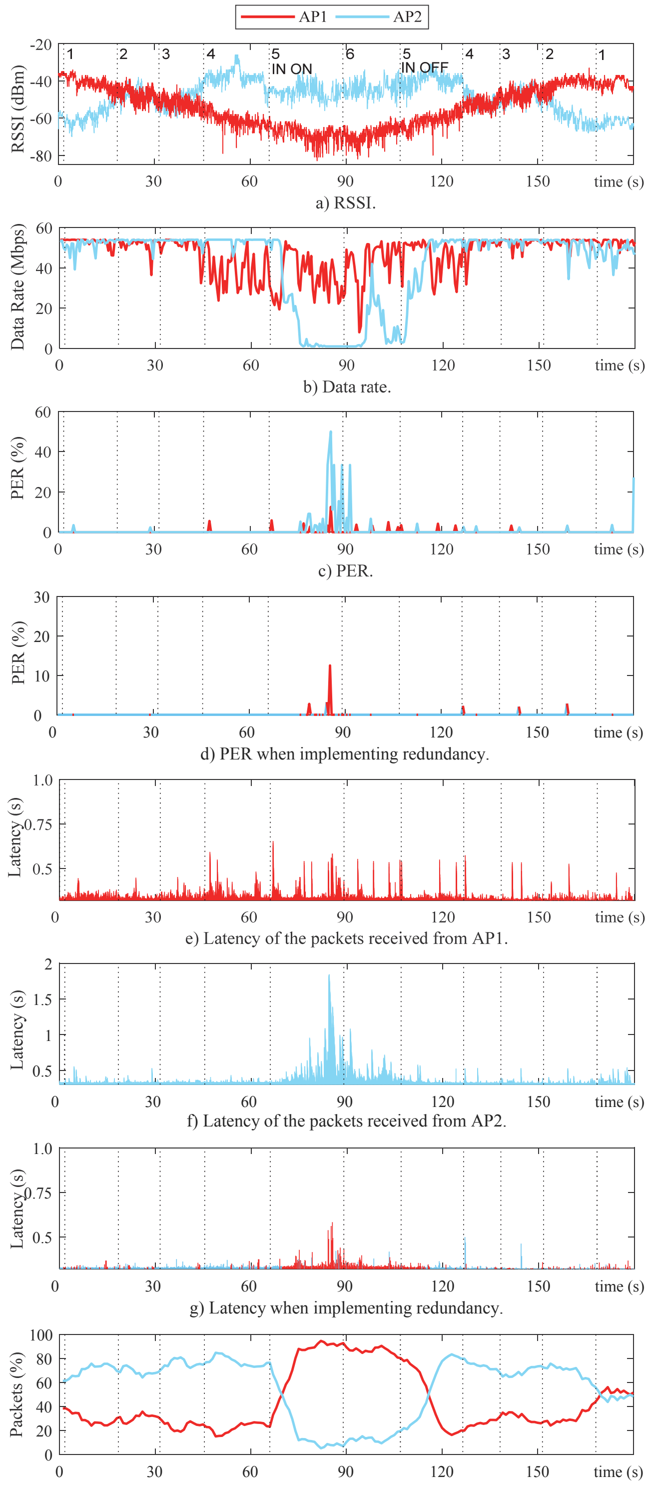

a) RSSI.
b) Data rate.
c) PER.
d) PER when implementing redundancy.
e) Latency of the packets received from AP1.
f) Latency of the packets received from AP2.
g) Latency when implementing redundancy.
h) Percentage of packets received through each AP.

Fig. 10. Comparison of redundant wireless communications with connections to a single AP under the presence of interference.

received through AP1 between t=69 s and t=110 s even if the RSSI at AP1 is lower than at AP2 (Fig. 10.a). The interference experienced in AP2 results in more packet errors due to packet collisions even if AP2 experiences a higher RSSI than AP1.

It should be noted that the interference also affects the permanent IEEE 802.11 networks deployed at the premises. Fig. 4 showed how the permanent networks change their channel when IN is switched on. The networks operating in channels 1 and 6 detect the interference and change to channel 11 that is the least affected channel by the interference. These changes increase the load on channel 11 and ultimately affect the performance with AP1 since it also operates in channel 11. However, the redundant configuration can overcome the increase in channel congestion by exploiting the variability of the radio channel and use AP2 to receive some packets during the high congestion period in channel 11 (Fig. 10.h). This can be appreciated in Fig. 11. The figure shows that the redundant solution reduces the percentage of packets that experience higher latencies compared to when only using AP2 (this AP is affected by interference) or AP1 (this AP is affected by congestion). For example, 59.2% and 13.1% of the packets experience latency levels higher than 350 ms when only AP2 and AP1 are used respectively. This percentage reduces to 7.1% considering the redundant configuration.

Fig. 10.e and Fig. 11 have shown that redundancy reduces the end-to-end latency compared to only communicating through AP1 or AP2. This is actually observed even during the time period when IN is switched on. Overall, the results in Fig. 10 show that redundancy can provide reliable and low latency industrial communications for mobile industrial applications even under the presence of interference.

C. *Service availability*

The wireless prototype supports communications between a dual-arm robot and a mobile robot. This use case is identified by the 3GPP as a standard mobile robot application, and is a relevant use case for Factories of the Future or Industry 4.0. The 3GPP establishes that standard mobile robot applications will require maximum latency levels of up to 500 ms. It is important that these latency levels are guaranteed during all or most of the connection between the dual-arm and mobile robots. Following [5], we define the communication service availability as the percentage of time that the end-to-end connection satisfies a target QoS level. We consider as target QoS level the correct reception of data packets before a maximum end-to-end deadline equal to 500 ms following 3GPP requirements for standard mobile robot applications [5]. Table I depicts the communication service availability when the mobile robot connects to a single AP (AP1 or AP2) and

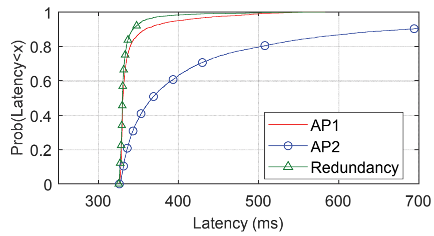

Fig. 11. CDF of the end-to-end latency experienced under the presence of interference.

when it utilizes MPTCP[11] or redundant connections. The results in Table I show that redundancy provides the highest availability under the evaluated conditions. Fig. 6.d and Fig. 6.e showed that MPTCP reduces link outage probability and increases the reliability of the wireless connection compared to connecting through a single AP. MPTCP also requires lower bandwidths than redundancy since packets are not duplicated. However, the reactive approach that MPTCP utilizes to select the link for transmitting packets affects the end-to-end latency and slightly degrades the service availability compared to the redundant solution.

Table I shows that using only AP2 can achieve high service availability levels that are close to those obtained with the redundant solution when interferences are not present[12]. Relying on the connection to a single AP reduces the deployment and bandwidth cost. However, network resilience can be compromised if the connection to that AP experiences any degradation. This is for example the case of the interference scenario analyzed in this study and also reported in Table I. Scenarios with variable signal quality and interference levels can be more frequent in future reconfigurable manufacturing plants. The results reported in Table I show that the service availability is significantly degraded when only relying on AP2 compared to when considering redundancy.

TABLE I
COMMUNICATION SERVICE AVAILABILITY (%)

Scenario	AP1	AP2	Redundancy	MPTCP
Without interference	61.38	99.99	99.996	99.83
With interference[13]	98.92	79.60	99.51	-

VIII. CONCLUSION

The development of the Industry 4.0 or Factory of the Future vision significantly relies on the availability of robust and low latency industrial communication networks. These must include wireless networks in order to support mobile industrial applications. This study has presented first experimental trials that demonstrate how redundancy can improve the reliability and latency of mobile industrial wireless communications even under the presence of interference. The study has demonstrated that redundancy achieves better reliability and latency performance than diversity when implemented using MPTCP. This protocol dynamically selects the link with the lowest RTT between the two communicating robots. This selection process entails some processing delay that ultimately affects the end-to-end latency. This is avoided with redundancy that transmits duplicated packets through different links. These simultaneous transmissions help combat channel variations and link outages without the processing delay incurred when having to detect and select the best possible link. The results obtained show that redundancy applied over IEEE 802.11 networks can guarantee relevant communication requirements identified for mobile Industry 4.0 applications.

[11] Results are not shown for MPTCP under interference conditions since redundancy achieved higher performance levels without interference and was hence selected as the reference solution to also combat interference.

[12] This shows that APs were adequately deployed to serve the target area. Extensive field measurements were conducted to select the most adequate locations for the APs. Such planning is necessary prior to any deployment.

[13] We only consider the packets transmitted during the time period the IN was activated.

REFERENCES

[1] European Factories of the Future Association (EFFRA), “Factories 4.0 and Beyond”, September 2016.
[2] 5G Alliance for Connected Industries and Automation, “5G for Connected Industries and Automation (White Paper-2nd Edition),” Feb. 2019. Available at: https://www.5g-acia.org/publications/5g-for-connected-industries-and-automation-white-paper/ (accessed on Oct. 2019).
[3] European Factories of the Future Association (EFFRA), “EFFRA vision for a manufacturing partnership in Horizon Europe 2021-2027”, 2019.
[4] S. Montero, J. Gozalvez and M. Sepulcre, “Neighbor discovery for industrial wireless sensor networks with mobile nodes”, *Computer Communications*, vol. 111, pp. 41-55, Oct. 2017.
[5] 3GPP, Technical Specification Group Services and System Aspects; Service requirements for cyber-physical control applications in vertical domains; Stage 1; (Release 16), 3GPP TS 22.104, v16.3.0, Sept. 2019.
[6] J. Gozalvez, M. Sepulcre, J.A. Palazon, "On the feasibility to deploy mobile industrial applications using wireless communications", *Computers in Industry*, vol. 65 (8), pp. 1136-1146, Oct. 2014.
[7] P. Popovski et al., "Wireless Access for Ultra-Reliable Low-Latency Communication: Principles and Building Blocks", *IEEE Network*, vol. 32, no. 2, pp. 16-23, March 2018.
[8] G. Cena, S. Scanzio, A. Valenzano, "Improving Effectiveness of Seamless Redundancy in Real Industrial Wi-Fi Networks", *IEEE Trans. on Industrial Informatics*, vol. 14, no.5, pp. 2095-2107, May 2018.
[9] M. Mohiuddin, et al., “Experimental validation of the usability of Wi-Fi over redundant paths for streaming phasor data”, in *Proc. 2016 IEEE International Conf. on Smart Grid Comms. (SmartGridComm)*, Nov. 2016, Sydney, Australia.
[10] G. Cena, S. Scanzio and A. Valenzano, "Seamless Link-Level Redundancy to Improve Reliability of Industrial Wi-Fi Networks", *IEEE Trans. on Industrial Informatics*, vol. 12, no.2, pp. 608-620, 2016.
[11] IEC 62439-3 Ed.02: Industrial Communication Networks - High availability automation networks - Part 3: Parallel Redundancy Protocol (PRP) and Hig-availability Seamless Redundancy (HSR), IEC, 2012.
[12] J.J. Nielsen, R. Liu, P. Popovski, "Optimized Interface Diversity for Ultra-Reliable Low Latency Communication (URLLC)", in *Proc. IEEE Global Comms. Conf. (GLOBECOM2017)*, Singapore, 2017, pp. 1-6.
[13] A. Ford, et al., "TCP extensions for multipath operation with multiple addresses", *Internet Requests for Comments RFC6182*, January 2013.
[14] M. Prakash, A. Abdrabou and W. Zhuang, "An Experimental Study on Multipath TCP Congestion Control With Heterogeneous Radio Access Technologies", *IEEE Access*, vol. 7, pp. 25563-25574, 2019.
[15] G. Prytz, “Network recovery time measurements of RSTP in an ethernet ring topology”, in *Proc. IEEE Conf. Emerg. Technol. Fact. Autom. (ETFA)*, Sep. 2007, pp. 1247–1253.
[16] A. Frommgen, et al., "ReMP TCP: Low latency multipath TCP", in *Proc. 2016 IEEE International Conference on Communications (ICC)*, May 2016, Kuala Lumpur, pp. 1-7.
[17] IEEE Standard for Information technology--Telecommunications and information exchange between systems Local and metropolitan area networks--Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007), March 2012.
[18] C. Paasch, S. Barre, et al., Multipath TCP in the Linux Kernel. Available at https://www.multipath-tcp.org.