

Improving Effectiveness of Seamless Redundancy in Real Industrial Wi-Fi Networks

Gianluca Cena, *Senior Member, IEEE*, Stefano Scanzio, *Member, IEEE*, and Adriano Valenzano, *Senior Member, IEEE*

Abstract—Reliability and determinism of Wi-Fi can be tangibly improved by means of seamless redundancy, to the point of making this technology suitable for industrial environments. As pointed out in recent papers, the most benefits can be achieved when no phenomena can simultaneously affect transmissions on all channels of a redundant link.

In this paper several aspects are analyzed which, if not properly counteracted, may worsen seamless redundancy effectiveness. Effects they cause on communication have been experimentally evaluated in real testbeds, which rely on commercial Wi-Fi devices. Then, practical guidelines are provided, which aim at preventing joint interference through a careful system design. Results show that measured communication quality can be made as good as expected in theory.

Index Terms—IEEE 802.11, parallel redundancy protocol (PRP), seamless redundancy, Wi-Fi, wireless channel independence, design guidelines.

I. INTRODUCTION

The past decade witnessed a dramatic growth of the IEEE 802.11 [1] technology (best known as Wi-Fi) as the standard wireless local area network (WLAN) solution in many application areas, including industrial environments [2]. Notable exceptions are real-time control applications running at the shop-floor, as Wi-Fi is currently deemed unable to satisfy determinism and reliability requirements demanded by time- and safety-critical systems. For this reason, other wireless solutions like WirelessHART [3] are mostly being considered in these scenarios. Nonetheless, a number of proposals have lately appeared which concern the use of Wi-Fi in distributed real-time control systems. For instance, [4] and [5] rely on deterministic overlays placed atop the IEEE 802.11 Medium Access Control (MAC) layer, those in [6], [7], and [8] foresee modifications to either the MAC retransmission scheme or the backoff mechanism, while [9] investigates the use of IEEE 802.11e Quality-of-Service (QoS). Besides the high throughput, the main advantage of Wi-Fi is that it can be interconnected with Ethernet directly at the data-link layer, which means that hybrid wired-wireless networks can be easily

This work was partially supported by Regione Piemonte and the Ministry of Education, University, and Research of Italy (MIUR) in the framework of the Call “Fabbrica Intelligente”, Project HuManS “Human centered Manufacturing Systems” (application number 312-36). Copyright (c) 2017 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org. The authors are with the National Research Council of Italy, Istituto di Elettronica e di Ingegneria dell’Informazione e delle Telecomunicazioni (CNR-IEIIT), I-10129 Torino, Italy (e-mail: gianluca.cena@ieiit.cnr.it, stefano.scanzio@ieiit.cnr.it, adriano.valenzano@ieiit.cnr.it).

set up [10]. Unlike [11], packet losses in wired portions of such systems are typically negligible, and upper bounds can be often enforced on the related transmission delays. This is no longer true for Wi-Fi links, which become the weakest point in the network and negatively affect the overall reliability.

As shown in recent literature, the adoption of seamless redundancy at the data-link layer of wireless networks improves their behavior significantly, under the basic assumption that disturbance and interference on physical channels are statistically independent. In [12] and [13], such a hypothesis was experimentally checked for redundant links based on real Wi-Fi equipment. While tangible improvements were achieved in typical operating conditions, channels were found not to be completely independent. Not surprisingly, this was not due to nearby wireless stations (STA), whose operations were actually uncorrelated, but depended on the design and configuration of the software and hardware modules of *redundant stations* (RSTA) in the testbeds.

In this paper, the most relevant design aspects that may cause joint and mutual interference among channels in a Redundant Basic Service Set (RBSS), i.e., a Wi-Fi network exploiting seamless redundancy, are identified. Then, suitable guidelines are provided, which can be used when industrial devices, communicating over an RBSS, have to be implemented and deployed, so as to maximize seamless redundancy benefits. With respect to [14], which only reported on preliminary and partial results, this paper offers a comprehensive and structured overview of all main causes of interference. Effects of the network infrastructure and deeper insights on the ways transmissions on air from nearby antennas may affect each other are also included. Moreover, a thorough assessment is made about the residual channel dependence, after such phenomena have been tackled by means of above guidelines.

The paper is organized as follows: Section II provides a quick survey of Wi-Fi seamless redundancy, while Section III describes the experimental setup we employed for evaluating the quality of communication. Section IV presents a general view of phenomena which may affect transmissions on redundant links, whereas Sections V and VI focus on interference that takes place inside devices and on air, respectively. Performance of a target solution, where all causes of interference have been dealt with properly, is experimentally assessed in Section VII. Finally, Section VIII summarizes practical design guidelines and draws conclusions.




Fig. 1. Practical implementation of PRP over Wi-Fi (PoW).

II. RELATED WORK

A number of solutions aimed at improving fault tolerance in Industrial Ethernet networks are based on seamless redundancy. For example, in the Parallel Redundancy Protocol (PRP) [15] two copies of each packet are sent by the originator on a pair of distinct, similar networks at the same time. The copy reaching the recipient first is retained, while the late copy is discarded. This approach can be profitably applied to wireless links (in particular, those based on Wi-Fi) in order to improve reliability and timeliness of data exchanges. In fact, the *packet loss ratio* (PLR) on the redundant link is noticeably lower than on either single physical channel, and also transmission latencies decrease.

Several proposals have appeared in the past few years where seamless redundancy is exploited in specific application scenarios. In [16], a hybrid wired-wireless redundancy scheme was envisaged to improve reliability in electrical substations, while redundant wireless paths were used in [12] for streaming phasor data over Wi-Fi. The IP parallel redundancy protocol (iPRP), introduced in [17], applies the same concepts to geographic networks for use with smart grids. In [18] seamless redundancy was proposed for clock synchronization protocols in WLANs, in order to both increase robustness against packet losses and mitigate the effects of malicious security attacks. In [19], PRP was applied to 4G Long Term Evolution (LTE) cellular networks in the context of railway wireless communications. Finally, in [20] Ultra WideBand (UWB) technology was considered, possibly on redundant channels, as a backup network for avionic systems, in order to reduce cables and weight.

In [21] and [22], an interesting implementation of seamless redundancy applied to IEEE 802.11, explicitly aimed at factory automation, is described. Such an arrangement, which is sketched in Fig. 1, relies on commercial off-the-shelf (COTS) Wi-Fi access points (APs), some of which operate in client mode, and special-purpose PRP devices (RedBoxes). In the following, it will be referred to as *PRP over Wi-Fi* (PoW). This solution is particularly appealing, as it can be easily employed to provide wireless extensions to industrial networks based on switched Ethernet, like EtherNet/IP and SafetyNET p. It is worth noting that, in PoW, packet duplication and deduplication at both ends of a redundant wireless link can also be carried out in software, with significantly lower implementation costs. The *Wi-Fi Redundancy* (Wi-Red) proposal [23] is quite similar to PoW, but it defines specific duplicate avoidance mechanisms, preliminarily introduced in




Fig. 2. Experimental setup for evaluating communication quality.

[24] and [25], which are aimed at reducing network traffic [26] and enhancing the overall system behavior.

As shown in [12] and [13], the most benefits can be obtained from seamless redundancy when phenomena that affect transmissions on distinct physical channels are completely independent. In this case, in fact, the probability that a packet is lost (or that it misses a certain deadline) on the redundant link can be obtained by multiplying the relevant probabilities, evaluated separately on each physical channel. For example, if the PLR on each single channel is 1%, the overall PLR when redundancy is exploited can be as low as 0.01%.

In theory, allocating channels of a redundant link to non-overlapping frequencies is enough to ensure their statistically independent behavior, as sources of interference (nearby STAs) and disturbance (narrowband noise) are uncoordinated. However, as pointed out in [14], achieving true independence also requires a number of practical aspects to be taken into account. Two main factors may worsen the performance of real redundant Wi-Fi networks: the first one is due to a number of protocol functions, and sometimes depends on the specific implementation of *software* modules of the protocol stack, whereas the second is caused by physical phenomena affecting *hardware* components, including radio transceivers and antennas. The related interference can occur either *inside* nodes (in both end devices and network equipment) or *on air*.

III. EXPERIMENTAL TESTBED

To investigate what aspects affect the quality of communication on redundant links, a simplified testbed was developed, which retains the same features as full-fledged PoW implementations like the one in [21]. Then, an experimental campaign was carried out and the effects of a number of implementation choices were analyzed. The aim is drawing a set of practical guidelines for system designers to maximize benefits achievable with seamless redundancy in Wi-Fi networks. The setup sketched in Fig. 2 is an extension of those described in [13] and [7]. Two identical Wi-Fi adapters are installed in a Linux PC and associated to two identical APs. The APs are in turn connected to the Gigabit Ethernet port of the PC through a switch. Equipment of different brands was used in the experiments, so as to ensure that results do not depend on the behavior of a specific device.

A measurement task running on the PC generates a cyclic stream of packets, and transmits them concurrently on the two channels. The PC acts as both source and destination of packets, which can be set to travel either from Wi-Fi adapters

to the Ethernet port or in the opposite direction. Two types of streams can be produced, unicast and multicast: the former is directed from STAs to APs, while the latter flows in the opposite direction. This is because multicast frames are sent by the originating STA to the AP as unicast. In the following, subscripts refer to stream types, while superscripts are used to identify channels (channels 1...13 belong to the 2.4 GHz band, while channels 36...165 lie in the 5 GHz band). For example, M_u^{c11} denotes a unicast stream transmitted on channel 11, whereas M_m^{c165} is a multicast stream on channel 165. As typical for process data in industrial applications, packets sent by the measurement task were quite small (the payload was set equal to 50 B). Unless otherwise specified, and to better adhere to the behavior of real-world networks, we left the rate adaptation mechanism enabled in the experiments. Therefore, the modulation and coding scheme, and hence the bit rate, were dynamically selected by the originating STA.

By taking timestamps on every single transmission and reception, the latency experienced by each packet (or, in case of errors, an indication that the packet went lost) can be determined on any channel. Besides transmission on air, latency measured on packet arrival at the Ethernet port, termed *end-to-end latency* (d), includes processing times in the PC, forwarding delays in both the AP and the switch, as well as transmission times on Ethernet cables. With the exception of delays introduced by the AP, which are analyzed in Section V-B, all these contributions were confirmed to be negligible. For example, processing times in commercial PCs [27], which constitute the second largest contribution to delays after the AP, are at least one order of magnitude smaller than the transmission latencies we measured. In addition, other timestamps were acquired upon reception of the relevant acknowledgment by each wireless adapter. They include both the interframe space (IFS) and the ACK frame, whose duration can be computed easily. By subtracting these durations from timestamps taken on ACKs, frame arrival times to the APs can be determined. In this way, a second latency can be defined for packets originating from wireless adapters, termed *link latency* (d'), which does not include contributions due to the network infrastructure.

All measures are collected at runtime in the main memory of the PC, saved to disk at the end of each experiment, and processed offline to evaluate the quality of communication. Starting from data acquired for each physical channel and by applying PRP rules (i.e., considering the copy of each packet which arrives to destination first), relevant performance indexes can be obtained for the redundant link too. In practice, transmission latencies d_i and d'_i for both physical channels and the redundant link were computed for every sample i concerning a successfully delivered packet in the experiment log file. Typical statistical indexes were evaluated for d_i , e.g., average value and standard deviation (\bar{d} and σ_d), minimum and maximum (worst-case) values (d_{\min} and d_{\max}), and percentiles (e.g., $d_{p99.9}$ and $d_{p99.99}$), and for d'_i as well. Depending on the experiment, statistics on either d_i or d'_i were used. More details on the testbed can be found in [13] and [7].

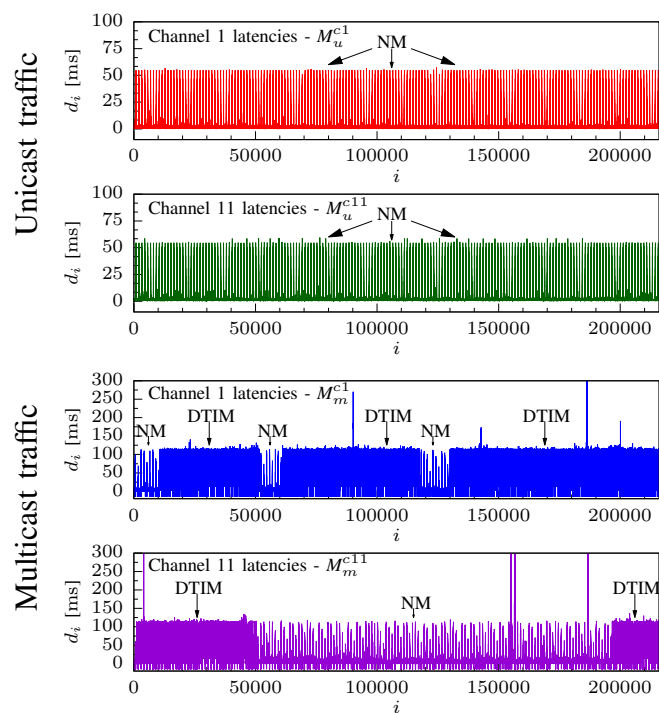


Fig. 3. End-to-end transmission latencies and packet losses in preliminary experiments for both unicast streams (top) and multicast streams (bottom).

IV. CONFIGURATION AND PRELIMINARY OPTIMIZATION

A first set of experiments were carried out by using the testbed “out-of-the-box”, i.e., without any specific optimization. As a matter of fact, conventional software modules and default equipment configuration resulted unsuitable for direct use in industrial applications. In particular, experiments highlighted the presence of interference affecting both channels of the redundant link, thus reducing expected performance.

A. Preliminary Evaluation

Two experiments, each one lasting one week, were performed to analyze the behavior of Wi-Fi seamless redundancy in a default configuration, using unicast and multicast services, respectively. They were considered as the baseline. Transmissions on the redundant link were modeled as a pair of concurrent synchronous streams, sent on channels 1 and 11. The generation period T_c was set to 100 ms. Fig. 3 summarizes experimental data in two significant intervals. The two topmost timing diagrams (M_u^{c1} , M_u^{c11}) concern the unicast case, whereas those at the bottom (M_m^{c1} , M_m^{c11}) refer to multicast. Each diagram spans over 6 hours of uninterrupted transmissions and shows the end-to-end latency d_i experienced by packet i on the related channel. In order to provide as much details as possible about transmissions, lost packets are visualized in the diagrams as negative latency values.

Two distinct interfering phenomena can be observed, which worsen communication determinism by causing packet latencies and losses. The first one is due to the IEEE 802.11 Delivery Traffic Indication Map (DTIM) mechanism, which concerns energy saving. Its observable effects, corresponding




Fig. 4. Example of multicast frame exchange (with and without DTIM).

to zones in diagrams for multicast streams M_m^{c1} and M_m^{c11} where interference is more dense, appear to be randomly distributed, because DTIM is activated when a STA that can enter the power-save state associates with the AP. In our experiments, since no authentication was required for the Basic Service Sets (BSS), any such STA entering the BSS range (e.g., mobile phones) might accidentally set DTIM on. Since this mechanism only affects multicast traffic, experimental results for unicast streams M_u^{c1} and M_u^{c11} did not exhibit this behavior. DTIM is analyzed in detail in Section IV-B. The second phenomenon is periodic and noticeably less intense than DTIM. However, it affected both multicast and unicast transmissions for the whole experiment. Investigations allowed us to conclude that it is due to a software module, running in the originating RSTA, which carries out a number of non-trivial functions, such as managing association/reassociation of the STAs with the AP and migration between adjacent BSSs (roaming). More details on this aspect are provided in Section V-A.

B. DTIM mechanism

According to the IEEE 802.11 standard, the DTIM mechanism is activated when at least one STA associated to the BSS enters the power-save mode. To ensure that all STAs are awake at the time the AP sends the multicast packets, they are buffered inside the AP and sent at specific time instants. The AP notifies the awakening time to the STAs by means of two fields, *DTIM period* (p) and *DTIM count* (c), included in the *beacon* frame, which is cyclically sent by the AP every T_{beac} (the typical value is 102.4 ms) and received by all STAs. Buffered multicast packets are sent out by the AP cyclically, with a period equal to $T_{\text{beac}} * p$. The value p is a configurable parameter of the AP (as T_{beac}), and allows to set the sleeping period of the STAs. The value c is read by STAs to know how many beacons have still to elapse before the AP sends the buffered packets. The value c is decreased by 1 by the AP on every beacon frame transmission and re-initialized to $c = p - 1$ every time it reaches value 0. In particular, $c = 0$ means that multicast packets transmission will follow the current beacon. Operation of the DTIM mechanism is summarized in Fig. 4.

1) *Effects of DTIM on communication:* End-to-end latencies and losses of two multicast streams M_m^{c1} and M_m^{c11} have been logged during an additional experiment lasting one hour. To provide higher resolution, packets were generated by the measurement task every $T_c = 10$ ms. Experimental data,




Fig. 5. Effects of the DTIM mechanism on end-to-end transmission latencies for multicast streams (unicast streams are not affected).

concerning a time interval where DTIM was active, are shown in Fig. 5, with timing diagrams 5-A and 5-B providing two different zoom levels. Latencies varied between ~ 15 ms and ~ 115 ms, according to a saw-tooth law whose period is about 100 ms. This behavior is compatible with the DTIM mechanism: multicast packets, sent on the Ethernet port, reach the switch and then the AP, where they are buffered until the next beacon. At the time of beacon transmission, a number of packets equal to about T_{beac}/T_c (~ 10 in our experiment) are queued in the AP. They are sent on air back-to-back just after the beacon frame, and this results in close arrival times (typically, consecutive received packets are spaced by less than 1 ms), whereas their transmission times are displaced by T_c . This explains the shape of diagram 5-B and why, between adjacent beacons, latency decreases by ~ 9 ms per sample.

DTIM may significantly worsen determinism of multicast packet streams. When it is active, their measured transmission latencies suffered from a jitter that, when $p = 1$, is as high as one hundred ms, and grows linearly for higher p values. This is hardly acceptable in real-time industrial applications. Even worse, the probability that both channels of a redundant link are affected by DTIM at the same time is not negligible (see, e.g., bottom diagrams in Fig. 3). As a consequence, effectiveness of seamless redundancy may be severely impaired.

2) *Possible remedies to DTIM issues:* A first, trivial rule to improve communication quality on the redundant link is to improve the performance of each physical channel separately. In particular, a simple remedy to circumvent the DTIM problem is to hide the Service Set Identifier (SSID), a string embedded in beacon frames that identifies the BSS. In this way, STAs not belonging to the control system are prevented from unintentionally associating with the AP. Unfortunately, this method only provides loose guarantees: in fact, a number of security attacks are known which target Wi-Fi networks configured this way. A better and safer approach relies on the use of authentication. Currently, the *de facto* standard is IEEE 802.11i, which standardizes the security protocol popularly known as Wi-Fi Protected Access II (WPA2). By using WPA2, STAs not involved in the control network cannot authenticate with the AP, hence preventing DTIM interference. Of course, only authentication is mandatory to avoid DTIM effects. However, since security attacks on industrial traffic may lead to harmful consequences, encryption is highly recommended.

In [14], experiments were carried out for measuring to what extent latencies d_i are practically affected by WPA2 when using recent commercial adapters. Results show that minimum

and mean values (d_{\min} and \bar{d}) were only marginally larger than an open system, the difference being in the order of a few μs . Worsening was slightly higher in terms of percentiles, but still negligible if compared to latency values. Hence, it can be safely assumed that the overhead introduced for encryption is mostly irrelevant in the light of the benefits, which justifies the use of WPA2 to prevent unwanted DTIM effects.

V. INTERFERENCE INSIDE DEVICES

Even when channels are set to operate on distinct frequencies (hence making frame transmissions on air totally independent, in theory), other sources of interference may affect communications in a joint fashion. This is because redundant transmission paths are not completely disjoint, but include overlapping portions located inside devices. In general, interference in the communication hardware (physical and MAC layers) is negligible, as wireless adapters are replicated and operate independently. However, a small amount of residual interference still exists, because operations involved in packet transmission and reception are (in part) managed in software by the same hardware components, that is, CPU, memory (RAM), and system bus (e.g., PCIe).

Exploiting concurrent tasks on multi-core CPUs may lessen this problem. Moreover, when fast CPUs are employed, as in modern PCs, this interference only causes short additional delays and jitters (in the order of few tens μs). For this reason, these contributions have not been considered explicitly in the remaining part of this paper. The same is no longer true when APs are taken into account, as they are usually based on embedded microcontrollers characterized by lower performance. Hence, internal interference in APs can affect the quality of communication much more than in the PC.

Besides the hardware components shared between redundant paths, software tasks may also cause non-negligible joint interference, when they repeatedly act on wireless adapters at the same time, e.g., to carry out management activities. A significant example of such a kind of behavior is described in the following section.

A. Internal interference due to network management services

To achieve independence between channels in a redundant link, wireless adapters should never be engaged in concurrent activities triggered by tasks other than the originator and recipient of packets. Unfortunately, this condition may occasionally occur in real systems. For instance, Linux includes a *network manager* (NM) service, which periodically scans all wireless channels to discover nearby BSSs (from a practical viewpoint, it looks for APs the STA may connect to). Similar services are foreseen in most other operating systems as well. The NM service is essential to ensure an adequate level of reliability. In fact, it allows STAs to reconnect to the AP when, because of disturbance conditions lasting long enough, they get disassociated. Moreover, NM also supports mobility between BSSs, hence achieving larger coverage areas. The ability to automatically reconnect STAs is practically mandatory in automation systems, to avoid situations where communications are permanently broken and human intervention is required.

On the contrary, mobility is needed only in specific industrial scenarios, when movements of STAs may exceed the coverage area of a single BSS, e.g., in the case of automated guided vehicles (AGV) and drones.

Periodic channel scanning, carried out by NM for all Wi-Fi adapters installed in the device, causes non-negligible performance degradation, as experimental results in Fig. 3 clearly highlight. In fact, while this task is in progress, application data cannot be sent and not even received. This implies that packets to be transmitted are temporarily buffered, and sent on air only when scanning is concluded. Unavoidably, this introduces delays for each single adapter.

In systems exploiting seamless redundancy, this issue can be prevented by purposely displacing in time the NM operations on different channels. Unfortunately, our experiments showed that, in conventional platforms like Linux PCs, scanning on different adapters took place at the same time. This makes interference on channels no longer independent, and worsens the communication quality of the redundant link consequently.

To evaluate NM effects, two experiments were performed, where the DTIM issue was prevented, by configuring the adapters to operate on channels 1 and 11, and using both unicast (M_u^{c1}, M_u^{c11}) and multicast (M_m^{c1}, M_m^{c11}) streams. Each experiment lasted one hour and, in both cases, $T_c = 10$ ms. Fig. 6 shows end-to-end latencies and losses for the adapter tuned on channel 1. Timing diagrams in the upper part of the figure concern unicast frames, while those in the lower part are about multicast. Results for the second adapter are similar, and hence they are not shown in the figure.

1) *Unicast transmission*: Diagram 6-A (top-left) refers to a 10-minute-wide interval. Periodic interference, in the form of peaks of transmission latency d_i , appeared every two minutes for the whole duration of the experiment. They are due to the NM service running in the PC. Diagram 6-B (top-middle) shows a single scan performed by NM across all channels. Its duration, measured between the outermost peaks, is ~ 3 s (308 packets are affected). It includes 13 peaks, corresponding to the probe operations NM carried out on each one of the 13 channels legally allowed for Wi-Fi in Europe, according to the European Telecommunications Standards Institute (ETSI) regulations [28]. Finally, diagram 6-C (top-right) shows an expanded view of a single peak, and depicts the effects of a specific probing operation.

When NM tunes the RF module of the adapter used to transmit stream M_u^{c1} on a new channel (e.g., 5), the measurement task temporarily loses its ability to send the related packets, which are temporarily buffered in the STA's adapter. If the number of packets generated during NM activity exceeds the transmission buffer capacity, some packets may be possibly discarded. As soon as NM switches the RF module back to channel 1 and re-enables communication, buffered packets are sent on air, in the same order as they were queued. Probing in diagram 6-C affected 6 consecutive packets: the "older" packet suffers from the largest latency (~ 60 ms), and delays decrease on each following packet by about T_c (10 ms) minus the time to transmit one frame on air, which in our setup was in the order of one hundred μs . Not all peaks in diagram 6-B have exactly the same height, because the generation process

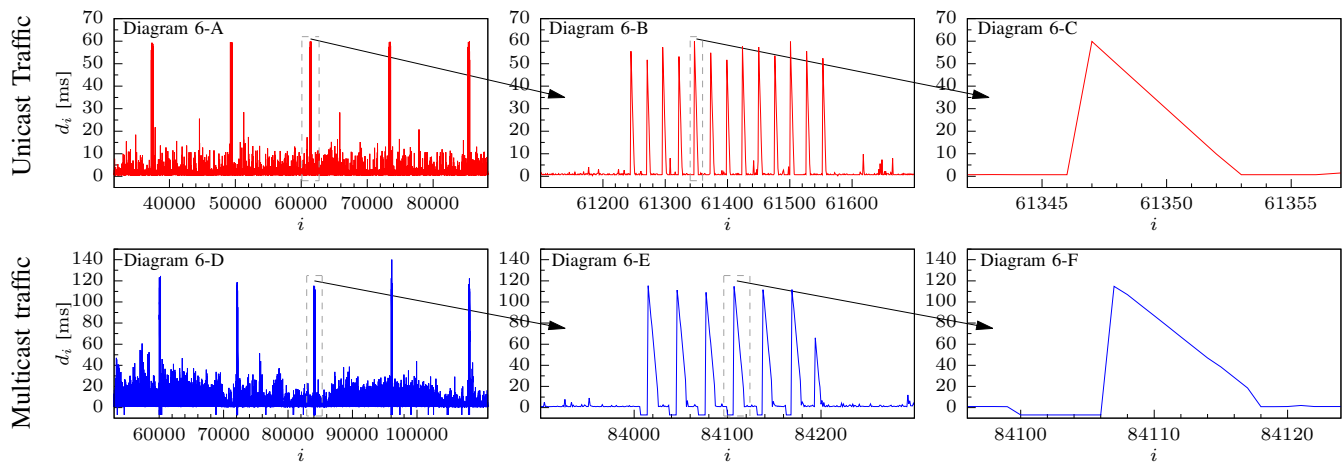


Fig. 6. Effects of the network manager service on end-to-end transmission latencies and losses for both unicast streams (top) and multicast streams (bottom).

of M_u^{c1} and the activation interval of NM are not synchronized.

2) *Multicast transmission:* Results for multicast packets are similar to the unicast case, with some peculiar differences due to the fact that, in our setup, they flow from APs to the STAs. When NM is active, packets are dropped because the RF module in the STA is listening to a different channel. Unlike the unicast case, where this issue is counteracted by the AP using retransmissions, no retry is carried out for multicast frames as they are unconfirmed. As for the baseline case, in diagrams 6-D, 6-E, and 6-F, dropped packets are represented as negative latency values.

Diagram 6-E concerns a complete scan of all channels. Only 7 peaks are observed instead of 13: this is because, before switching the RF module frequency, NM temporarily activates the DTIM mechanism in the AP, in an attempt not to lose received multicast packets. This causes packets, generated by the measurement task in a time interval comprising scans on several channels (two, in the case of diagram 6-F), to be buffered by the AP and sent immediately after a beacon. For this reason the height of peaks in diagrams referred to multicast streams is about twice the unicast case. The last peak is shorter (~ 60 ms) because only one channel is concerned. Although the particular shape of diagrams depends on the specific implementation of NM in Linux, similar amounts of interference are nevertheless expected in different operating systems to carry out the related functionality.

3) *Possible remedies to NM interference:* Although a network scanning service has to be mandatorily implemented in IEEE 802.11, scanning sequences are not typically triggered automatically by the adapter, but are invoked on demand by the NM. Disabling NM [29] is the easiest solution to avoid above described issues. This has been done in all experiments reported in the remaining part of this paper, so that results are not affected by NM interference. However, the NM service is fundamental for industrial Wi-Fi networks, where a permanent loss of connectivity cannot be tolerated. In fact, when a STA experiences a disconnection from the AP, it is unable to reassociate without the NM intervention. Both sides of a wireless link (STA and AP) may force a disconnection, if

no communication takes place for a certain lapse of time. This condition occurs frequently for mobile nodes (i.e., those allowed to migrate to areas covered by other APs), but it may also happen for unmovable nodes (including those whose movements are confined inside the coverage area of a single AP), due, e.g., to prolonged interference or disturbance on air. Solutions other than disabling NM are thus needed.

If properly exploited, network redundancy can offer tangible advantages for NM operations. In particular, it can lessen the impact this service has on data transfers, as well as the likelihood that all channels become disconnected at the same time. To this aim, specific NM implementations, tailored to such application contexts, have to be developed, which take into account their specific characteristics and, in particular, whether devices are unmovable or mobile. Channel scanning is not strictly required for unmovable STAs, in which case network parameters can be configured statically. As an alternative, the discovery procedure can be executed only once, at network startup. In both cases, NM has only to monitor the STA connection state, and reassociate to the same AP in the shortest possible time when a disconnection occurs.

When mobility is required, NM has to periodically scan all channels involved in the redundant link to acquire information about nearby APs. Importantly, this operation must never be carried out on the two adapters at the same time. In a similar way, when a device wishes/needs to migrate to another AP, NM has to carry out reassociation on one channel at a time. Only when this operation is concluded on one adapter, it can be started on the other adapter too. This ensures that, during migration, communication is never interrupted. A number of heuristics can be devised to manage roaming, aimed at deciding timings and the channel order used for migration, based on information provided by the MAC, e.g., the signal-to-noise ratio or the PLR on each channel. These peculiar aspects fall outside the scope of this paper and are left for future investigations.




Fig. 7. CCDF of end-to-end latencies on channels 1 (d^{c1}) and 165 (d^{c165}).

TABLE I
EFFECTS OF NETWORK INFRASTRUCTURE ON SEAMLESS REDUNDANCY:
ONE SIMULTANEOUS DUAL-BAND AP VS. TWO SINGLE-BAND APs.

Exp.	Stream	\bar{d}	σ_d	$d_{p99.9}$	$d_{p99.99}$	d_{max}	PLR [%]
		[ms]					
IAP	M_u^{c1}	1.267	1.685	20.064	27.775	41.044	0.0005
	M_u^{c165}	0.346	1.179	18.639	20.425	21.466	0.0
	$M_u^{c1 \cup c165}$	0.346	1.179	18.639	20.425	21.466	0.0
2APs	M_u^{c1}	1.236	1.730	20.071	27.483	40.939	0.0005
	M_u^{c165}	0.304	1.173	18.457	20.241	20.430	0.0
	$M_u^{c1 \cup c165}$	0.210	0.133	2.077	5.110	10.839	0.0

B. Internal interference in network equipment

As shown in Fig. 2, the simplest way to set up an infrastructure RBSS is deploying a pair of APs (configuration *2APs*). As an alternative, a single simultaneous dual-band AP (an AP provided with two radio blocks and two MAC entities, operating on distinct channels at the same time) can be used (configuration *IAP*). This permits the roles of the two APs in Fig. 2 to be functionally merged in the same device. Unlike RSTAs implemented on PCs, where the code of the protocol stack is often available, and can be customized to cope with phenomena causing joint interference on redundant paths, AP firmware is seldom accessible, especially for COTS devices.

To evaluate the effects of internal interference in the network equipment, we configured a single AP (NETGEAR WAC120) to operate simultaneously on channel 1 in the 2.4 GHz band and 165 in the 5 GHz band. A cyclic unicast traffic (M_u^{c1} , M_u^{c165}) with period $T_c = 100$ ms and directed to the AP was generated by the PC. The experiment lasted 6 hours and results are reported in the top rows of Table I, labeled *IAP* (clearly, end-to-end latencies were considered). The very same experiment was repeated using one AP per channel (two identical NETGEAR WAC120 were employed in this case). Results are included in the second set of rows, labeled *2APs*.

1) *Forwarding delays*: M_u^{c1} statistical indexes are similar in the two experiments, and the same holds for M_u^{c165} . Differences were due in part to the fact that the wireless medium is not stationary (i.e., channel characteristics change over time). Main causes are disturbance (caused by electromagnetic noise, moving objects, multipath fading, attenuation, etc.) and interfering traffic generated by nearby Wi-Fi networks, both of which are out of the experimenter's control. For the *IAP* configuration, only 0.18% of the packets arrived first on channel 1, which causes statistical indexes of stream $M_u^{c1 \cup c165}$

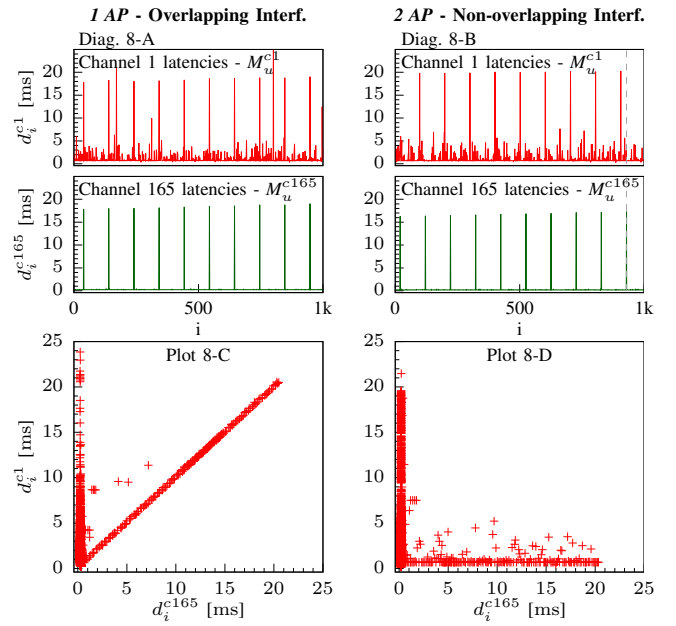


Fig. 8. End-to-end transmission latency on channels 1 (d_i^{c1}) and 165 (d_i^{c165}), shown using timing diagrams (top) and scatter plots (bottom), for *IAP* and *2APs* configurations.

(which models the redundant link) to be almost the same as for channel 165 (see second and third lines of Table I).

Fig. 7 shows the *complementary cumulative distribution functions* (CCDF), measured in experiments *IAP* and *2APs*, for both channels 1 and 165. Overall, latencies in the 2.4 GHz band are worse than in the 5 GHz band, mostly because of the higher interference on air. More in detail, the CCDFs of d^{c1} (plot on the left) are roughly the same for the two experiments. Instead, the CCDF of d^{c165} (plot on the right) in *IAP* is shifted to the right by $\sim 40 \mu s$ with respect to *2APs*, likely because of the additional overhead for managing operations in both bands. Thus, when a single AP is employed, the time taken to forward packets received on the 5 GHz band increases, on average, by the same amount, which directly worsens seamless redundancy performance. This is also confirmed by mean latency values on channel 165 (\bar{d}^{c165}), equal to $0.346 \mu s$ and $0.304 \mu s$ for configurations *IAP* and *2APs*, respectively.

2) *Joint interference*: A more serious phenomenon that affects seamless redundancy is shown in the timing diagrams in the upper part of Fig. 8. As can be seen, some internal activity is cyclically executed in our APs every 10 s, which interferes with packet forwarding. Packets reaching a busy AP suffer from additional delays that, in the worst case, can be as long as 20 ms. If the same AP manages both wireless networks (*IAP*), this interference always occurs at the same time on both channels (see diagrams 8-A), hence affecting the redundant link to the same extent. Conversely, interference seldom occurs in a joint manner when two distinct APs are employed (*2APs*), since their operation is not coordinated (see diagrams 8-B, where peaks are displaced in time).

Statistical dependence can be better analyzed by means of the scatter plots in the bottom part of Fig. 8, where each

point represents a packet and its XY coordinates correspond to latencies d_i^{c165} and d_i^{c1} experienced on channels 165 and 1, respectively. Most packets suffer from negligible delays, and the related points fall close to the plot origin. As plot 8-D shows, in the $2APs$ configuration several points are placed near the X and Y axes. This is because d_i^{c165} and d_i^{c1} are independent and, as shown in diagrams 8-B, time intervals where interference occurs seldom overlap, since the time bases in the two APs are not synchronized. Points lying near the identity line ($X = Y$) in plot 8-C, which concerns the IAP configuration, correspond to packets that experienced similar, non-negligible delays on both channels (i.e., when $d_i^{c165} \simeq d_i^{c1}$). They confirm the presence of joint interference, which delays the copies of the same packet by the same amount of time on both channels. When this behavior is prevented, as in the $2APs$ case, statistical indexes about the latency on the redundant link improve noticeably. For example, as shown in Table I, the 99.9 and 99.99 percentiles decrease from 18.639 and 20.425 ms in the IAP configuration to 2.077 and 5.110 ms for the $2APs$ configuration, respectively.

3) *Possible remedies to AP internal interference*: The above examples put into evidence some aspects, due to infrastructure network equipment, which might impair the effectiveness of seamless redundancy. An obvious suggestion is that, in (soft) real-time applications, devices have to be properly characterized before their actual deployment. Unfortunately, this can result in a complex task and takes significant time and effort. A slightly more expensive alternative is relying on distinct APs to reduce the common parts of redundant paths, so that the likelihood that interfering phenomena affect all copies of the same packet decreases.

VI. INTERFERENCE ON AIR

Experiments with real devices revealed that another kind of interference exists for physical channels, which could affect frame transmissions on air. An in-depth analysis highlighted that a residual statistical dependence was caused by the relative positions of antennas which, when placed close to one another, induced a phenomenon known as Adjacent Channels Interference (ACI) [30]. In our setup, wireless adapters were fastened to two PCIe slots of the same PC, so that antennas were located a couple of centimeters apart. This condition is not so unrealistic, since real industrial devices are typically provided with small-sized enclosures, and antennas can not always be separated by means of cables.

A. Evaluation of Adjacent Channel Interference

To evaluate ACI, experiments based on unicast traffic were carried out using the setup with two APs of Fig. 2, after NM interference in the PC was prevented. Link latencies d' measured on ACK arrivals were considered. The packet loss ratio was determined both on the PC's Ethernet port (PLR) and by counting packets for which no confirmation (ACK frame) was received by Wi-Fi adapters (PLR'). Since ACK is not returned for lost frames, $PLR' \geq PLR$. However, in typical operating conditions, the likelihood that, for a given packet transmission, all the related ACK frames are corrupted

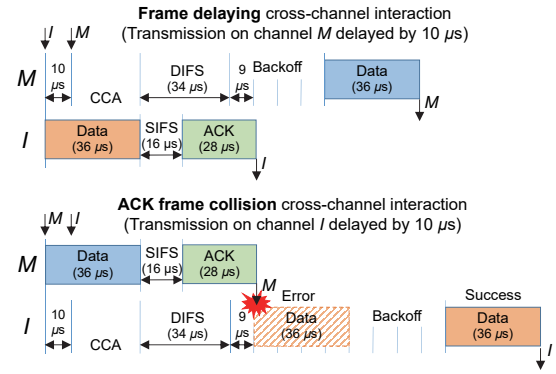


Fig. 9. ACI effects on transmission: frame delaying and ACK frame collision.

is very low, and so $PLR' \simeq PLR$. To make results clearer, we disabled MAC layer retransmissions (i.e., we relied on one-shot acknowledged frames) and selected a fixed 54 Mbit/s bit rate.

In the following, the two physical channels are conventionally referred to as *channel-under-test* (M), on which measurements about packets were collected, and *interfering channel* (I), while the related adapters are denoted STA_M and STA_I , respectively. STA_M was set to channel 165 in the 5 GHz band, while the operating frequency of STA_I was varied. Packets were repeatedly transmitted on M by the measurement task with period $T_c = 100$ ms. On every second packet on M, a similar packet was also sent on I. The sequence of packets on M, defined as $\{m_1^A, m_2^A, m_3^A, m_4^A, \dots\}$, consists of two distinct subsets, namely, $\bar{A} = \{m_1^A, m_3^A, \dots\}$, which contains packets sent only on the channel-under-test, and $A = \{m_2^A, m_4^A, \dots\}$, which includes those packets concurrently sent on both channels. Interleaving simplex and duplex streams allows to reliably compare ACI effects in almost the same environmental conditions, despite the wireless spectrum varies unpredictably over time. Statistics on experimental samples in the two sets were compared to determine the impact of ACI on latencies and losses.

B. ACI Effects on Seamless Redundancy

Not all ACI effects on packet transmission can be easily predicted, because some of them depend on very specific behaviors of the physical layer. Generally speaking, interfering frames sent on non-overlapping adjacent channels by nearby antennas may interact with the Clear Channel Assessment (CCA) mechanism in such a way that the originator could wrongly sense the channel as busy, even when it is idle. More rarely, they might even corrupt ongoing frame receptions. Two relevant (although not exhaustive) examples of cross-channel interaction, we termed *frame delaying* and *ACK frame collision*, have been experimentally analyzed.

1) *Frame delaying*: To make this phenomenon more apparent, transmissions on I were set to occur slightly in advance ($10 \mu s$) with respect to the corresponding ones on M. This situation is sketched in the topmost diagram of Fig. 9. When a request is issued for a frame transmission on M, STA_M finds the channel busy, due to the ongoing transmission on

TABLE II
 ADJACENT CHANNEL INTERFERENCE BY VARYING THE INTERFERING CHANNEL: FRAME DELAYING AND ACK FRAME COLLISION EFFECTS.

Frame delaying cross-channel interaction (transmissions on channel M delayed by 10 μ s)																								
Ch. M	Ch. I	Samples suffering from ACI (channel 165, A set)							Samples unaffected by ACI (channel 165, \bar{A} set)															
		\bar{d}'	$\sigma_{d'}$	d'_{\min}	d'_{p50}	d'_{p99}	$d'_{p99.9}$	d'_{\max}	PLR'	PLR	\bar{d}'	$\sigma_{d'}$	d'_{\min}	d'_{p50}	d'_{p99}	$d'_{p99.9}$	d'_{\max}	PLR'	PLR					
		[ms]							[%]							[ms]							[%]	
165	161	0.265	0.645	0.059	0.176	5.369	5.526	5.728	0.089	0.089	0.070	0.059	0.066	0.068	0.080	0.289	4.376	0.133	0.133 (a)					
165	157	0.246	0.536	0.060	0.177	3.702	5.630	5.913	0.056	0.056	0.070	0.049	0.061	0.068	0.074	0.285	3.926	0.278	0.267					
165	153	0.088	0.128	0.057	0.070	0.240	2.662	4.801	0.050	0.050	0.069	0.014	0.061	0.068	0.077	0.263	1.104	0.011	0.011					
165	149	0.113	0.349	0.058	0.069	2.578	5.017	5.311	0.028	0.017	0.069	0.030	0.061	0.068	0.074	0.255	3.627	0.017	0.017					
165	48	0.073	0.057	0.059	0.070	0.085	0.274	5.456	0.011	0.011	0.069	0.018	0.066	0.068	0.077	0.301	1.403	0.011	0.006					
165	36	0.073	0.081	0.059	0.069	0.085	0.319	5.248	0.022	0.022	0.069	0.015	0.062	0.068	0.073	0.272	1.232	0.039	0.039					
165	1	0.070	0.013	0.066	0.069	0.085	0.262	1.003	0.039	0.039	0.070	0.013	0.066	0.069	0.085	0.229	1.066	0.122	0.122					

ACK frame collision cross-channel interaction (transmissions on channel I delayed by 10 μ s)																								
Ch. M	Ch. I	Samples suffering from ACI (channel 165, A set)							Samples unaffected by ACI (channel 165, \bar{A} set)															
		\bar{d}'	$\sigma_{d'}$	d'_{\min}	d'_{p50}	d'_{p99}	$d'_{p99.9}$	d'_{\max}	PLR'	PLR	\bar{d}'	$\sigma_{d'}$	d'_{\min}	d'_{p50}	d'_{p99}	$d'_{p99.9}$	d'_{\max}	PLR'	PLR					
		[ms]							[%]							[ms]							[%]	
165	161	0.069	0.033	0.060	0.068	0.075	0.355	4.902	0.334	0.012	0.069	0.032	0.065	0.068	0.074	0.285	4.889	0.019	0.014					
165	36	0.070	0.011	0.058	0.069	0.077	0.251	1.484	0.192	0.014	0.069	0.011	0.057	0.068	0.075	0.249	1.685	0.014	0.014					
165	1	0.069	0.135	0.058	0.068	0.075	0.244	44.425	0.007	0.007	0.069	0.014	0.057	0.068	0.075	0.246	2.994	0.016	0.016					

I. As soon as STA_M senses the channel idle, it waits for a DIFS and the following backoff before starting transmission on air. As a consequence, frame transmission on M is postponed by the duration of the frame on I plus a DIFS and the random backoff selected by STA_M. Interestingly, channel M is sensed idle by STA_M upon completion of the interfering data frame, and not the related ACK, because the latter comes from the AP, which is located some meters apart from the PC.

The experiment consisted in a number of trials, each one lasting one hour, carried out by changing the interfering channel I (161, 157, 153, 149, 48, and 36 in the 5 GHz band, and 1 in the 2.4 GHz band). This permitted to evaluate how much the channel-under-test M (165) is affected by ACI due to channel I vs. the offset between channel frequencies. Results are reported in the upper part of Table II, where each row concerns a trial. Results clearly show that ACI effects appear, to varying degrees, for almost all channels in the 5 GHz band.

The closer the channels are, the more the latency increases. For instance, when I is set to channel 161, the average link latency \bar{d}'_A^M for packets on M in the A set is 0.265 ms, much larger than in the \bar{A} set, where $\bar{d}'_A^M = 0.070$ ms. Even bigger differences between the two sets can be evinced from percentiles. The CCDF and the *probability density function* (PDF) of link latencies \bar{d}'_A^M and $\bar{d}'_{\bar{A}}^M$ for this specific experiment are reported in Fig. 10. They show two important facts: first, the CCDF changed sensibly because of ACI, with a tangible increase in latencies; second, the PDF reveals that, as expected, the majority of the packets in the A set are delayed by the backoff interval, whose effects correspond to the peaks in the zoomed area [31].

ACI effects gradually decrease when the frequency offset between channels M and I increases, and disappear completely when I is set to operate on a different band (2.4 GHz).

2) *ACK frame collision*: To better highlight the ACK frame collision issue, transmission requests on I were delayed by 10 μ s with respect to M. The interaction between frames sent




Fig. 10. CCDF and PDF of measured link latencies for experiment (a) in Table II (channel 165 vs. 161) to highlight the frame delaying effect.

on the two channels is exemplified in the lower time diagram of Fig. 9. In particular, STA_I senses the channel busy and defers its transmission (situation is specular with respect to the previous experiment). After one DIFS plus the backoff interval, computed from the end of the data frame on M, a copy of the same packet is also sent on I. If, depending on the random backoff selected by STA_I, frame transmission on I overlaps with the ACK frame on M, reception of the latter by STA_M might sometimes fail because of proximity of antennas.

Results of experiments are reported in the lower part of Table II. As expected, transmission latencies on M are unaffected. Conversely, ACI directly worsens PLR' (but not PLR). In fact, because of missed ACK reception, STA_M is wrongly led to believe that the packet was not correctly delivered to the AP, which in turn triggers a retransmission attempt. However,

TABLE III
 RESULTS AFTER OPTIMIZATION (UNICAST TRANSMISSIONS).

Experiment	Stream	Statistics derived from the PC Ethernet port						Statistics derived from ACK frame arrivals						
		\bar{d}	σ_d	$d_{p99.9}$	$d_{p99.99}$	d_{max}	PLR	\bar{d}'	$\sigma_{d'}$	$d'_{p99.9}$	$d'_{p99.99}$	d'_{max}	PLR'	
		[ms]						[ms]						
No Load	Night	M_u^{c1}	1.040	1.442	19.567	25.523	72.141	0.00250	0.335	0.860	7.950	24.992	71.512	0.00250
		M_u^{c165}	0.302	1.160	18.416	20.234	20.378	0.0	0.079	0.052	0.255	0.409	29.220	0.0
		$M_u^{c1Uc165}$	0.209	0.099	1.394	3.864	8.249	0.0	0.077	0.050	0.182	0.319	29.220	0.0
	Day	M_u^{c1}	1.684	5.514	41.478	178.129	808.210	0.00111	0.994	5.640	41.326	177.989	843.950	0.00111
		M_u^{c165}	0.303	1.160	18.425	20.250	20.453	0.0	0.076	0.018	0.254	0.392	1.854	0.0
		$M_u^{c1Uc165}$	0.213	0.211	2.455	10.123	19.943	0.0	0.074	0.013	0.176	0.324	1.854	0.0
Low Load	Night	M_u^{c1}	1.049	1.469	19.632	24.980	85.837	0.000278	0.342	0.853	8.304	22.519	85.194	0.000278
		M_u^{c165}	0.437	1.250	18.959	21.087	22.955	0.0	0.220	0.312	2.486	6.087	14.920	0.0
		$M_u^{c1Uc165}$	0.304	0.220	1.953	4.246	14.687	0.0	0.099	0.114	1.261	1.993	5.280	0.0
	Day (b)	M_u^{c1}	1.236	1.750	20.207	29.519	86.122	0.00250	0.534	1.306	13.486	28.441	85.511	0.00278
		M_u^{c165}	0.436	1.240	19.015	21.042	22.304	0.0	0.219	0.302	2.411	3.924	17.140	0.0
		$M_u^{c1Uc165}$	0.309	0.245	2.388	5.645	17.878	0.0	0.108	0.136	1.468	2.235	3.730	0.0
High Load	Night	M_u^{c1}	1.125	1.566	19.823	27.197	46.916	0.000833	0.423	1.047	10.178	27.294	46.294	0.000833
		M_u^{c165}	1.955	5.370	79.741	123.427	195.071	0.0	1.746	5.236	79.555	123.332	194.985	0.0
		$M_u^{c1Uc165}$	0.606	0.523	5.322	13.218	27.041	0.0	0.204	0.425	4.195	6.869	26.436	0.0
	Day (c)	M_u^{c1}	1.285	4.521	20.306	99.607	814.861	0.00250	0.584	4.367	13.372	99.030	812.695	0.00250
		M_u^{c165}	1.746	4.554	65.367	119.014	160.995	0.0	1.539	4.395	65.105	118.924	160.891	0.0
		$M_u^{c1Uc165}$	0.624	0.555	5.696	11.893	22.976	0.0	0.236	0.477	4.572	8.351	22.866	0.0

since in the experiments we disabled automatic retransmission, a discrepancy appears concerning the transmission outcomes as seen on the originating and recipient sides. For the \bar{A} set, the values of PLR and PLR' are practically identical. Instead, when frequencies of channels are close, for the \bar{A} set there is a sensible increase of the number of dropped ACK frames (PLR' = 0.334% vs. PLR = 0.012%). Once again, this phenomenon completely disappears when channels are set to operate on different bands.

3) *Possible remedies to ACI*: In the above experiments, transmission times on M and I were displaced by $\pm 10 \mu s$ in order to always induce either frame delaying or ACK frame collision phenomena. In seamless redundancy implementations where packets are queued in the adapters of the originator at the same time, a mix of the above phenomena (plus, possibly, others) may take place, which worsens both latencies and losses. Variability of the related effects drastically increases due to retransmissions and rate adaptation techniques, which are typically enabled in real-world networks. As a consequence, to increase predictability and performance of seamless redundancy, ACI has to be prevented by setting redundant adapters on widely spaced channels or, even better, on different bands. Alternatively, transmissions on redundant physical channels by the same RSTA must be adequately displaced in time, so as not to overlap.

VII. PERFORMANCE OF OPTIMIZED REDUNDANT LINKS

A final experimental campaign was performed after all the guidelines described in the previous sections were applied to the testbed. Its aim is evaluating benefits that optimizations provide on seamless redundancy performance, by making communication paths practically independent. In detail:

- 1) DTIM effects were prevented for multicast traffic,
- 2) the NM service was disabled in the PC,
- 3) two distinct APs were used as network infrastructure, and,
- 4) channels were set to operate on different bands.

Experiments were carried out both at *Night* and by *Day*, and lasted 10 hours each. Two packet streams M^{c1} and M^{c165} were sent on channels 1 and 165, respectively, and generation period was set to $T_c = 100$ ms. Since no STAs were found on channel 165, besides those involved in the experiments (in the PC and the AP), load conditions on the two channels were pretty much unbalanced. This is the worst operating condition for links based on seamless redundancy, because the resulting communication quality tends to approach that of the best physical channel. For this reason, besides the *No Load* condition, we also defined the *Low Load* and *High Load* conditions, for which additional traffic was purposely injected on channel 165. Both interfering loads are based on packet bursts, where packets with 1500 B payload size are generated every 400 μs . The number of packets in a burst and the gap between consecutive bursts are modeled according to exponential distributions, with mean values 300 packets and 200 ms, respectively. The difference between *Low Load* and *High Load* conditions lies in the number of interfering nodes, i.e., 1 and 3, respectively.

A. Quality of Communication on the Redundant Link

Both unicast and multicast traffics were taken into account. In the former case, we evaluated statistics on both the end-to-end latency d measured on packet arrival at the PC Ethernet port (after traversing the AP) and the link latency d' obtained from ACK frame reception on wireless adapters. For multicast traffic, only results about end-to-end latency d are available, because this kind of traffic is not confirmed.

1) *Unicast transmission*: Results are summarized in Table III. As expected, tangible improvements were achieved by seamless redundancy in all experimental conditions and for every statistical index. As an example, for the *High Load-Day* condition, percentile $d_{p99.99}$ decreased by about one order of magnitude compared to single channels. Slightly lower advantages were achieved when individual channels are highly unbalanced, as in the *No Load-Day* condition.

Statistics on the end-to-end latency d (left side of the table) mostly concern the point of view of applications. Think, e.g., to the communication quality perceived by a device (wireless sensor, AGV, drone, etc.) that communicates on air, by means of redundant APs, with a programmable logic controller (PLC) connected to the wired backbone. In this case, results are biased by the characteristics of the specific APs used in the experiments. For example, in our setup, every 10 s each AP causes an additional delay on frame forwarding between the wireless and wired segments, which may exceed 20 ms (see Section V-B). The maximum latency recorded for single channels over all experiments was 814.861 ms, which significantly decreases to 27.041 ms for the redundant link.

Instead, statistics on the link latency d' (right side of the table) do not depend in any way from APs. Hence, results are more accurate when evaluating the redundant link alone. Surprisingly, the worst-case latency d'_{\max} occasionally exceeded d_{\max} in the same experiment. The explanation is that, the originator takes timestamps on ACK frame arrivals, but sometimes only the ACK frame, and not the related data frame, is lost, which means that a later ACK frame will be used for calculating d' , thus making it possibly larger than d .

2) *Multicast transmission*: Experimental conditions are the same as for unicast packets. Results, reported in Table IV, show that the number of multicast packets lost on physical channels is much larger than in the unicast case. In fact, they are unconfirmed, and hence the retransmission mechanism cannot be exploited by the MAC layer to improve reliability. While improvements provided by seamless redundancy on latency are not as good as for unicast transmissions, a dramatic reduction can be seen on the number of losses. For instance, in the *High Load–Day* condition, only 1.658 % of the packets were lost on the redundant link, while losses on individual channels amounted to 9.398 % and 17.940 %, respectively.

Remarkably, the worst-case latency on the redundant link may sometimes coincide with physical channels. For instance, in the *High Load–Night* condition, it is the same as channel 1, that is 133.312 ms. This situation was explicitly checked, and corresponds to a packet that suffered from the maximum delay on channel 1 and went lost on channel 165.

B. Assessment of Independence between Paths

To perform such an assessment we exploited the property that, when two events (either delays or losses) are statistically independent, their joint probability equals the product of individual probabilities.

1) *Unicast transmission*: In order to evaluate the effectiveness of guidelines to reduce dependence between channels, CCDFs can be considered [13]. If paths are independent, the CCDF of latencies measured on the redundant link can be theoretically estimated from those of individual channels (CCDF^{c1} and CCDF^{c165}) by means of Eq. (7) in [13], i.e., $\widehat{\text{CCDF}}^{c1 \cup c165} = \frac{1}{1 - \text{PLR}^{c1} \cdot \text{PLR}^{c165}} \cdot [\text{PLR}^{c1} \cdot (1 - \text{PLR}^{c165}) \cdot \text{CCDF}^{c165} + \text{PLR}^{c165} \cdot (1 - \text{PLR}^{c1}) \cdot \text{CCDF}^{c1} + (1 - \text{PLR}^{c165}) \cdot (1 - \text{PLR}^{c1}) \cdot \text{CCDF}^{c1} \cdot \text{CCDF}^{c165}]$, where $\widehat{\text{CCDF}}$ is the *estimated* CCDF.

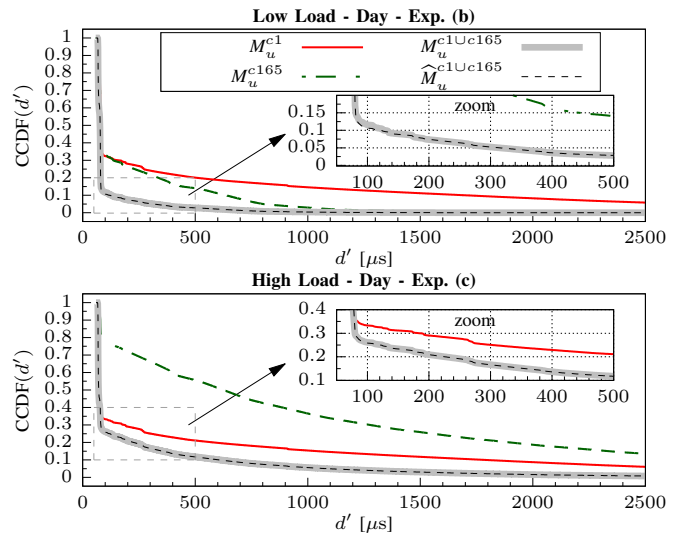


Fig. 11. CCDFs for experiments (b) and (c) in Table III (latency derived from ACK frame arrival).

As an example, for two specific experiments, identified in Table III with labels (b) and (c), the measured CCDF of link latencies d' on the redundant link is reported in Fig. 11, along with its theoretical estimate. Since they practically coincide, paths can be considered independent. In other words, the communication quality offered in practice by seamless redundancy is the same as what can be expected from theory, and can be modeled with excellent approximation using statistics acquired on individual Wi-Fi channels.

2) *Multicast transmission*: Quantity $\widehat{\text{PLR}}$ in Table IV is the *estimated* packet loss ratio, which in case channels are statistically independent can be obtained by multiplying their measured PLRs, i.e., $\widehat{\text{PLR}}^{c1 \cup c165} = \text{PLR}^{c1} \cdot \text{PLR}^{c165}$ [13]. As can be noted from results, the assumption of independence between paths reasonably holds for all the considered conditions. In fact, estimated and measured PLR values were always similar, the only exception being the *No Load–Night* condition, which is statistically unreliable because very few packets went lost. Similarity between estimated and measured results, not reported here for space reasons, also holds for CCDFs of latencies for multicast packets.

VIII. CONCLUSIONS

Seamless redundancy according to PRP can be adopted to tangibly improve communication quality in Wi-Fi networks. However, to obtain the performance boost one may expect from theory, care has to be taken so that concurrent transmissions on physical paths of the same redundant link do not suffer from either joint or mutual interference.

In this paper, several phenomena have been considered that, if not properly addressed, may impair effectiveness of seamless redundancy. They depend on aspects such as management activities performed in network equipment and end nodes, which may cause joint interference, antenna placement, which could lead to mutual interference on air between adjacent

TABLE IV
 RESULTS AFTER OPTIMIZATION (MULTICAST TRANSMISSIONS).

Exp.	Stream	Statistics derived from the PC Ethernet port					PLR [%]	PLR	
		\bar{d}	σ_d	$d_{p99.9}$ [ms]	$d_{p99.99}$	d_{max}			
No Load	Night	M_m^{c1}	1.382	1.553	19.770	27.163	47.791	6.131	0.000017
		M_m^{c165}	0.862	1.162	18.932	20.749	86.359	0.000278	
		$M_m^{c1 \cup c165}$	0.772	0.328	3.105	17.181	86.359	0.000278	
	Day	M_m^{c1}	1.873	3.816	30.004	84.203	796.201	9.731	
		M_m^{c165}	0.862	1.163	18.920	20.723	54.372	0.0	
		$M_m^{c1 \cup c165}$	0.778	0.419	5.941	19.119	54.372	0.0	
Low Load	Night	M_m^{c1}	1.570	5.513	39.353	147.394	1145.965	5.574	0.305
		M_m^{c165}	0.982	1.209	19.189	21.321	22.548	5.477	
		$M_m^{c1 \cup c165}$	0.866	1.791	10.713	28.365	947.181	0.288	
	Day	M_m^{c1}	1.594	2.280	20.988	60.583	329.424	7.748	
		M_m^{c165}	0.981	1.216	19.351	21.247	22.912	5.469	
		$M_m^{c1 \cup c165}$	0.874	0.738	9.737	20.196	230.299	0.413	
High Load	Night	M_m^{c1}	1.365	1.609	19.887	29.653	133.312	6.483	1.252
		M_m^{c165}	1.622	1.494	20.311	22.592	62.383	19.318	
		$M_m^{c1 \cup c165}$	1.073	0.890	13.810	21.210	133.312	1.252	
	Day	M_m^{c1}	2.061	6.501	46.164	301.973	896.643	9.398	
		M_m^{c165}	1.630	1.500	20.528	22.774	26.773	17.940	
		$M_m^{c1 \cup c165}$	1.260	2.424	20.011	58.843	579.104	1.658	

channels, and mechanisms like DTIM, which indirectly affect redundancy. To practically analyze the effects of such phenomena on communication quality in redundant wireless network implementations, a thorough experimental campaign was carried out in a real setup, consisting of conventional Linux PCs and COTS Wi-Fi equipment.

Experiments enabled us to define suitable countermeasures. From a practical viewpoint, the following guidelines should be followed when designing a networked control system communicating on redundant Wi-Fi links:

- 1) Mechanisms to support power saving in IEEE 802.11 battery-powered devices (e.g., DTIM) may severely interfere with timely delivery of multicast packets. To prevent this behavior, devices not belonging to the control system, like mobiles, should not be allowed to accidentally associate to the AP. This should be preferably obtained by enabling authentication and, possibly, encryption (e.g., WPA2). Although relevant, this aspect does not specifically concern redundant communications.
- 2) Software modules typically exist in end-nodes, which may jointly affect transmissions on all paths of the redundant link. As far as possible, this has to be prevented. A relevant example is the network manager (NM), which is used to cope with device (re)association to the AP. Disabling NM is usually not tolerable, as reliability would be severely compromised. A viable solution is to modify the NM code (or to rewrite it from scratch), so that channel scanning is never carried out concurrently on redundant adapters.
- 3) Joint interference caused by internal processing in network equipment (e.g., the AP) has to be prevented as well. Unfortunately, bringing modifications to the AP firmware is typically a quite complex task (and, sometimes, not possible at all). In the case of commercial APs, interference can typically be mitigated by using a pair of distinct devices in place of a single simultaneous dual-band AP.
- 4) RF modules may cause mutual interference between channels, when antennas are located too close to each other.

In such conditions, adjacent channel interference (ACI) may occur, which prevents carrier sensing from operating correctly and affects communication on both paths at the same time. To prevent ACI, proper frequency selection is needed. While channels should be preferably chosen in distinct frequency bands (2.4 GHz and 5 GHz), widely-spaced channels in the 5 GHz band may also be used.

Experimental evidences we collected confirm that the above guidelines make Wi-Fi channels statistically independent, to the point that communication quality on the redundant link matches the theoretical figures with very good approximation. Results, obtained from real testbeds configured in this way, fully support this claim. Of course, several other minor aspects can be also considered, concerning both the protocol and its implementation. They will be the subject of our future work.

REFERENCES

- [1] "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 802.11-2016 (Rev. of IEEE Std 802.11-2012)*, pp. 1–3534, Dec 2016.
- [2] A. Willig, "Recent and emerging topics in wireless industrial communications: A selection," *IEEE Trans. Ind. Informat.*, vol. 4, no. 2, pp. 102–124, May 2008.
- [3] S. Petersen and S. Carlsen, "WirelessHART Versus ISA100.11a: The Format War Hits the Factory Floor," *IEEE Industrial Electronics Magazine*, vol. 5, no. 4, pp. 23–34, Dec 2011.
- [4] R. Moraes, F. Vasques, P. Portugal, and J. Fonseca, "VTP-CSMA: A Virtual Token Passing Approach for Real-Time Communication in IEEE 802.11 Wireless Networks," *IEEE Trans. Ind. Informat.*, vol. 3, no. 3, pp. 215–224, 2007.
- [5] H. Trsek and J. Jasperneite, "An isochronous medium access for real-time wireless communications in industrial automation systems - A use case for wireless clock synchronization," in *IEEE Int. Symp. on Precision Clock Synchronization for Measurement Control and Communication (ISPCS)*, Sept 2011, pp. 81–86.
- [6] G. Gamba, F. Tamarin, and A. Willig, "Retransmission Strategies for Cyclic Polling Over Wireless Channels in the Presence of Interference," *IEEE Trans. Ind. Informat.*, vol. 6, no. 3, pp. 405–415, Aug 2010.
- [7] L. Seno, G. Cena, S. Scanzio, A. Valenzano, and C. Zunino, "Enhancing Communication Determinism in Wi-Fi Networks for Soft Real-Time Industrial Applications," *IEEE Trans. Ind. Informat.*, vol. 13, no. 2, pp. 866–876, April 2017.
- [8] G. Tian, S. Camtepe, and Y. C. Tian, "A Deadline-Constrained 802.11 MAC Protocol With QoS Differentiation for Soft Real-Time Control," *IEEE Trans. Ind. Informat.*, vol. 12, no. 2, pp. 544–554, April 2016.
- [9] G. Cena, L. Seno, A. Valenzano, and C. Zunino, "On the Performance of IEEE 802.11e Wireless Infrastructures for Soft-Real-Time Industrial Applications," *IEEE Trans. Ind. Informat.*, vol. 6, no. 3, pp. 425–437, 2010.
- [10] G. Cena, A. Valenzano, and S. Vitturi, "Hybrid wired/wireless networks for real-time communications," *IEEE Industrial Electronics Magazine*, vol. 2, no. 1, pp. 8–20, March 2008.
- [11] T. D. Chung, R. B. Ibrahim, V. S. Asirvadam, N. B. Saad, and S. M. Hassan, "Adopting EWMA Filter on a Fast Sampling Wired Link Contention in WirelessHART Control System," *IEEE Trans. Instrum. Meas.*, vol. 65, no. 4, pp. 836–845, April 2016.
- [12] M. Mohiuddin, M. Popovic, A. Giannakopoulos, and J. Y. L. Boudec, "Experimental validation of the usability of Wi-Fi over redundant paths for streaming phasor data," in *Int. Conference on Smart Grid Communications (SmartGridComm)*, Nov 2016, pp. 533–538.
- [13] G. Cena, S. Scanzio, and A. Valenzano, "Experimental Evaluation of Seamless Redundancy Applied to Industrial Wi-Fi Networks," *IEEE Trans. Ind. Informat.*, vol. 13, no. 2, pp. 856–865, April 2017.
- [14] —, "Design guidelines to improve reliability of seamless redundancy in Wi-Fi networks," in *IEEE Int. Conf. on Emerging Technologies and Factory Automation (ETFA)*, Sept 2016, pp. 1–8.

- [15] IEC, "Industrial communication networks - High availability automation networks - Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)," *IEC Std 62439-3 ed 2.0*, pp. 1–177, 2010.
- [16] P. Castello, P. Ferrari, A. Flammini, C. Muscas, P. A. Pegoraro, and S. Rinaldi, "A Distributed PMU for Electrical Substations With Wireless Redundant Process Bus," *IEEE Trans. Instrum. Meas.*, vol. 64, no. 5, pp. 1149–1157, May 2015.
- [17] M. Popovic, M. Mohiuddin, D. C. Tomozei, and J. Y. L. Boudec, "iPRP—The Parallel Redundancy Protocol for IP Networks: Protocol Design and Operation," *IEEE Trans. Ind. Informat.*, vol. 12, no. 5, pp. 1842–1854, Oct 2016.
- [18] A. Mahmood, R. Exel, H. Trsek, and T. Sauter, "Clock Synchronization Over IEEE 802.11; A Survey of Methodologies and Protocols," *IEEE Trans. Ind. Informat.*, vol. 13, no. 2, pp. 907–922, April 2017.
- [19] M. M. Awad, H. H. Halawa, M. Rentschler, R. M. Daoud, and H. H. Amer, "Novel system architecture for railway wireless communications," in *IEEE Int. Conference on Smart Technologies (EUROCON)*, July 2017, pp. 877–882.
- [20] D. K. Dang, A. Mifdaoui, and T. Gayraud, "Design and analysis of UWB-based network for reliable and timely communications in safety-critical avionics," in *IEEE Workshop on Factory Communication Systems (WFCS)*, May 2014, pp. 1–10.
- [21] M. Rentschler and P. Laukemann, "Towards a reliable parallel redundant WLAN black channel," in *IEEE Int. Workshop on Factory Communication Systems (WFCS)*, 2012, pp. 255–264.
- [22] M. Rentschler, O. A. Mady, M. T. Kassis, H. H. Halawa, T. K. Refaat, R. M. Daoud, H. H. Amer, and H. M. ElSayed, "Simulation of parallel redundant WLAN with OPNET," in *IEEE Conference on Emerging Technologies Factory Automation (ETFA)*, Sept 2013, pp. 1–8.
- [23] G. Cena, S. Scanzio, and A. Valenzano, "Seamless Link-Level Redundancy to Improve Reliability of Industrial Wi-Fi Networks," *IEEE Trans. Ind. Informat.*, vol. 12, no. 2, pp. 608–620, April 2016.
- [24] G. Cena, S. Scanzio, A. Valenzano, and C. Zunino, "An enhanced MAC to increase reliability in redundant Wi-Fi networks," in *IEEE Int. Workshop on Factory Communication Systems (WFCS)*, May 2014, pp. 1–10.
- [25] —, "Dynamic duplicate deferral techniques for redundant Wi-Fi networks," in *IEEE Int. Conf. on Emerging Technology and Factory Automation (ETFA)*, Sept 2014, pp. 1–8.
- [26] G. Cena, S. Scanzio, and A. Valenzano, "Duplication avoidance mechanisms to reduce bandwidth usage in redundant Wi-Fi networks," in *IEEE Int. Workshop on Factory Communication Systems (WFCS)*, May 2017, pp. 1–10.
- [27] —, "A software-defined MAC architecture for Wi-Fi operating in user space on conventional PCs," in *IEEE Int. Workshop on Factory Communication Systems (WFCS)*, May 2017, pp. 1–10.
- [28] "Electromagnetic compatibility and Radio spectrum Matters (ERM); Wideband transmission systems; Data transmission equipment operating in the 2.4 GHz ISM band and using wide band modulation techniques; Harmonized EN covering the essential requirements of article 3.2 of the R&TTE Directive," *ETSI EN 300 328 V1.9.1 Std.*, pp. 1–91, Feb. 2015.
- [29] F. Tramarin, S. Vitturi, M. Luvisotto, and R. Parrozzani, "Performance assessment of an IEEE 802.11-based protocol for real-time communication in agriculture," in *IEEE Int. Conf. on Emerging Technology and Factory Automation (ETFA)*, Sept 2014, pp. 1–6.
- [30] J. Nachtigall, A. Zubow, and J. P. Redlich, "The Impact of Adjacent Channel Interference in Multi-Radio Systems using IEEE 802.11," in *Int. Wireless Communications and Mobile Computing Conference*, Aug 2008, pp. 874–881.
- [31] G. Cena, I. C. Bertolotti, A. Valenzano, and C. Zunino, "Evaluation of Response Times in Industrial WLANs," *IEEE Trans. Ind. Informat.*, vol. 3, no. 3, pp. 191–201, Aug 2007.



Gianluca Cena (SM'09) received the Laurea degree in electronic engineering and the Ph.D. degree in information and system engineering from the Politecnico di Torino, Italy, in 1991 and 1996, respectively.

Since 2005 he has been Director of Research with the Institute of Electronics, Computer and Telecommunication Engineering, National Research Council of Italy (CNR–IEIIT), Turin. His research interests include wired and wireless industrial communication systems, real-time protocols, and automotive networks. In these areas he has coauthored about 130 technical papers, three of which awarded as Best Papers of the 2004, 2010, and 2017 editions of the IEEE Workshops on Factory Communication Systems, and one international patent.

Dr. Cena served as a Program Co-Chairman for the 2006 and 2008 editions of the IEEE International Workshop on Factory Communication Systems, and as a Track Co-Chairman in six editions of the IEEE International Conference on Emerging Technologies and Factory Automation. Since 2009 he has been an Associate Editor of the IEEE Transactions on Industrial Informatics.



Stefano Scanzio (S'06-M'12) received the Laurea and Ph.D. degrees in Computer Science from Politecnico di Torino, Torino, Italy, in 2004 and 2008, respectively.

He was with the Department of Computer Engineering, Politecnico di Torino, from 2004 to 2009, where he was involved in research on speech recognition and, in particular, he has been active in classification methods and algorithms. Since 2009, he has been with the National Research Council of Italy (CNR), where he is a tenured Researcher with

the Institute of Electronics, Computer and Telecommunication Engineering (IEIIT), Turin.

Dr. Scanzio teaches several courses on Computer Science at Politecnico di Torino. He has authored and co-authored of more than 50 papers in international journals and conferences, most of them in the area of industrial communication systems, real-time networks, wireless networks and clock synchronization protocols. He received the Best Paper Awards for the papers he presented at the 8th and 13th IEEE Workshops on Factory Communication Systems (WFCS 2010 and WFCS 2017).



Adriano Valenzano (SM'09) received the Laurea degree in electronic engineering from Politecnico di Torino, Torino, Italy, in 1980.

He is Director of Research with the National Research Council of Italy (CNR). He is currently with the Institute of Electronics, Computer and Telecommunication Engineering (IEIIT), Torino, Italy, where he is responsible for research concerning distributed computer systems, local area networks, and communication protocols. He has coauthored approximately 200 refereed journal and conference papers in the area of computer engineering.

Dr. Valenzano is the recipient of the 2013 IEEE IES and ABB Lifetime Contribution to Factory Automation Award. He also received, as a coauthor, the Best Paper Award presented at the 5th, 8th and 13th IEEE Workshops on Factory Communication Systems (WFCS 2004, WFCS 2010 and WFCS 2017). He has served as a technical referee for several international journals and conferences, also taking part in the program committees of international events of primary importance. Since 2007, he has been serving as an Associate Editor for the IEEE Transactions on Industrial Informatics.