

Multi-Link Operation and Wireless Digital Twin to Support Enhanced Roaming in Next-Gen Wi-Fi

Stefano Scanzio*, Matteo Rosani*[†], Gabriele Formis*[†], Dave Cavalcanti[‡],
Valerio Frascolla[§], Guido Marchetto[†], Gianluca Cena*,

*National Research Council of Italy (CNR-IEIIT), Italy. [†]Politecnico di Torino, Italy.

[‡]Intel Labs, Intel Corporation, Hillsboro, OR, USA. [§]Intel Labs, Intel Deutschland, Neubiberg, Germany.

Email: name.surname@{cnr.it, intel.com, polito.it}

Abstract—The next generation of Wi-Fi is meant to achieve ultra-high reliability for wireless communication. Several approaches are available to this extent, some of which are being considered for inclusion in standards specifications, including coordination of access points to reduce interference.

In this paper, we propose a centralized architecture based on digital twins, called WiTwin, with the aim of supporting wireless stations in selecting the optimal association according to a set of parameters. Unlike prior works, we assume that Wi-Fi 7 features like multi-link operation (MLO) are available. Moreover, one of the main goals of this architecture is to preserve communication quality in the presence of mobility, by helping stations to perform reassociation at the right time and in the best way.

I. INTRODUCTION

Due to its high performance and low-cost, Wi-Fi is becoming one of the most pervasive communication technologies, which will enable new production paradigms in the factories of the future. The recent enhancements to IEEE 802.11 specifications make Wi-Fi suitable, in terms of reliability and communication latency, for many application contexts related to Industry 4.0 and the forthcoming Industry 5.0. Some of these improvements, especially those aimed at enhancing network throughput and multi-link operation (MLO), have been included in the IEEE 802.11be amendment (also known as Wi-Fi 7), while others are being defined by the IEEE 802.11bn Task Group for Wi-Fi 8 [1].

When mobility is considered, e.g., for Automated Guided Vehicles (AGV) and Autonomous Mobile Robots (AMR), which are increasingly used in modern factories, Wi-Fi does not perform as well as cellular technologies like 5G. Roaming still constitutes a critical operation in Wi-Fi because, when a wireless station (STA) reassociates to a different access point (AP), it experiences a short interval (up to a few seconds) where communication quality is either unsatisfactory (excessive latency and jitters) or prevented altogether. These gaps are annoying for soft real-time applications, but may lead to serious consequences in firm/hard real-time ones (control loops), where the lack of communication is not tolerated.

This work has been partially funded by SoBigData.it “SoBigData.it receives funding from European Union – NextGenerationEU – National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) – Project: “SoBigData.it – Strengthening the Italian RI for Social Mining and Big Data Analytics” – Prot. IR0000013 – Avviso n. 3264 del 28/12/2021.”, and partially funded by the European Commission Horizon Europe SNS JU PREDICT-6G (GA 101095890) Project.

The problem of channel selection in a dense deployment has been explored in [2] using reinforcement learning techniques. In particular, a Parallel Transfer Reinforcement Learning (PTRL) algorithm has been presented to improve convergence of MLO-enabled networks in the event of a channel selection. In [3] the channel allocation problem is explored in the Internet of Things context and a multi-armed-bandit-based channel allocation method is proposed to deliver higher frame delivery ratio (FDR) and faster response for AP channel selection. Lastly, a channel selection method that relies on the network load is presented in [4], which estimates the load of the network and uses this information to make intelligent decisions. Another critical point is the handover procedure because the interruption of service is a challenging problem in time-sensitive scenarios [5]. In [6] a new technique is proposed to reduce the handover delay to a bounded value.

In this work, a digital twin (DT) of the wireless communication environment is defined that permits moving STAs with MLO capability to perform seamless roaming between APs. In particular, the proposed DT exploits the capability of multi-link devices (MLD) to have multiple links active at the same time. This means that, during reassociation, links can be moved one at a time to the new AP, so that at any time there is always at least one link that is operating properly and supports communication with adequate quality for time-aware applications. The paper is structured as follows: Section II sketches an overview of the system architecture, while its main operations and components are analyzed in Section III, just before the conclusive remarks are drawn in Section IV.

II. SYSTEM ARCHITECTURE

The architecture of Fig. 1, which is based on a centralized approach, is being proposed to support reliable communication in presence of STA mobility, e.g., to interconnect fleets of AGVs and AMRs. In the proposed architecture, APs are connected to a wired infrastructure (typically an Ethernet network), which permits them to communicate with a centralized entity named *WiTwin*. We assume that the wired network behaves ideally in terms of frame losses. If fault-tolerance is demanded, solutions like parallel redundancy can be exploited.

Since we are witnessing the dawn of Wi-Fi 7, as specified by IEEE 802.11be [7], we assume that every AP is an MLD that supports MLO [8], [9], i.e., it may have two or

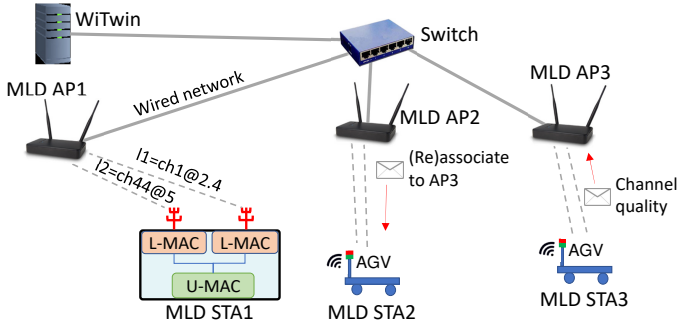


Fig. 1. Proposed architecture based on WiTwin.

more active associations to different APs and transmit (quasi) simultaneously on different links tuned on distinct channels. Conversely, STAs could be either MLDs or legacy Wi-Fi devices. In an MLD, the media access control (MAC) layer is split in two “sub-layers”: at the bottom L-MAC entities (two or more, each one related to a specific link) interface with the PHYs and manage aspects like channel access and retransmissions. On top the U-MAC (a single entity for every STA) coordinates all the local L-MACs and interface them with the data-link user. Every STA cyclically sends the perceived *channel quality*, as well as its position, to the WiTwin. Starting from the information acquired from STAs, the WiTwin builds (and then keeps updated) a spatial model of the radio environment, described in terms of the communication quality for every channel. Mobile STAs can exploit this model, named *WiTwinModel*, to optimize roaming between APs. However, the same model can be also exploited by static STAs to balance the load among APs and improve the overall communication quality on air.

Roaming in legacy Wi-Fi is managed by every single STA separately, and decisions about reassociation rely only on locally observable quantities. Examples are receive signal strength indicator (RSSI) and FDR. For the *network-driven roaming* we propose, the WiTwin notifies the STAs the information about the new (best) AP to reassociate using specific messages. In this case, all the available information, as described by the *WiTwinModel* (including the current position and mobility pattern of the STAs), can be exploited to make optimal decisions. Among the metrics that can be used to select the new AP, we deem that latency could be one of the most appropriate in the context of mobility in industrial environments. In fact, for this class of applications, ensuring bounded transmission delays (even from a probabilistic point of view) is essential for the correct operation of machinery and equipment. Reassociation can be triggered either directly by the WiTwin (*proactive* behavior) or by the STA when the observed communication quality on the link suddenly falls below a given threshold (*reactive* behavior). In addition to network-driven roaming, legacy *STA-driven roaming* is still needed to permit STAs to reassociate also when communication with the AP is prevented. In this case, the STA autonomously decides the best AP, possibly exploiting prior information that was

made available by the WiTwin and cached in the STA, or by using machine learning (ML) algorithms aimed at predicting channel quality in the near future [10].

Concerning the above architecture, an important aspect is the *reassociation procedure*. Specifically, reassociation of MLD STAs should be done one L-MAC at a time to keep latency bounded and preserve reliable communication with nodes connected to the wired network, e.g., a programmable logic controller (PLC) [11].

III. MAIN OPERATIONS AND COMPONENTS

Our proposal is characterized by four main aspects that are analyzed independently: feature acquisition, *WiTwinModel*, roaming, and MLD reassociation.

A. Feature acquisition

This aspect encompasses the acquisition and transfer of information from STAs, useful to build the *WiTwinModel*. The i -th sample obtained on transmissions from the source STA s to the destination AP d (or vice-versa) on channel c , can be described as a tuple $x_i^{(s,d,c)} = \langle \nu_1, \nu_2, \dots, \nu_N \rangle$, where N is the dimension of the feature space, which depends on the measuring device’s capabilities (i.e., on the ability of the node to measure certain physical quantities). For example, an MLD STA with three L-MACs operating on channels in the 2.4 GHz, 5 GHz, and 6 GHz bands will acquire features related to any of its three channels, and information on frame transmissions between STA 3 and AP 5 (unique identifiers are needed, e.g., MAC addresses) on channel 1 in the 2.4 GHz band will be denoted $x_i^{(3,5,ch1@2.4)}$. For instance, in the sample $x_i^{(s,d,c)} = \langle t, x, y, \nu_4, \dots, \nu_N \rangle$, t is a timestamp taken using a common (network-wide) time base, x and y (plus, possibly, z) represent the position of the device, while ν_4, \dots, ν_N are other features related to channel quality (e.g., FDR, RSSI, number of transmission attempts for a packet). For APs, the position needs not be refreshed often because they are static. In the initial version of the architecture we suppose that nodes are localized by means of either the IEEE 802.11az protocol or other less precise techniques [12].

Delivering sample $x_i^{(s,d,c)}$ from the STA to the WiTwin can be done by means of specifically-defined messages (that may include one or more samples, possibly compressed) or by using messages defined in IEEE 802.11k (after the required adjustments). Moreover, the AP may also infer some quantities by analyzing message exchanges on air (e.g., from the reception of beacon, data, and ACK frames). APs exploit the wired network to deliver information to the WiTwin.

B. *WiTwinModel*

The environment is modeled by WiTwin using the *WiTwinModel*. While exact details about it must still be finalized, it is based on a heatmap that, in its most straightforward implementation, is defined on a two-dimensional space (3D extensions can be easily obtained). For every point in the plane the *WiTwinModel* provides the estimated transmission channel quality \hat{h} (or, possibly, its forecast in the immediate




Fig. 2. Example of network-driven roaming.

future) to each MLD AP in the area of interest. Depending on the roaming management algorithm (introduced in the next section), the estimate \hat{h} can refer to different quantities. At present, RSSI is widely used in Wi-Fi STAs to decide whether and where to roam [13]. In this work, we propose instead to use the estimated FDR, as it is directly related to the probability to succeed in sending a frame to destination, and consequently to latency.

The WiTwinModel is constantly (re)trained using the features acquired at runtime, and implements the function

$$\hat{h}(x, y, id_{AP}, c), \quad (1)$$

that, given the STA position (x, y or x, y, z), a tag id_{AP} that unambiguously identifies an AP, and a specific channel c , returns an estimate of the link quality. A possible extension of \hat{h} is to add the future time t_f at which to perform estimation, obtaining a slightly different prototype $\hat{h}(x, y, id_{AP}, c, t_f)$. Doing so permits to better handle mobility because the time when reassociation is done can be selected in advance based on spectrum dynamics and the mobility characteristics of the node (direction and speed), but it also makes the implementation of the WiTwinModel noticeably more complex.

We wish to point out that, in this architecture, MLD APs are not enabled to change at runtime their configuration (in terms of the used channels). In fact, this operation is quite complex in practice, as all nodes already associated with an AP would need to be migrated to other channels, which unavoidably impacts on the real-time characteristics of traffic.

As an example, for an MLD AP with three configured links (e.g., channel 5 in 2.4 GHz band, channel 44 in 5 GHz band, and channel 73 in 6 GHz band), the WiTwinModel may return for a given point in the plane (or in the space) a tuple of three elements that model the estimated channel quality related to a future position of a node. In practice, this information can be obtained by invoking function \hat{h} thrice, and results can be used by the WiTwin to compute and send information to the STA about the AP to which it must reassociate, plus additional information about the optimal roaming sequence. These aspects are discussed in the next section. Concerning the previous example, the tuple might be as follows

$$\langle \hat{h}(\cdot, ch5@2.4), \hat{h}(\cdot, ch44@5), \hat{h}(\cdot, ch73@6) \rangle, \quad (2)$$

where the first three arguments of $\hat{h}(x, y, id_{AP})$ have been replaced with a dot because they are the same in all the calls.

Estimation carried out by \hat{h} can be implemented using either statistical methods or ML. Both methods are trained on the set of feature $\{x_i^{(s,d,c)}\}$, where more weight should be clearly given to the most recent samples (the latter is typically handled automatically by ML algorithms). Concerning the use of ML for predicting the future channel quality, a good survey is presented in [14], including, e.g., the use of combinations of exponential moving averages [15] and artificial neural networks [10].

C. Roaming

In this paper we only focus on network-driven roaming, even if the services provided by the WiTwin could greatly improve communication latency also for STA-driven roaming. In network-driven roaming, the WiTwin sends a specific message to the STA to enforce the reassociation procedure when the node enters (or is expected to) an area where another AP is available that is characterized by a higher FDR (switching to it may potentially improve communication latency too). In this context, the WiTwin uses the WiTwinModel to determine in real-time when the channel quality of the new candidate AP overtakes the current AP. The new AP $id_{AP'}$ is selected as follows:

$$id_{AP'} = \arg \min_{a \in \mathcal{A}} \sum_{c \in \mathcal{C}^a} (1 - \hat{h}(x, y, a, c)), \quad (3)$$

where \mathcal{A} is the set of all the APs on which search is performed while \mathcal{C}^a is the set of all the channels configured on AP a . The message sent by the WiTwin to the STA to enforce reassociation is generated only when $id_{AP'}$ differs from the current AP and improvements, in terms of the aggregate FDR (which considers all the links of the MLD AP) exceed a given threshold, to prevent oscillations and network instability.

This message is delivered to the STA either through the AP it is associated to or, possibly, via the beacons of neighboring APs operating on the same channel, and specifies the target AP it must/should reassociate and what is the migration order of the affiliated L-MACs. With reference to the example of Fig. 2, when the WiTwin notices that the STA is going to enter into an area where a better communication quality (i.e., higher FDR) is available, AP1 sends a message instructing the STA to reassociate to AP2.

D. MLD reassociation

Likely, managing the reassociation procedure of MLDs is one of the most interesting points discussed in this work. To ensure bounded (and minimal) latency to the time-sensitive applications that involve a moving MLD STA, e.g., a PLC connected to the wired network controlling an AMR over the air, reassociation is assumed to always involve a single L-MAC at a time. During reassociation, for a certain time interval the MLD STA is associated with two distinct MLD APs (the older and the newer ones), with the U-MAC selecting frame-by-frame the best option. In this case, packets could be possibly sent on paths involving both APs. By doing so, connection is virtually never interrupted. The possibility of

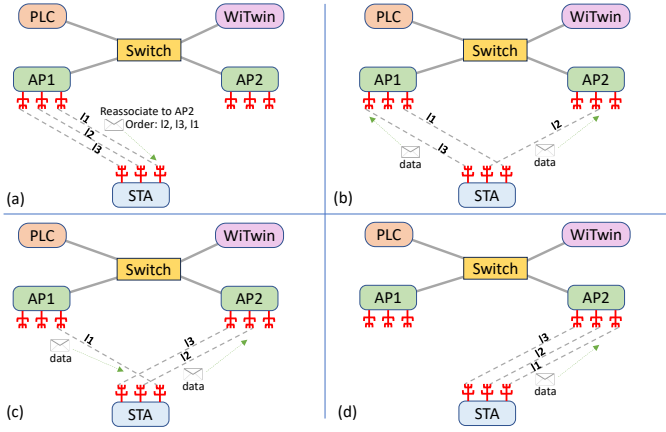


Fig. 3. Sequence of steps performed during MLD roaming.

sending two copies of the same packet on multiple links at the same time permits to drastically reducing latency during roaming operations in time-critical Wi-Fi applications [16]. However, it requires a more detailed analysis to determine which network entity is in charge of removing duplicates.

The message sent by WiTwin to the MLD STA that triggers reassociation also specifies the migration order. This sequence is selected to probabilistically maximize the combined link quality in terms of FDR. The MLD roaming process is shown in Fig. 3. Let us assume (Fig. 3.a) that the moving STA is associated to AP1 with three links (l_1, l_2, l_3) to transmit/receive data to/from the PLC located in the wired network. When the WiTwin detects that, according to the law of motion, AP2 has (likely) become able to provide better communication quality than AP1, it sends a message to the STA that contains the indication of this new AP and the order in which links must be moved from the current to the new AP (e.g., l_2, l_3, l_1).

At this point the STA starts moving link l_2 to the new AP2 (Fig. 3.b), but in the meanwhile, links l_1 and l_3 can be still used for communicating with the PLC. Then, it is the turn of l_3 (Fig. 3.c), and links l_2 and l_1 can be used for communication. In this case, part of the messages flow through AP1 and part through AP2. Finally, also the link l_1 is reassociated with AP2 (Fig. 3.d), and links l_2 and l_3 can be used in the meanwhile, after which the MLD roaming procedure is completed and the STA can use all the three links now associated with AP2.

IV. CONCLUSIONS

Wi-Fi is currently the most popular wireless communication technology that enables wireless extensions to Ethernet network infrastructures. One of the most effective approaches to increase its reliability is to coordinate APs operations to reduce the interference among them. To this aim several proposals, which select associations between STAs and APs to maximize communication quality, have recently appeared.

In this paper we consider two additional aspects that are expected to gain attention in the close future. First, we rely on Wi-Fi 7, whose specifications are going to be published soon and for which commercial equipment has started to be

available. Second, we focus on mobility of STAs, and strive to optimize roaming (i.e., the handover procedure) by improving the way reassociation is carried out.

In particular, we exploit a DT for the radio environment that finds: a) the best time to perform reassociation, before communication quality starts to worsen sensibly; b) the best option for the next AP to which reassociate; and, c) in which order the different links of an MLD must be moved to the new AP. To this purpose we are working on the definition of a suitable architecture, which we present here in its draft form.

REFERENCES

- [1] L. G. Giordano, G. Geraci, M. Carrascosa, and B. Bellalta, "What Will Wi-Fi 8 Be? A Primer on IEEE 802.11bn Ultra High Reliability," *arXiv*, vol. 2303.10442, 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2303.10442>
- [2] P. E. Iturria-Rivera, M. Chenier, B. Herscovici, B. Kantarci, and M. Erol-Kantarci, "Channel Selection for Wi-Fi 7 Multi-Link Operation via Optimistic-Weighted VDN and Parallel Transfer Reinforcement Learning," in *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2023, pp. 1–6.
- [3] S. Hasegawa, R. Kitagawa, A. Li, S.-J. Kim, Y. Watanabe, Y. Shoji, and M. Hasegawa, "Multi-Armed-Bandit Based Channel Selection Algorithm for Massive Heterogeneous Internet of Things Networks," *Applied Sciences*, vol. 12, no. 15, 2022. [Online]. Available: <https://www.mdpi.com/2076-3417/12/15/7424>
- [4] D. Saliba, R. Imad, and S. Houcke, "Wifi channel selection based on load criteria," in *2017 20th International Symposium on Wireless Personal Multimedia Communications (WPMC)*, 2017, pp. 332–336.
- [5] L. Di Gregorio and V. Frascolla, "Handover optimality in heterogeneous networks," in *2019 IEEE 5G World Forum (5GWF)*, 2019, pp. 365–370.
- [6] P. Avila-Campos, J. Haxhibeqiri, X. Jiao, B. Van Herbruggen, I. Mörnerman, and J. Hoebeke, "Unlocking mobility for wi-fi-based wireless time-sensitive networks," *IEEE Access*, pp. 1–1, 2024.
- [7] C. Deng, X. Fang, X. Han, X. Wang, L. Yan, R. He, Y. Long, and Y. Guo, "IEEE 802.11be Wi-Fi 7: New Challenges and Opportunities," *IEEE Commun. Surv. Tutor.*, vol. 22, no. 4, pp. 2136–2166, 2020.
- [8] M. Carrascosa-Zamacois, G. Geraci, E. Knightly, and B. Bellalta, "Wi-Fi Multi-Link Operation: An Experimental Study of Latency and Throughput," *IEEE/ACM Trans. Netw.*, vol. 32, no. 1, pp. 308–322, 2024.
- [9] B. Bellalta, M. Carrascosa, L. Galati-Giordano, and G. Geraci, "Delay Analysis of IEEE 802.11be Multi-Link Operation Under Finite Load," *IEEE Wireless Commun. Letters*, vol. 12, no. 4, pp. 595–599, 2023.
- [10] S. Scanzio, F. Xia, G. Cena, and A. Valenzano, "Predicting Wi-Fi link quality through artificial neural networks," *Internet Technology Letters*, vol. 5, no. 2, p. e326, 2022. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/itl2.326>
- [11] S. Sudhakaran, I. Ali, M. Eisen, J. Perez-Ramirez, C. Cazan, V. Frascolla, and D. Cavalcanti, "Zero-Delay Roaming for Mobile Robots Enabled by Wireless TSN Redundancy," in *2023 IEEE 19th International Conference on Factory Communication Systems (WFCS)*, 2023, pp. 1–8.
- [12] P. Roy and C. Chowdhury, "A survey on ubiquitous WiFi-based indoor localization system for smartphone users from implementation perspectives," *CCF Transactions on Pervasive Computing and Interaction*, vol. 4, no. 3, pp. 298–318, Sep 2022.
- [13] S. Feirer and T. Sauter, "Seamless handover in industrial WLAN using IEEE 802.11k," in *2017 IEEE 26th International Symposium on Industrial Electronics (ISIE)*, 2017, pp. 1234–1239.
- [14] S. Szott, K. Kosek-Szott, P. Gawlowicz, J. T. Gómez, B. Bellalta, A. Zubow, and F. Dressler, "Wi-Fi Meets ML: A Survey on Improving IEEE 802.11 Performance With Machine Learning," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 3, pp. 1843–1893, 2022.
- [15] G. Formis, S. Scanzio, G. Cena, and A. Valenzano, "Linear Combination of Exponential Moving Averages for Wireless Channel Prediction," in *2023 IEEE 21st International Conference on Industrial Informatics (INDIN)*, 2023, pp. 1–6.
- [16] G. Cena, S. Scanzio, D. Cavalcanti, and V. Frascolla, "Seamless Redundancy for High Reliability Wi-Fi," in *2023 IEEE 19th International Conference on Factory Communication Systems (WFCS)*, 2023, pp. 1–4.