Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility

Fink, Andreas; Mogensen, Rasmus Suhr; Rodriguez, Ignacio; Kolding, Troels E.; Karstensen, Anders; Pocovi, Guillermo

Published in:
European Wireless 2021; 26th European Wireless Conference

Publication date:
2021

Document Version
Early version, also known as pre-print

Citation for published version (APA):
Fink, A., Mogensen, R. S., Rodriguez, I., Kolding, T. E., Karstensen, A., & Pocovi, G. (2021). Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility. In *European Wireless 2021; 26th European Wireless Conference* (pp. 125-132). VDE Verlag GMBH. https://ieeexplore.ieee.org/document/9657101

Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility

Andreas Fink*, Rasmus Suhr Mogensen*, Ignacio Rodriguez*, Troels Kolding†,
Anders Karstensen*, Guillermo Pocovi†
*Wireless Communication Networks Section, Department of Electronic Systems, Aalborg University, Denmark
†Nokia Bell Labs, Aalborg, Denmark

***Abstract*—This paper presents empirical latency measurements of enterprise-grade Wi-Fi 6 in an industrial setting with focus on handover performance. The basic mechanisms of Wi-Fi handover are evaluated along with improvements from several IEEE 802.11 amendments. Measurements are done for both idle and loaded networks using either dedicated frequency channels or frequency re-use. The benefits of using IEEE 802.11r and optimising scanning parameters are determined. It was found that optimising channel-related scanning parameters significantly reduces latency at the 99.9%-ile, whereas IEEE 802.11r shows improvements to a lesser degree on loaded networks. The observed latency values exceed the typical requirements assumed for IIoT use cases.**

I. INTRODUCTION

Industrial Internet of Things (IIoT) will be a major aspect in the push for Industry 4.0, where both adaptability and improved machine-to-machine communication play critical roles [1]. As the demand for smarter and more advanced IIoT systems increases, so do the requirements for the wireless communications of the devices, where some need low-latency communication while others require higher throughput. One example of IIoT with strict latency requirements is the use of Autonomous Mobile Robots (AMR), where on-demand transport of resources and supplies can be applied to a highly dynamic and configurable environment [2], but may require constant communication with the network for rapid decision-making in a flexible environment. In order to support these use cases, current wireless technologies need to be enhanced with these properties taken into consideration. One example is 5G NR, where a significant effort is being made to provide support for industrial time-sensitive networking and ultra-reliable low latency communication (URLLC). IEEE 802.11 Wi-Fi is another case of this, with its latest iteration, Wi-Fi 6, aiming to bring improved performance such as increased throughput using multi-user MIMO and dedicated resource allocation through orthogonal frequency-division multiple access (OFDMA).

Mobility is a key concern when addressing IIoT use cases in order to ensure full flexibility of industrial devices, and to support it, Wi-Fi and cellular technologies each have different approaches. While in 5G NR, the handovers are managed by the network; Wi-Fi relies on the Station (STA) itself to determine when and how to handle roaming events when leaving the service area of a particular Access Point (AP). This is typically done in a non-seamless manner, where the STA will dissociate with its current AP to connect to a new nearby AP, causing a brief period without any connectivity. The duration of this period can, in worst-case scenarios, cause some applications on the network to fail, as they may expect the device to be reachable. This will depend on a multitude of factors, such as the environment in general, the overall coverage, the line-of-sight conditions between STA and AP, and the interference from other Wi-Fi sources. Under these considerations, we intend to investigate the performance of enterprise Wi-Fi deployments in industrial settings.

Different aspects of the handover performance in Wi-Fi have been previously investigated in related work. In [3], various causes for the data interruption time were identified, along with parameters for handover decisions, such as the received signal strength and latency. In [4], an experimental evaluation of the impact of IEEE 802.11r Fast BSS Transition (FT) found that significant improvements in terms of minimising the handover interruption time could be obtained for Wi-Fi networks utilising IEEE 802.1X authentication, since the communication with an authentication server can be avoided when roaming between APs connected to the same network. Improved handover performance using the IEEE 802.11k amendment was investigated by [5], in which experimental results revealed a significant decrease in the duration of handovers by minimising the scanning time. A solution to optimising the choice of AP for which a STA should connect to has been proposed in [6], where the direction of a mobile node is used for this decision making, with simulations showing a clear reduction in the handover latency.

It is clear that the mobility performance in Wi-Fi have been studied based on empirical analysis. However, there is a lack in terms of empirical studies into this topic, with regards to reliable latency (i.e., the latency that is achieved with a certain probability, e.g. 99.9%) in the communication. Thus, this paper presents an experimental analysis of Wi-Fi performance based on the latest commercial iteration (Wi-Fi 6) in an industrial setting, with focus on latency and reliability. Furthermore, the impact of IEEE 802.11 handover-specific amendments aiming at seamless roaming are also investigated.

The remaining of this paper is structured as follows: Section II presents an overview of Wi-Fi with focus on mobility, including approaches to improve its performance. Section III details aspects related to tests of the performance with different network setups. Section IV presents the results of real-world measurements and identifies areas with potential for improvements. Finally, Section V concludes the paper.

II. HANDOVER IMPROVEMENTS IN WI-FI

Using Wi-Fi in infrastructure mode requires the STA to be connected to an AP in order to communicate. Because the connection between a STA and an AP may be degraded due to effects such as interference, scattering from nearby moving objects or because the STA itself is moving away from the AP, establishing a new connection to a different AP may be necessary. This is known as a handover and, as depicted in Fig. 1, can be divided into 4 main stages: scanning, authentication, association, and handshake. During scanning, the device will search for available APs to connect to using active or passive probing. This can, depending on the number of channels, take a significant amount of time as e.g. active probing requires waiting for responses to a probe from any AP on the frequency channel. The choice of which AP to connect to is made by the end device itself and can be customised to prioritise certain APs, but will, in most cases, only be based on which available AP has the highest received signal strength. The authentication stage is used to initiate the connection between the device and the selected AP. During the association stage, the device is registered to the AP directly. Device/AP capabilities are also exchanged at this point. Finally, a handshaking process is used to agree upon a selected encryption method such as WPA2-PSK. After this, data transfer between the device and AP can begin.

The device can be disassociated with its current serving AP through a number of means, but when considering mobility, two main cases can be expected: roaming due to a low Received Signal Strength Indicator (RSSI) value or by detecting that the serving AP is out of range (disconnection). If low RSSI is detected, as illustrated in Fig. 1a, the scanning phase is initiated, but with small interruptions to allow for data transfer. This is possible as the device is still associated with its current AP, so while this may extend the scanning period slightly and introduce increased latency for the communication, it is not a complete disconnect from the network. If the device however detects that it cannot reach the serving AP, it may initiate a short scan on channels at which it had previously found suitable APs to minimise the handover impact, as shown in Fig. 1b. If no APs are found during this short scan, a regular scan is initiated to scan other channels. To avoid reconnecting to the previous AP upon the first search, this one is blacklisted until another connection is established. The device will not be able to transfer data like in the previous case, as it is fully disconnected at this point. Even though an RSSI threshold is used to avoid this, these disconnections can still occur, such as when a large amount of interference is present or due to the characteristics of the environment. Nonetheless, once the STA initiates the connection to a new AP, the data transfer is halted until the handshake stage is completed.

As the handover procedure introduces lapses in the connection, it is desirable to minimise these periods as much as possible. In this respect, several amendments have been made to the IEEE 802.11 standard in an effort to improve the handover-related performance.

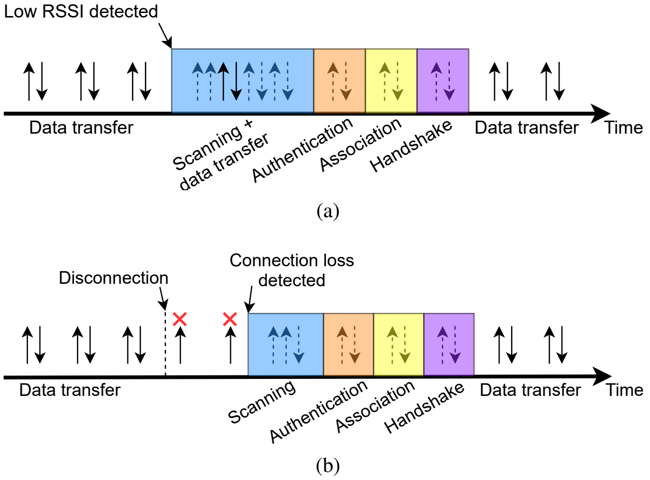

Fig. 1: Simplified illustration of the 4 stages which occur during a handover for the two different situations: a) handover triggered by low RSSI, and b) handover triggered by detecting that the AP cannot be reached (disconnection).

A. IEEE 802.11k

The IEEE 802.11k amendment allows for the STA and AP to generate and share information about the radio environment. Instead of simply choosing the most optimal AP in terms of RSSI, this can allow for STAs to request information regarding APs in the environment to optimise the choice of AP when roaming [5]. The amendment also enables the STA to request its current AP for a neighbour report containing information about other APs in the same Extended Service Set (ESS) and serve equal network settings. Based on the extra information available, the STA can therefore roam to an underutilised network. Although using IEEE 802.11k in Wi-Fi 6 networks supporting OFDMA may not yield improvements to the same degree due to increased scalability performance, it will still be sensible to distribute the overall load. The processing of neighbour reports is an implementation-specific feature that requires full compatibility with the STA.

B. IEEE 802.11v

This amendment builds upon IEEE 802.11k to perform active load balancing through BSS Transition Management (BSS-TM). The network can send suggestions to the STA in order to steer it to another AP in which it may have better service [7]. It can likewise be utilised to redirect poorly connected clients. This allows the network to also contribute to the decision of which APs a STA chooses to connect to, which would otherwise be decided solely by the STA. This in turn helps to mitigate the effect of the Listen-Before-Talk (LBT) mechanism of Wi-Fi, where the medium access is highly dependent on the number of active users.

C. IEEE 802.11r

The IEEE 802.11r amendment enables the use of Fast BSS Transition. After the STA initially connects to an AP following the normal stages, the following handovers to other APs will contain fewer handshake messages, allowing for a faster handover to the new AP. This however requires both the network and the device to support this, and furthermore only works when a STA roams between APs in the same ESS and using either Pre-Shared Keys (PSK) or IEEE 802.1X authentication. IEEE 802.11r also allows for AP-assisted roaming through a Distributed System (DS), known as over-the-DS roaming as opposed to the traditional over-the-air roaming. Here, the STA can communicate with a target AP through the current serving AP in case both APs are connected through the same backend. By offloading some of the steps to the APs which can communicate using a contention-free medium, the handover can in general be improved in terms of duration.

III. TEST SETUP

Performance evaluation of various Wi-Fi configurations was performed at the AAU 5G Smart Production Lab at Aalborg University [2]. This industrial environment is equipped with three ceiling-mounted CISCO MR36 Enterprise Wi-Fi 6 access points [8] distributed throughout the lab, as indicated in Fig. 2. This equipment represents an off-the-shelf enterprise grade Wi-Fi 6 deployment, as opposed to specialised Wi-Fi solutions where same-vendor STA and AP devices are optimised to meet stringent IIoT requirements. Thus, the work seeks to evaluate the achievable performance using off-the-shelf Wi-Fi 6 solution together with flexible choice of STAs. Because OFDMA was not supported at the time of writing, the impact of this could not be investigated. This is, however, not expected to have a significant impact on the tests and results, as only a few low-throughput devices will be connected to an AP at any given time. The CISCO Client Balancing feature was enabled throughout all the tests to enable the IEEE 802.11v BSS-TM functionalities. A MiR200 AMR was used to enable the mobility aspect of the setup. In the mobility tests, the AMR was configured to follow a specified route through the lab, bringing the measurement STA setup through each AP coverage area. The automated route was chosen to maximise the number of handovers to better determine its impact on the link performance. The robot moves with a maximum speed of 1.5 m/s and provides simultaneously positioning information data through an internal mapping system with 5 cm accuracy.

Measurements were collected using the STA described in Table I. The STA was configured to utilise wpa_supplicant v2.9 [9], which is commonly used among a wide variety of platforms. The software communicates with the driver of the Wi-Fi card and handles roaming and key negotiation. It is furthermore used to obtain statistics regarding connection state, signal strength, and communication throughput. The route to the target device on the network is added as a static entry to the STAs link-level routing table to avoid overhead from discovery protocols. The round-trip time (RTT) latency is measured by utilising the Linux ping functionality with a packet size of 64 B and an inter-packet interval of 50 ms, communicating with a network edge-cloud device connected to the different APs through Ethernet. This allowed for emulation of an overall application data rate of 10.2 kbit/s, which is comparable to that of typical IIoT processes such as the fleet manager-based control of AMRs or the control of PLCs in production lines [2]. When the STA detects an RSSI of -85 dBm, a scan is requested through the wpa_supplicant, which then triggers a roaming event and checks whether another AP with significantly better RSSI is nearby. This RSSI-based roaming is necessary to enable the IEEE 802.11r roaming functionality, which does not trigger for timeout-based roaming, i.e. when the STA loses connection to its current AP. Because this RSSI threshold also has an impact on the supported data rate, it should be chosen with the given IIoT application in mind. In this case, as our IIoT application is low data rate, a lower RSSI could be set.

TABLE I: Details of the measurement STA hardware and software setup.

HW/SW	Details
Device Model	Intel NUC Board NUC5i3MYBE
CPU	Intel i3-5010U @ 2.10 GHz
RAM	8GB @ 1600 MHz
OS	Ubuntu 20.04.1 LTS
Kernel	5.4.0-52-generic
Wi-Fi Network Card	Intel Wi-Fi 6 AX200

Fig. 2: Overview of the industrial environment and the measurement setup including the AMR with the STA used in the testing and one of the ceiling-mounted Wi-Fi 6 APs.

Under the current configuration and measurement route (see Fig. 3 depicting the lab layout, and where the APs are indicated with red dots), a single handover event occurs every two minutes on average, corresponding mainly to when the robot roams between the two labs. A total of 45 handovers will occur for each test. The impact of an unoptimized handover is estimated to last ∼520 ms including the scanning period based on the measurements, resulting in 1% of the measurements. We can, therefore, expect to see the difference in terms of handover performance around the 99%-ile.

In order to determine the impact of handovers on the link latency performance, four Wi-Fi configuration schemes were considered:

1) Baseline: IEEE 802.11v features enabled. This is not expected to provide any notable benefit for the given setup, as the load on all APs will be comparable.
2) Optimised Scanning: the list of channels from which a STA can scan for APs will be reduced to only include the frequencies of present APs. This will reduce the number from a default of 38 to 3.
3) IEEE 802.11r: over-the-air roaming features will be utilised to reduce the handover time itself.
4) Optimised Scanning and IEEE 802.11r: both features are enabled simultaneously.

As described in Section II, the mobility features from IEEE 802.11k and IEEE 802.11v (which builds on top of IEEE 802.11k) are implementation-specific and require fully compatible STAs to operate them. Unfortunately, these elements are not supported in our current setup and thus its evaluation is left for future work.

To get insight on different deployment situations, the four Wi-Fi 6 configuration schemes enumerated about were examined over the following network configurations:

1) Idle network (single STA under test) with dedicated frequency channels at each AP.
2) Network with controlled load background traffic and dedicated frequency channels at each AP.
3) Network with controlled load background traffic and frequency re-use across APs.

The Wi-Fi spectrum at the lab is fully controlled. Each AP operates on their own 5 GHz frequency channel with 20 MHz bandwidth, except for the last test, where the APs will be configured to use the same channel. Of course, larger bandwidths are supported, but this allocation is enough for the aim presented in this paper. For the first test, only a controlled load, dedicated frequency channel will be connected to an AP at any given time. Although other nearby STAs from a different network in the area may choose to scan the channel for APs (this is an ISM band), the interference experienced in this setup is negligible. For the remaining two tests, two additional STAs will be connected to each AP, with each either sending or receiving 10 Mbit/s UDP traffic generated using iperf3 [10], resulting in 10 Mbit/s uplink and downlink interference traffic per AP. The location of these STAs is shown with green dots in Fig. 3. This traffic load was chosen to reflect a low-medium usage of the network, with sufficient traffic to impact the communication while not reaching congestive conditions.

For completeness, reference measurement tests were also performed for intra-AP static (non-mobility) and intra-AP mobility situations for the idle and the controlled load cases.

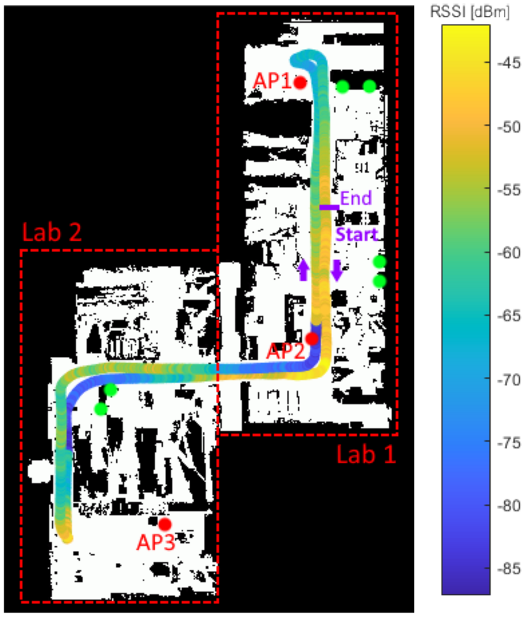

Fig. 3: RSSI heatmap for a single measurement lap for baseline scheme and idle network configuration. The location of the APs and the controlled traffic load source/sink STAs are indicated by the red dots and green dots, respectively.

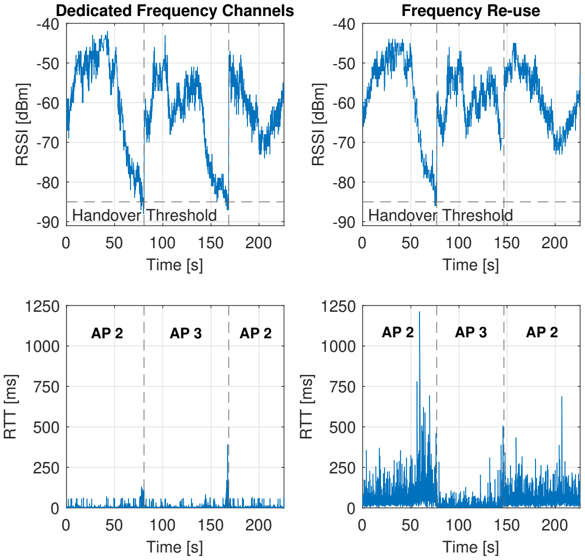

Fig. 4: RSSI and RTT measurements for a single measurement lap for the baseline scheme and controlled load network configurations with dedicated frequency channels (on the left), and with frequency re-use (on the right).

IV. TEST RESULTS

A heatmap of the RSSI for a single measurement lap under idle network conditions is illustrated in Fig. 3. Here it is shown that, in this particular example, the STA did not roam to AP 1 since the RSSI was approximately -70 dBm. This was, however, not the case for all of the measurements, as the STA would occasionally roam to it when different propagation conditions applied or higher interference was present. The heatmap further shows that there is a clear overlap in terms of coverage area between AP 2 and AP 3. Although these coverage areas can be better planned by changing the deployment positions of the APs or changing their transmit power, this is left out of the focus of this study, as it has

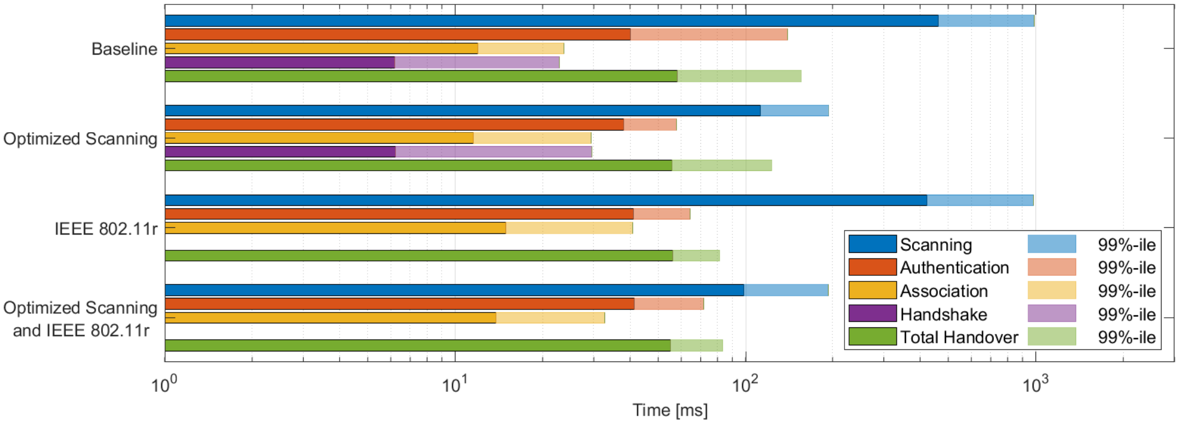

Fig. 5: Mean (dark solid) and 99%-ile (light solid) duration of the different handover stages for the idle network station with a single STA for the different Wi-Fi optimised configurations. Total handover = authentication + association + handshake.

no impact on our handover measurements (e.g. handovers will still happen between these two APs). Nonetheless, the heatmap helps us to determine the location of the handover regions, and most handovers were found to be concentrated in the same area.

Fig. 4 shows the correlation between RSSI and RTT for a single lap starting from AP 2 and moving towards AP 3 in idle network conditions without mobility optimizations. On the left part of the figure, it is shown that the STA will initiate the scanning process once -85 dBm is reached for the dedicated multi-channel configuration. During this phase, there is a slight increase in latency due to the scanning, followed by a large latency spike from the handover itself. When using a single-channel frequency re-use configuration where the three APs overlap, the overall latency is much higher (with exception of the instances where the STA is connected to AP 3). This is shown on the right part of the figure. This is due to the fact that in our industrial scenario, Lab 2, is separated from Lab 1 by a thick high-isolation wall, which blocks a significant part of the interference. When the STA moves back to Lab 1, we observe a handover occurring earlier than for the dedicated frequency channel case at -73 dBm caused by the timeout-based roaming event described in Fig. 1b. The reason that we do not see a spike in terms of RTT during this handover is that the impact of the interference is, in this case, much more significant than the handover itself.

By using the data from wpa_supplicant, it can be determined how much time is spent during each stage of the handover, which is detailed in Fig. 5 for the different Wi-Fi 6 configuration schemes in idle network conditions. For the baseline configuration, the scanning time is significantly larger than the handover time (authentication, association and handshake) itself. Since the STA may only need to scan a single channel for timeout-based roaming, the period without data transfer can be minimal. However, if an AP is not found and a full scan is required, the time until connectivity is restored will correspond to the scanning time and the total handover duration combined, assuming that an AP is found in the second search. Reducing the number of channels to scan significantly reduces this duration for the general scans, as well as the 99%-ile due to the variance of the time spent scanning each channel. The mean duration of the handover itself, i.e. excluding the scan, is constant across all configurations. While IEEE 802.11r skips the handshake stage (normally lasting around 6 ms), a slight increase of approximately 3 ms in the duration association stage was observed. The 99%-ile of the total handover duration is nonetheless improved due to the removal of the handshake stage. The benefit of using this feature is, therefore, seemingly negligible for the mean duration, but it should be noted that this is for best-case conditions without background traffic and by using WPA2-PSK encryption. Note that using enterprise encryption (e.g. with IEEE 802.1X), where a separate server may be contacted to obtain access, further gains by using IEEE 802.11r are expected as some of these steps may be skipped.

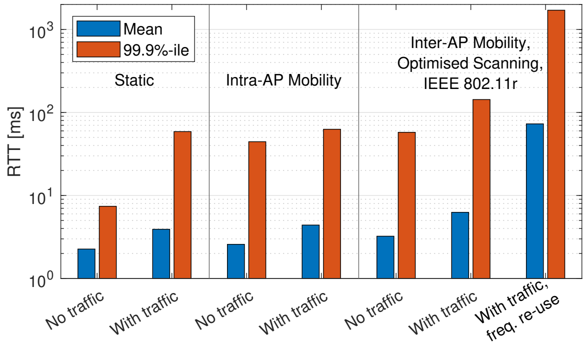

Fig. 6: Mean and 99%-ile RTT latency measurement results for the intra-AP static, intra-AP mobility and inter-AP mobility cases. Dedicated frequency channels were used unless otherwise specified.

Fig. 6 summarises the mean and 99%-ile RTT values for the different inter-AP mobility measurements (with handovers), as well as for the static and intra-AP mobility for reference (without handovers). Measurements for intra-AP mobility were gathered in the area around AP 2, while for

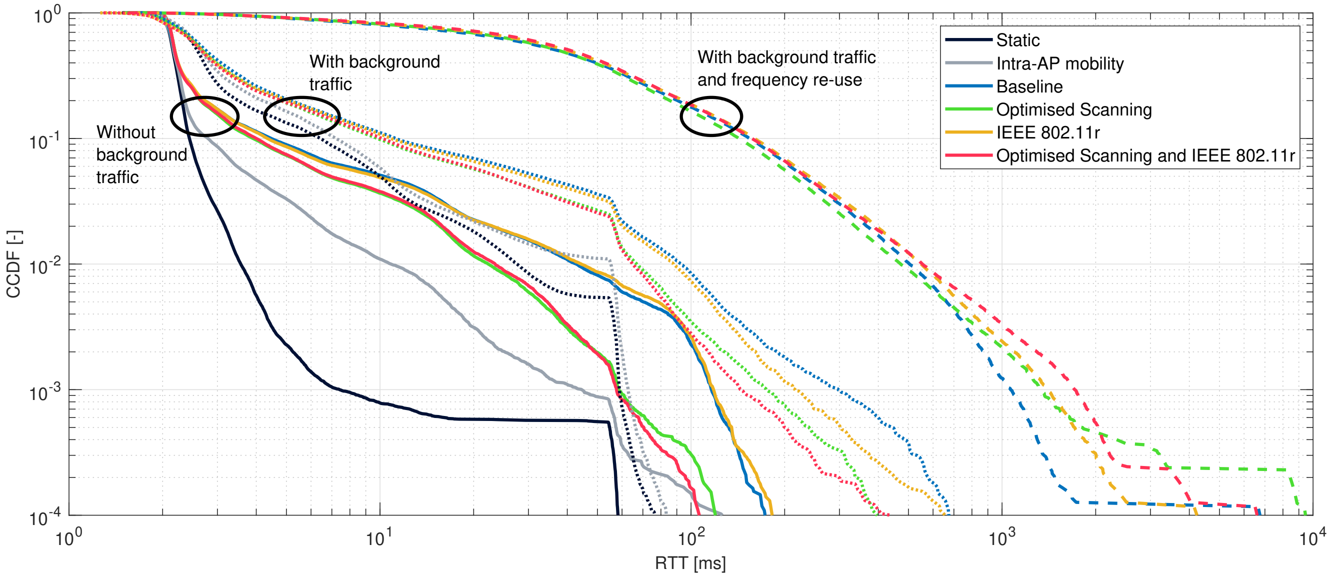

Fig. 7: RTT Empirical CCDFs for all the different Wi-Fi schemes and network configurations explored in the measurements. Dedicated frequency channels were used unless otherwise specified.

TABLE II: Summary of RTT and PER measurement results for the different Wi-Fi schemes and network configuration setups.

Setup		**Min**	**Avg**	**99.9%-ile**	**Jitter**	**PER**
Static Reference	Idle network (single STA)	1.5 ms	2.3 ms	7.4 ms	0.3 ms	0%
	Background traffic, dedicated frequency channels	1.5 ms	3.9 ms	58.8 ms	2.9 ms	0%
Intra-AP Mobility Reference	Idle network (single STA)	1.6 ms	2.6 ms	44.4 ms	0.8 ms	0%
	Background traffic, dedicated frequency channels	1.6 ms	4.4 ms	62.7 ms	3.5 ms	0%
Baseline	Idle network (single STA)	1.8 ms	3.9 ms	116.0 ms	2.4 ms	0.043%
	Background traffic, dedicated frequency channels	1.6 ms	7.8 ms	297.0 ms	5.9 ms	0.066%
	Background traffic, frequency re-use	1.6 ms	67.2 ms	1062.0 ms	37.9 ms	0.119%
Optimised Scanning	Idle network (single STA)	1.6 ms	3.2 ms	58.3 ms	1.5 ms	0.045%
	Background traffic, dedicated frequency channels	1.6 ms	6.4 ms	174.0 ms	4.8 ms	0.071%
	Background traffic, frequency re-use	1.5 ms	66.6 ms	1320.0 ms	37.5 ms	0.103%
IEEE 802.11r	Idle network (single STA)	1.6 ms	4.0 ms	118.0 ms	2.4 ms	0.046%
	Background traffic, dedicated frequency channels	1.3 ms	7.3 ms	215.0 ms	5.8 ms	0.073%
	Background traffic, frequency re-use	1.6 ms	71.5 ms	1391.0 ms	39.0 ms	0.109%
Optimised Scanning and IEEE 802.11r	Idle network (single STA)	1.8 ms	3.2 ms	57.6 ms	1.6 ms	0.044%
	Background traffic, dedicated frequency channels	1.3 ms	6.3 ms	143.0 ms	4.8 ms	0.065%
	Background traffic, frequency re-use	1.6 ms	72.9 ms	1704.0 ms	39.6 ms	0.107%

the intra-AP static measurements, the data was obtained from four static locations close to the measurement route. The RTT was measured using the same configuration as for the inter-AP configurations, but without any handovers, naturally. Introducing mobility to a Wi-Fi connection results, even in the intra-AP case, in additional latency, albeit mainly in the lower percentiles. Nonetheless, the additional 36 ms for the 99.9%-ile for idle networks with a single STA and without background traffic is a notable impact that must be taken into account for IIoT applications. If the STA roams between APs, the latency is further increased by a considerable amount for both the mean and 99.9%-ile levels. The presence of background traffic will, moreover, increase latency in any setting regardless of mobility, which is expected from the LBT mechanism. However, if frequency re-use is utilised, the overlapping networks will cause much more severe delays in the communication compared to the other cases.

Empirical complementary cumulative distribution functions (CCDF) computed over more than 100,000 RTT latency samples per Wi-Fi and network configurations are shown in Fig. 7 with their key statistics and Packet Error Rate (PER) summarised in Table II. As detailed, the overall latency distribution is highly affected by the amount of background traffic present in the network due to impact on the LBT mechanism, increasing the latency for all percentiles. It is however also shown that improving aspects related to the handover will result in improved latency after the 90%-ile. This is especially evident by optimising the scanning stage, which further confirms that the scanning period is one of the main contributors to handover-related latency, both in cases with and without interference load on the network. With optimised scanning, the jitter is likewise reduced by 1 ms for all conditions. While the benefit of using IEEE 802.11r is negligible for idle network conditions, it has a notable impact on loaded networks around the 99%-ile. As stated previously, larger improvements can be expected in Wi-Fi deployments using enterprise-level authentication and IEEE 802.1X. If frequency re-use is utilised for all APs, it is evident that the performance is severely affected. In this case, the mobility optimization mechanisms do not exhibit as large gains as in the other cases, with the latency at the 99.9%-ile exceeding 1 second. The increased latency is generally caused by interference, but also due to the roaming being triggered by a timeout-mechanism shown in Fig. 4.

By comparing the inter-AP mobility distributions with the intra-AP mobility one, it is observed that the impact from handovers is more significant from the 90%-ile to 99%-ile. Close to the 99.9%-ile, the latency distributions converge and the performance of the inter-AP mobility with optimised handovers is similar to that of the intra-AP mobility, indicating that other environmental factors contribute to the latency when considering lower percentiles. When comparing the performance of the static case and the intra-AP mobility case without background traffic, it is clear that mobility itself introduces some additional latency. Because similar conditions in terms of radio channel variations are possible for static deployments (in case of e.g. scattering from other mobile objects), mobility itself cannot be seen as the only bottleneck towards high reliability for latency.

In terms of PER, it was observed that, in general, packet losses occurred mainly during the handover processes for all the cases. Background interference has a clear impact on the PER and an increased PER of $\sim$0.1% was found, in the worst case, when frequency-reuse was used, as compared to the $\sim$0.06-0.07% for the dedicated frequency channels case and to the $\sim$0.04% for the idle network single STA case.

V. CONCLUSION

In this paper we investigated the latency performance of Wi-Fi for static and mobile IIoT conditions with emphasis on the associated handover under mobility. The study was done experimentally in an industrial production environment at the AAU 5G Smart Production Lab at Aalborg University. The measurements were done on a commercial enterprise-grade Wi-Fi 6 system using a Linux-based STA device roaming in a predetermined path through three Wi-Fi coverage areas.

When the STA conducts a handover, the main delay contribution originates from scanning for new APs. This can last up-towards 1 second, in which the latency of data transfers is increased significantly. This is followed by the $\sim$55 ms of the handover itself in which no data communication is possible. If the STA disconnects completely from the previous AP before the handover is initiated, a faster single-channel scan is utilised. This scan is based on previously observed APs and allows for reducing the time required for the scan. If an AP is not found, a full scan will be required. Since the STA is disconnected from the AP, no data transfer can occur until a new connection is established.

A RTT latency of $\sim$110 ms was measured for the 99.9%-ile in an idle network using a non-optimised baseline configuration. By using optimisations targeting the handover, where the number of channels to scan was drastically lowered and IEEE 802.11r was utilised to shorten the handover itself, a latency of 58 ms was achieved for the same percentile. Utilising the same improvements for loaded networks resulted in a reduction in latency from 297 ms to 143 ms, with benefits of IEEE 802.11r being more notable due to the interference of present traffic, which would otherwise have introduced delays in the communication between the STA and AP. If a large amount of interference is present, such as when using frequency re-use among APs, the contribution in latency from handovers become negligible. The latency performance at the 99.9-99.999% level appears to be dominated by limitations in the Wi-Fi solution to capture dynamic channel changes as handovers have little impact beyond the performance seen with intra-AP mobility.

The achieved latency values with mobility are on the high end for many IIoT applications, often requiring $<$10-100 ms performance and at higher levels of reliability than used in this paper (ex. up to 99.999% reliability). As shown in the paper, using dedicated clean channels helps mobility performance and it is important that the load is managed in the network. Specialised IIoT Wi-Fi solutions optimised for latency will still be needed for such challenging applications.

REFERENCES

[1] A. Varghese and D. Tandur, “Wireless requirements and challenges in Industry 4.0,” in *International Conference on Contemporary Computing and Informatics (IC3I)*, Nov. 2014, pp. 634–638.

[2] I. Rodriguez *et al.*, “5G Swarm Production: Advanced Industrial Manufacturing Concepts Enabled by Wireless Automation,” *IEEE Communications Magazine*, vol. 59, no. 1, pp. 48–54, 2021.

[3] H. Cho, S. Shin, D. Han, and J. Chung, “Analysis of Wi-Fi data interruption and handover decision parameters,” in *18th IEEE International Symposium on Consumer Electronics (ISCE 2014)*, Jun. 2014, pp. 1–2.

[4] A. A. Tabassam, H. Trsek, S. Heiss, and J. Jasperneite, “Fast and seamless handover for secure mobile industrial applications with 802.11r,” in *34th IEEE Conference on Local Computer Networks*, Oct. 2009, pp. 750–757.

[5] S. Feirer and T. Sauter, “Seamless handover in industrial WLAN using IEEE 802.11k,” in *26th IEEE International Symposium on Industrial Electronics (ISIE)*, Jun. 2017, pp. 1234–1239.

[6] S. Chatterjee, D. Sarddar, J. Saha, S. Banerjee, A. Mondal, and M. K. Naskar, “An improved mobility management technique for IEEE 802.11 based WLAN by predicting the direction of the mobile node,” in *National conference on computing and communication systems*, Nov. 2012, pp. 1–5.

[7] IEEE, “IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Amendment 8: IEEE 802.11 Wireless Network Management,” *IEEE Std 802.11v-2011*, pp. 1–433, 2011.

[8] “CISCO MR36 Cloud-managed Wi-Fi 6 (802.11ax).” [Online]. Available: https://meraki.cisco.com/product/wi-fi/indoor-access-points/mr36/

[9] “Linux WPA Supplicant (IEEE 802.1X, WPA, WPA2, RSN, IEEE 802.11i).” [Online]. Available: https://w1.fi/wpa_supplicant/

[10] “iPerf - The TCP, UDP and SCTP network bandwidth measurement tool.” [Online]. Available: https://iperf.fr/