Enhanced Wi-Fi 6 Wireless Communication Systems for Industrial Use Cases Requiring Mobility

Communication Technology
- Aalborg University -

Master Thesis

Andreas Engelsen Fink

AALBORG UNIVERSITY

NOKIA Bell Labs

Aalborg University
Department of Electronic Systems
Fredrik Bajers Vej 7B
DK-9220 Aalborg

AALBORG UNIVERSITY

STUDENT REPORT

Department of Electronic Systems
Fredrik Bajers Vej 7
DK-9220 Aalborg Ø
http://es.aau.dk

Title:
Enhanced Wi-Fi 6 Wireless Communication Systems for Industrial Use Cases Requiring Mobility

Theme:
Communication Technology

Project Period:
Master Thesis, 3rd and 4th semester M.Sc. E.

Project Group:
CT10-925

Participant(s):
Andreas Engelsen Fink

Supervisor(s):
Ignacio Rodriguez Larrad

External Supervisor(s):
Troels Kolding (Nokia Bell Labs)

ECTS: 50

Number of Pages: 91

Date of Completion:
June 2, 2021

Abstract:

With more intelligent and flexible Industrial IoT (IIoT) emerging in industrial plants, the requirements for the communication are becoming more demanding. As mobility is a key enabler for Industry 4.0 applications, the integration of wireless technologies with mobile industrial applications is gaining momentum. However, the performance over a wireless medium may not be sufficient for applications with critical latency requirements. The objective of this thesis is to explore the boundaries of wireless technologies, specifically Wi-Fi, to meet these use-cases, as well as to develop new methods to improve them with specific emphasis on mobility and high levels of reliability.
To determine how well Wi-Fi can support a given latency requirement, its performance with emphasis on reliable latencies was measured in a realistic industrial environment using commercially available Wi-Fi equipment. The main contribution to latency was found to be the handovers between Access Points, increasing the 99.9%-ile latency from 63 ms to 297 ms.
To both improve the overall performance using Wi-Fi as well as to fully mitigate the impact of handovers, a solution using multi-connectivity was pursued. This resulted in a radio-aware multi-connectivity layer-4 scheduling mechanism. Using this, a latency of 77 ms was observed at the 99.9%-ile. With a 74% reduction compared to the non-optimized Wi-Fi, this solution can better support critical IIoT applications.

AALBORG UNIVERSITET
STUDENTERRAPPORT

Institut for Elektroniske Systemer
Fredrik Bajers Vej 7
DK-9220 Aalborg Ø
http://es.aau.dk

Titel:
Optimerede Wi-Fi 6 Trådløse Kommunikationssystemer for Industrielle Brugsscenarier med Mobilitet

Tema:
Communication Technology

Projektperiode:
Kandidatspeciale, 3. og 4. semester M.Sc. E.

Projektgruppe:
CT10-925

Deltager(e):
Andreas Engelsen Fink

Vejleder(e):
Ignacio Rodriguez Larrad

Ekstern Vejlder(e):
Troels Kolding (Nokia Bell Labs)

ECTS: 50

Sidetal: 91

Afleveringsdato:
2. juni 2021

Abstract:

Med mere intelligente og fleksible Industrielt IoT (IIoT) i industrianlæg bliver det vanskeligt at opfylde kravene for den underliggende kommunikation at kunne understøtte dette. Eftersom mobilitet bliver et centralt aspekt i Industri 4.0-applikationer, vil trådløse teknologier i stigende grad integreres i industrielle omgivelser. Ydelsen over trådløse medier er imidlertid muligvis ikke tilstrækkelig til applikationer med kritiske krav til latenstid. Formålet med denne afhandling er at undersøge grænserne for trådløse teknologier, især Wi-Fi, for at opfylde kravene til disse brugstilfælde og udvikle nye metoder til at forbedre dem med særligt fokus på mobilitet og høj pålidelighed.
For at fastslå hvor godt Wi-Fi kan opfylde et givet krav til latenstid blev dets ydeevne med vægt på pålidelige latenser målt i et realistisk industrielt miljø med kommercielt Wi-Fi udstyr. Det blev her konstateret at det største bidrag til latenstiden var overgangen mellem to netværk som øgede latenstiden ved 99.9%-ilen fra 63 ms til 297 ms.
Multi-konnektivitet blev undersøgt for at både forbedre den overordnede ydelse samt reducere påvirkningen af overgange. Den primære tilgang til at forbedre den overordnede ydelse samt reducere påvirkningen af overgangene er gennem multikonnektivitet. Dette resulterede i designet af en radiofølsom lag-4 pakkeskeduleringsmekanisme som yderligere sænkede latenstiden til 77 ms. Med en reduktion på 74 % kan denne løsning kan derved bedre understøtte kritiske IIoT-applikationer.

Preface

This long master's thesis was written during the 3rd and 4th semester of the study Communication Technology with specialization in Networks and Distributed Systems at the department of Electronic Systems at Aalborg University. The project started September 2nd 2020 and ended June 3rd 2021.

The topic of this project is empirical evaluation of enterprise Wi-Fi for industrial use cases requiring mobility. This project is collaborating with Nokia Bell Labs, Aalborg. The author would like to thank the supervisors Ignacio Rodriguez Larrad and Troels Kolding for their guidance and support throughout the project, as well as Guillermo Pocovi from Nokia Bell Labs, and Anders Karstensen and Rasmus Suhr Mogensen from Aalborg University.

The thesis is structured as a main part and appendix with worksheets. The main part includes the extended summary and the written papers. The worksheets and appendix are only used for supporting the main work.

Citations are marked in square-brackets with a number corresponding to a source in the bibliography, e.g. [1]. Citations in the papers contain their own respective bibliography.

Contents

Chapter 1

Extended Summary

As we approach Industry 4.0, Industrial Internet of Things (IIoT) is becoming increasingly integrated in industrial plants. Aspects such as improved adaptability to the current state of e.g., a production line, improved machine-to-machine communication and generally more powerful hardware play critical roles in automated environments. To enable these aspects, it is crucial that the underlying communication can support the requirements of the devices, such as increased throughput or real-time communication. Moreover, to improve flexibility in terms of set-up and connectivity while enabling new use cases such as a device requiring full mobility, integrating wireless communication technologies will be a key enabler. However, communicating in the wireless domain will result in degraded performance as compared to cabled, with increased latency and packet loss. It is therefore of interest to investigate these aspects in an industrial setting, which is the main focus of this project. Although technologies such as 5G NR are currently being developed targeting Ultra-Reliable Low Latency Communication (URLLC), it remains of interest to consider IEEE 802.11 Wi-Fi, as this is prominent in industrial plans today and does not depend on the use of licensed frequency bands. While this technology does not claim to support URLLC, it may be able to support a subset of time-sensitive applications. To evaluate the performance of the technology, the reliable latency will be used as one of the key metrics. This is the packet latency achieved at a certain reliability level, such as the 99.9%-ile representing the top 0.1% of measured latencies in a given dataset.

In order to raise a main problem statement for the project, an investigation into the Wi-Fi protocol was performed as to obtain an initial understanding of its design. This, in combination with preliminary measurements, led to an analysis of how mobility is handled in Wi-Fi in terms of how the Station (STA) roams between Access Points (AP). Because the technology is designed to use a 'break-before-make' approach, a disruption in the traffic will occur, which can be fatal for applications with critical communication requirements. Furthermore, with IIoT-devices requiring mobility, it becomes necessary to investigate the impact of these handovers and how their impact on the communication can be reduced.
Based on this, the main problem statement for this project is as follows:

"*In order for IEEE 802.11 Wi-Fi to better satisfy the requirements of time-sensitive IIoT in settings requiring mobility, the impact of handovers in communication needs to be mitigated.*"

To investigate the problem statement, the following research questions are defined:

RQ1: What is the current performance of enterprise Wi-Fi, and what are the main contributors to the latency?

RQ2: How can the latency and Packet Error Rate (PER) be reduced using non-proprietary mechanisms targeting mobility and latency?

RQ3: How can the network be deployed as to support devices with strict communication requirements?

RQ4: How can a standardized solution be designed such that the latency and PER of the communication is improved?

When a STA needs to roam to a new network due to e.g., low signal strength, it will initially scan for nearby APs which it can connect to. During this period, the performance of on-going communication will be degraded while the STA communicates with APs on other frequency bands. Once an eligible AP is found, the STA disconnects from its current AP and initiates the handover. During this period, communication is halted until the connection is established. To mitigate this issue, several amendments have been introduced to the IEEE 802.11 protocol with the goal of enabling seamless roaming, namely IEEE 802.11r, IEEE 802.11k and IEEE 802.11v.

To empirically determine the performance of Wi-Fi in real-life industrial settings, custom measurement software was designed and implemented. This software captures Wi-Fi properties such as RSSI and connection state, positioning data using an Autonomous Mobile Robot (AMR) and communication latency and PER statistics. Likewise, to identify the benefits of improving handover-specific properties the following improvements were investigated: 1) using IEEE 802.11r to reduce the communication required during handover through pre-authentication, and 2) simple radio tuning to reduce the number of channels to scan before the handover.

This led to the work presented in Paper 1, in which the performance of a single STA in an industrial setting was measured. To evaluate the impact of handovers, the reliability level of 99.9% is used, with 100.000 samples captured to ensure sufficient confidence. With the test setup configured to adequately capture the effects of the handovers, this led to ∼90 minutes of measurements per configuration. It was through these tests confirmed that handovers do indeed have a significant impact on the latency, with the latency of the 99.9%-ile (i.e., the latency of 0.01% of the communication) being increased from 62 ms to 297 ms for network conditions with background traffic, addressing RQ1. By applying the optimizations, the latency for the same percentile was decreased to 174 ms, successfully mitigating some of the impact from handovers, addressing RQ2. Specifically, reducing the number of channels during network scans would significantly reduce the latency impact of handovers. Furthermore, the performance during single-frequency deployments across APs was investigated, from which the effect of the Listen-Before-Talk mechanism of Wi-Fi was likewise found to substantially increase the latency

while doubling the PER as compared to deployments with dedicated frequency channels, addressing RQ3. Further details for the Wi-Fi protocol and the setup for the tests are available in the worksheets found in Sections 3.1 and 3.2.

After investigating the performance of a typical Wi-Fi STA in an industrial setting, it is clear that while the impact of the handovers is reduced, but further improvements are needed. A state-of-the-art approach for this is using transport-layer multi-connectivity, where traffic can be either routed using specific connections or duplicated over both. To address RQ4, this approach was adopted with the aim of including radio-level information to both improve the performance in general, and to allow for truly seamless handovers. This resulted in the design of a multi-STA solution, i.e. a device with multiple interfaces. Another objective of this multi-STA is to make it network-agnostic, with it being able to be integrated in current typical Wi-Fi deployments without affecting both devices on the network or the network itself, such as e.g. proprietary solutions would, which in turn could lead to vendor lock-in.

The multi-STA operates using a packet scheduler to handle transport-layer traffic based on radio-level properties, and a mobility coordinator to negate the impact of handovers and to improve signal diversity by connecting each STA to different APs. By both incorporating knowledge of the Wi-Fi signal and connection state as well as introducing radio-level control of the STAs, this is a novel approach to multi-connectivity in Wi-Fi. The multi-STA was designed with two main components: a transport-layer packet scheduler, from which the traffic would be duplicated over both STAs or only the STA with highest RSSI was used, and a mobility coordinator to ensure the STAs would connect to separate APs and avoid simultaneous handovers.

Paper 2 presents how having multiple APs overlap in terms of coverage area can allow for a multi-STA to achieve significantly improved performance over the optimized single-STA configuration. All of the evaluated multi-STA configurations were found to improve both latency and PER metrics, with the packet scheduler scheme and best path scheduling resulting in 99.9%-ile latencies of 80 ms and 77 ms, respectively. By including a mobility coordinator, the average time until medium access can be reduced in the packet duplication configuration. Furthermore, this ensures that the effect of handovers is completely negated, as opposed to only being decreased in single-STA configurations. More information regarding the design of the multi-connectivity solution is available in the worksheet in Section 3.3.

In conclusion and referring back to the main problem statement, it has been experimentally verified that roaming in Wi-Fi does indeed have a significant impact on the latency, with both network scans and the handover itself resulting in either periods of degraded performance or full interruptions in the communication. Although optimizations targeting this issue can be utilized in commodity Wi-Fi networks, further improvements would be required to achieve the same performance as that without handovers. To accomplish this, a multi-connectivity solution was designed and validated, in which the handovers

could be fully negated while improving the overall performance. This was found to yield the best performance across all tested configurations, with a decrease in the 99.9%-ile of 74% (297 ms to 77 ms), as well as almost completely negating packet drops.

Chapter 2

Papers

2.1 Paper 1

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

Andreas Engelsen Fink, Rasmus Suhr Mogensen, Ignacio Rodriguez, Troels Kolding, Anders Karstensen, Guillermo Pocovi

A. E. Fink, R. S. Mogensen, I. Rodriguez and A. Karstensen are with the Wireless Communication Networks Section, Department of Electronic Systems, Aalborg University, Aalborg Øst 9220, Denmark (email: aef16@student.aau.dk; rsm@es.aau.dk; irl@es.aau.dk; andka@es.aau.dk).
T. Kolding and G. Pocovi are with Nokia - Bell Labs, Research Center Aalborg, Aalborg Øst 9220, Denmark (email: troels.kolding@nokia-bell-labs.com, guillermo.pocovi@nokia-bell-labs.com)

The paper has been accepted for the
European Wireless Conference (EW) 2021.

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 improve-

ments. 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.

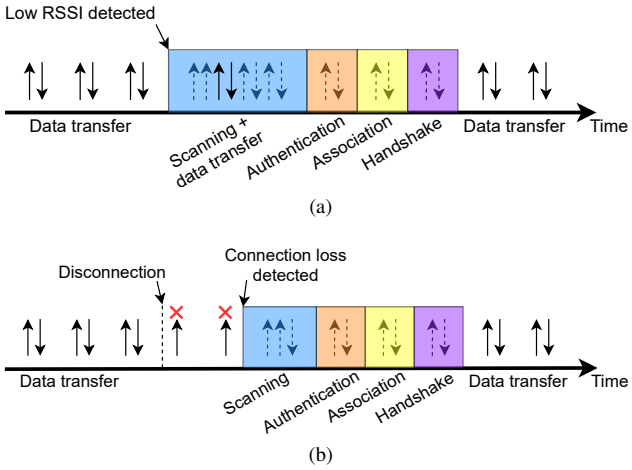

Fig. 1: Simplified illustration of the 4 stages which occur during a handover for the two different situations: a) handover triggered by low RSSI, and b) handover triggered by detecting that the AP cannot be reached (disconnection).

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.

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,

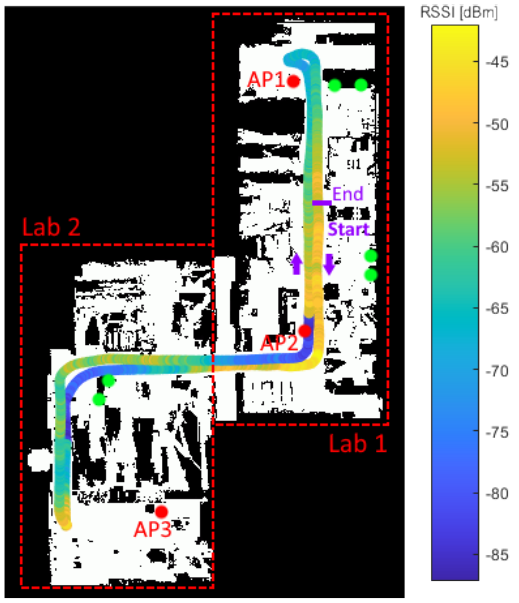

Fig. 3: RSSI heatmap for a single measurement lap for baseline scheme and idle network configuration. The location of the APs and the controlled traffic load source/sink STAs are indicated by the red dots and green dots, respectively.

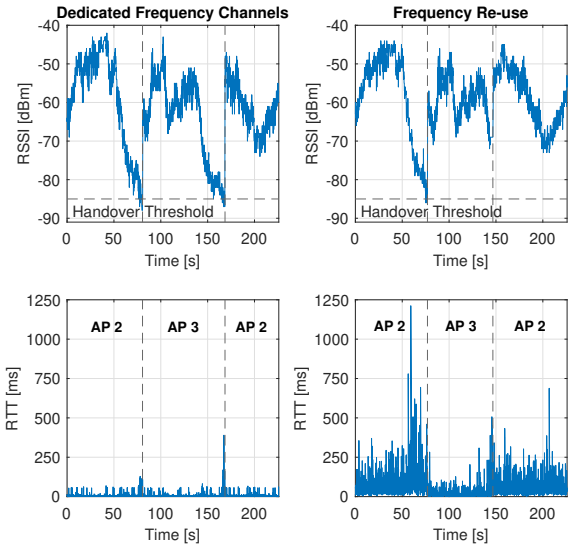

Fig. 4: RSSI and RTT measurements for a single measurement lap for the baseline scheme and controlled load network configurations with dedicated frequency channels (on the left), and with frequency re-use (on the right).

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

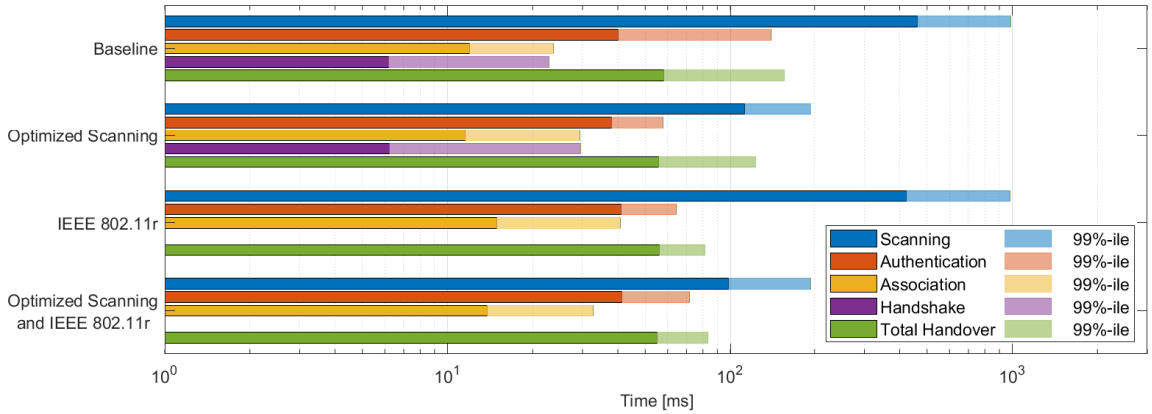

Fig. 5: Mean (dark solid) and 99%-ile (light solid) duration of the different handover stages for the idle network station with a single STA for the different Wi-Fi optimised configurations. Total handover = authentication + association + handshake.

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

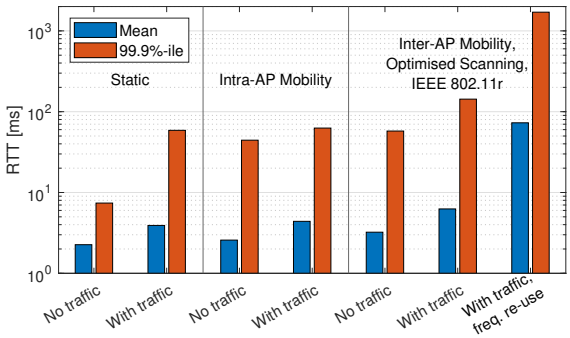

Fig. 6: Mean and 99%-ile RTT latency measurement results for the intra-AP static, intra-AP mobility and inter-AP mobility cases. Dedicated frequency channels were used unless otherwise specified.

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 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

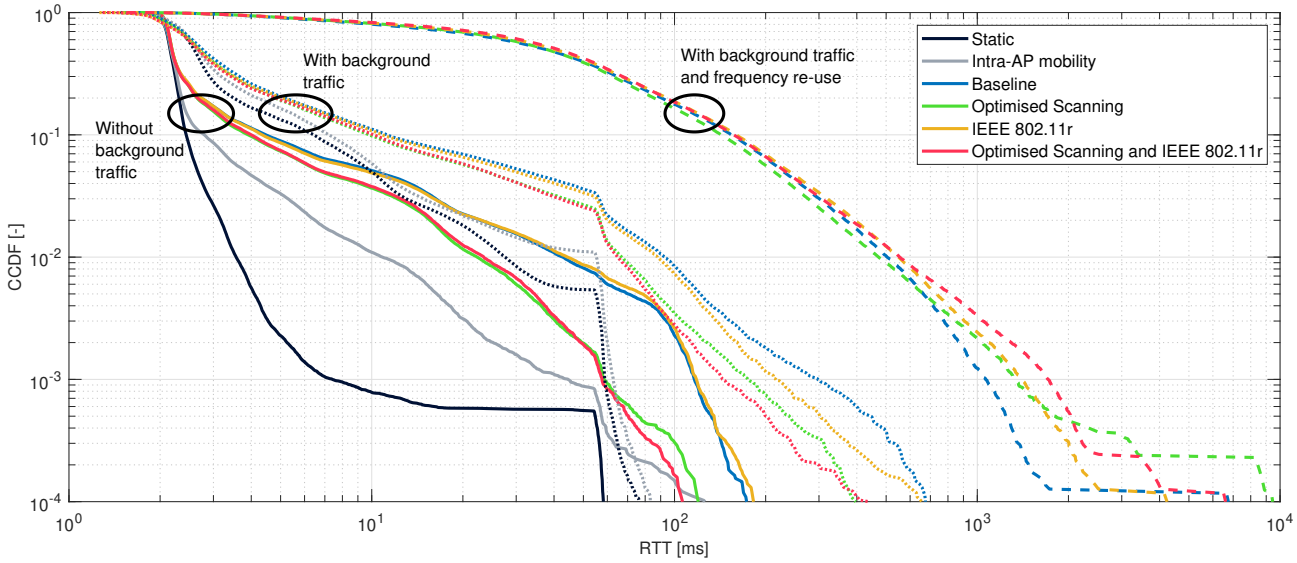

Fig. 7: RTT Empirical CCDFs for all the different Wi-Fi schemes and network configurations explored in the measurements. Dedicated frequency channels were used unless otherwise specified.

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/

2.2 Paper 2

Radio-Aware Multi-Connectivity Solutions based on Layer-4 Scheduling for Wi-Fi in IIoT Scenarios

Andreas Engelsen Fink, Rasmus Suhr Mogensen, Ignacio Rodriguez, Troels Kolding, Anders Karstensen, Guillermo Pocovi

A. E. Fink, R. S. Mogensen, I. Rodriguez and A. Karstensen are with the Wireless Communication Networks Section, Department of Electronic Systems, Aalborg University, Aalborg Øst 9220, Denmark (email: aef16@student.aau.dk; rsm@es.aau.dk; irl@es.aau.dk; andka@es.aau.dk).
T. Kolding and G. Pocovi are with Nokia - Bell Labs, Research Center Aalborg, Aalborg Øst 9220, Denmark (email: troels.kolding@nokia-bell-labs.com, guillermo.pocovi@nokia-bell-labs.com)

The paper has been submitted to the
IEEE Global Communications Conference (GLOBECOM) 2021.

Radio-Aware Multi-Connectivity Solutions based on Layer-4 Scheduling for Wi-Fi in IIoT Scenarios

Andreas Engelsen 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*—Due to mobility and interference, using enterprise Wi-Fi for communication in industrial networks might result in control loop latencies exceeding 100 ms for at least 0.1% of the time, even in those cases where Wi-Fi handover-specific parameters have been optimized, making the technology unfit for Industrial IoT (IIoT) with strict communication reliability requirements. To improve its performance, this paper presents a novel approach towards the design of a radio-aware multi-connectivity concept using a layer-4 scheduling mechanism. Two packet scheduling mechanisms are presented: packet duplication and best path scheduling. Moreover, a mobility coordinator scheme is used to improve the performance of the packet schedulers by preventing simultaneous handovers and ensures the STAs connect to different APs. By using this multi-connectivity solution, a significant performance improvement was observed, cutting down the latencies of the system to 30-80 ms at the 99.9%-ile of reliability (depending on the operational conditions). Furthermore, by applying the proposed schemes, Wi-Fi handovers delays can be fully mitigated allowing for true seamless roaming in mobile conditions.**

I. Introduction

One of the promises of Industry 4.0 is the increased coordination between various types of Industrial IoT (IIoT) equipment to both improve decision making and increase efficiency of the production in general. This functionality comes at the cost of having new and more sophisticated systems with more demanding communication requirements. With e.g. autonomous mobile robots (AMR) and other dynamic equipment utilizing high bandwidth to share data and receive latency-sensitive critical control traffic over a wireless interface, existing wireless technologies will be challenged to support this. With IEEE 802.11 Wi-Fi being commonplace among industrial plants, a large portion of IIoT is designed to utilize this. Previous studies have shown that Wi-Fi technology has challenges in meeting the requirements of latency-sensitive IIoT [1], especially due to the interruptions in communication in case of mobility across Access Points (APs). While the latest iteration of the technology, Wi-Fi 6, contains features that provides better support for critical traffic flows and QoS differentiation, it does not solve the issue with mobility.

Wi-Fi is most commonly deployed in infrastructure mode, with a single Wi-Fi Station (STA) interface on a device connecting to a nearby AP. If the signal between the STA and AP deteriorates, the STA will scan for nearby APs and perform a handover to establish a new connection. This results in a drastic increase of the latency until the new connection is established [1]. With mobility being a key aspect of new IIoT deployments, addressing such roaming events is essential as they might have a significant impact on latency-sensitive applications. Several amendments have been introduced to the IEEE 802.11 standard with the goal of reducing the handovers gaps and allow for seamless roaming, such as IEEE 802.11r which aims at reducing the duration of a handover and IEEE 802.11k which monitor nearby APs improving overall scanning time. However, the mobility latency levels achieved by applying these techniques might not be sufficient to support certain very demanding IIoT applications.

While a solution for this might be to implement a vendor-specific solution modifying the lower layers of the protocol, we consider instead a higher layer multi-connectivity approach [2] that utilizes multiple Wi-Fi STAs on the same device simultaneously (multi-STA configuration). Simply increasing the number of active Wi-Fi STAs per device is not ideal as if they are not properly managed, they might increase the collisions and network load, impacting the overall performance of the system [3]. Nonetheless, if increased reliability is desirable, using two Wi-Fi STAs has been estimated to significantly decrease the Packet Error Rate (PER) [4]. Furthermore, by having multiple STAs available, new possibilities emerge for the choice of which APs to connect to, thus minimizing the impact of handovers in performance. A similar concept was found to be successful in public LTE networks [5].

A common approach for multi-connectivity is to utilize Multipath TCP (MPTCP) [6]–[8], an extension to TCP allowing a single-connection to establish multiple subflows over different paths. However, this comes at a cost of degraded throughput [9] and lack of data multiplexing and prioritization [10]. Another approach is through packet duplication, which can significantly reduce high delays and jitter [11]. Duplicating traffic using different technologies have likewise been previously considered, where [12] and [13] both demonstrate how Wi-Fi and 4G LTE can improve the communication latency by making use of the separate medium access control schemes.

It is clear that utilizing multi-connectivity, either over multiple technologies, or simply using redundancy in a single technology, can deliver improved latency performance which is more suitable for critical IIoT devices. In this paper, we aim at leveraging low-latency mobile multi-connectivity implementations over Wi-Fi that are further enhanced by

considering information from the radio layer such as the signal strength or connection states. We present a novel approach with enterprise off-the-shelf Wi-Fi STAs and APs that utilizes contextual information from the Wi-Fi STAs to manage how the traffic is routed and to control to which APs the STAs are connected to. This paper addresses from the design and implementation and experimental validation of two schedulers and complementary mobility coordinators. To determine the benefits of the proposed solutions, we evaluate their performance in a realistic industrial environment with emphasis on reliable latency (i.e., the latency achieved at certain probabilities such as the 99.9%-ile). The paper is structured as follows: Section II details our Wi-Fi multi-connectivity solutions. Section III introduces the experimental environment, the setups, and the different radio configurations tested for the multiple multi-connectivity schemes. Section IV presents the latency performance measurements results. Section V contains the discussion of the results and highlights areas of potential improvements. Finally, section VI concludes the paper.

II. Radio-aware Scheduling Schemes for Wi-Fi

The proposed approach to multi-connectivity is based on a customized radio-aware layer-4 (transport-layer) packet scheduler to control the traffic flow through two Wi-Fi STAs. This scheduler bases its decisions on radio properties such as connection state and Received Signal Strength Indicator (RSSI) to improve the Wi-Fi system performance. To further enhance the performance, a Mobility Coordinator (MC) is introduced to ensure AP diversity (by preventing the STAs from connecting to the same AP) and to avoid simultaneous handovers. As a reference, the multi-STA components are illustrated in Fig. 1. Two schemes were designed and implemented based on these elements: 1) Packet Duplication (PD), and 2) Best Path Scheduling (BPS). These schemes implement different scheduler configurations which are assisted by the mobility coordinator. Further, the proposed schemes can benefit of making a proper network planning.

A. Wi-Fi Packet Scheduler

Two methods are considered for the transport-layer packet scheduler. The first method is using Packet Duplication, which, according to state-of-the-art, can yield significant performance improvements to communication latency and reliability. This however introduces a significant amount of redundant information to be transmitted over the medium, which in turn increases the average time to access the medium for all devices on the network due to the LBT channel access mechanism. With a goodput of less than 50% when also considering packet headers, this can severely harm the overall throughput, especially if multiple devices utilize this method. This issue will be mitigated with the assistance of the mobility coordinator.

In the second method, referred to as Best Path Scheduling, the device uses single-connectivity while taking advantage of the presence of a secondary STA. In contrast to MPTCP, where a shortest-RTT scheduler is used to determine the best path, our approach uses the RSSI of the STAs to steer the traffic. When either the primary STA being used for traffic is disconnected from its AP or the RSSI of the secondary STA exceeds the primary by a margin of 5 dB, the traffic is steered through the secondary STA. By allowing for seamless transition of the packet flow between the two STAs during runtime, the impact of e.g. handovers or abrupt disconnections can be mitigated. The margin of 5 dB is set to avoid excessive switches between the STAs if the RSSI is similar. The performance of this scheme is further enhanced by the logic of the mobility coordinator.

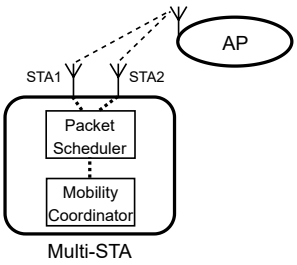

Fig. 1: Overview of multi-STA components: Primary STA, secondary STA, packet scheduler and mobility coordinator.

B. Wi-Fi Mobility Coordinator

When the connection between a STA and an AP is degraded by a significant amount (measured either through the RSSI or by detecting a connection loss), the STA will scan for other eligible APs nearby and then roam to the one with highest RSSI. If no effort is put into coordinating the two STAs and the locations of their antennas are close, they will experience very similar channel conditions and may choose to scan and roam between APs simultaneously. This will hurt the overall performance of either of the presented packet scheduler configurations, and is thus of interest to improve.

Both STAs contain a list of eligible Basic Service Set IDs (BSSIDs) which are used during network scans to choose the serving AP. To introduce coordination between the two STAs, a blacklist is maintained by the mobility coordinator to prevent both STAs from connecting to the same AP. The coordinator will furthermore periodically check if the STAs need to roam to a new AP, with a long enough periodicity (2.5 s in our case) to allow for a full scan and potential handover to another AP. In our implementation, roaming events will be triggered based on a RSSI threshold of -85 dBm. When the threshold is reached for a STA, the BSSID of the previous AP is added to the blacklist temporarily to force the disconnection.

The two algorithms for the combined packet scheduler and mobility coordinators are illustrated in Fig. 2 and 3 for the PD and BPS packet schedulers, respectively.

The objective of the mobility coordinator differs slightly depending on the scheduler. For the PD scheme, the coordinator prioritizes uptime on both interfaces while keeping track of association and disassociations for each STA to maintain the blacklist. For the BPS algorithm, the coordinator will only initiate roaming events for the secondary (i.e. idle) interface.

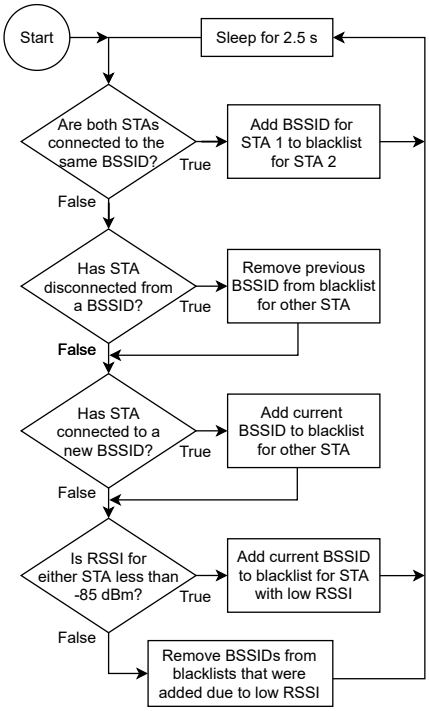

Fig. 2: Flowchart for the packet duplication scheme.

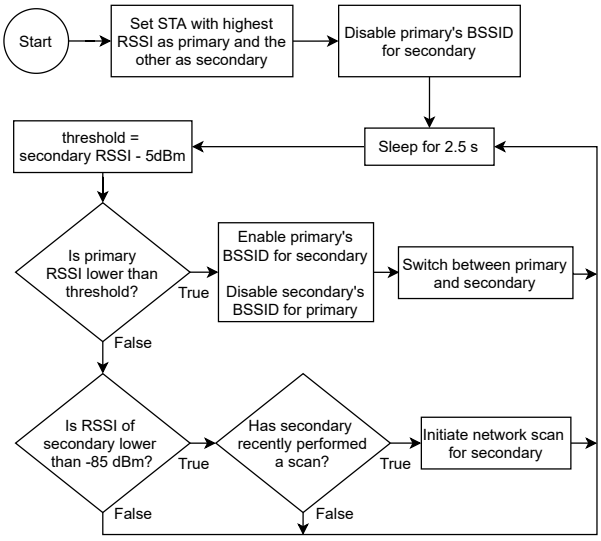

Fig. 3: Flowchart for the best path scheduling scheme.

C. Wi-Fi Network Planning

In order to fully utilize the mobility coordinator, deploying APs such that there are overlaps in coverage areas is necessary. Because of the LBT mechanisms of Wi-Fi, the average time until the medium can be accessed will be highly dependent on the number of active devices. The optimal performance of the mobility coordinator under the PD scheme will be achieved when two overlapping APs utilize different frequency channels. Under other spectrum configuration circumstances, both STAs might experience the same average time until medium access.

III. Experimental Test Setup

The performance evaluation of the designed and implemented Wi-Fi multi-connectivity schemes was performed at the AAU 5G Smart Production Lab at Aalborg University, Denmark [14]. This industrial environment (shown in Fig. 4) is equipped with three ceiling-mounted CISCO MR36 Enterprise Wi-Fi 6 APs [15] deployed throughout the lab as illustrated in Fig. 5. In order to trigger the mobility aspects of our Wi-Fi solution performance evaluation, a MiR200 AMR (also shown in Fig. 4) was used. The AMR was configured to follow a specific route through the lab as illustrated in Fig. 5, carrying the implemented multi-STA device around at a maximum speed of 1.5 m/s.

The multi-STA was configured using wpa_supplicant v. 2.9 [16] which is used to communicate with the driver for the Wi-Fi STAs and is furthermore used to handle roaming and key negotiation. The number of frequency channels which the STAs scan was optimized to match the number of APs in the testing environment (three). To enable the multi-connectivity aspect of the setup, an improved version of the multi-access gateway presented in [12] was used. In uplink, the gateway encapsulates traffic from an end-device (i.e. mobile robot) and transmits it through specified interfaces, in this case two Intel Wi-Fi 6 AX200 network cards, to another gateway on the network side from which the traffic is decapsulated and transmitted to another end-device (i.e. network server). The procedure is similar in downlink but with encapsulation happening at network-side and decapsulation happening at the device-side. This is illustrated in Fig. 6. During transmissions between gateways, an additional 44 bytes of headers and metadata from the encapsulation are added per frame, excluding technology-specific headers. The use of the gateways in this setup will furthermore introduce a calibrated delay of ∼0.12 ms per frame from the two Ethernet transmissions and processing in the gateway.

The multi-connectivity performance was evaluated between two end-devices (a NUC PC mounted on the AMR at the device-side and a server at the network-side) using the Linux ping functionality, providing an insight on round-trip time (RTT) and packet error rate (PER) statistics. A packet size

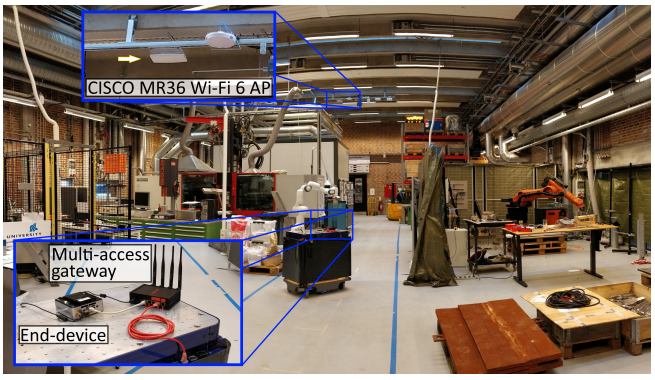

Fig. 4: Overview of the industrial environment and the multi-STA measurement setup, including one of the ceiling-mounted APs and the AMR used for mobility.

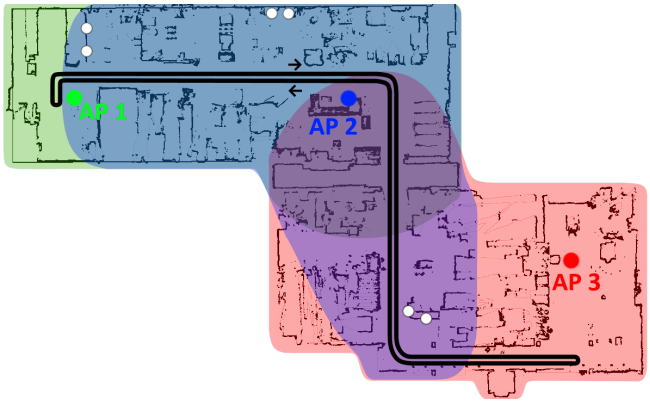

Fig. 5: Floor plan of the test environment, including AP locations (green, blue, and red dots) and coverage areas, the AMR measurement route and the location of the static traffic sources for the generation of background traffic (white dots).

of 64 B and an inter-packet interval of 50 ms was used. The packet interval was chosen to obtain sufficient samples while capturing the impact of handovers on the performance, for which typical AMR traffic models are unsuitable. Simultaneously, RSSI was monitored. The following four configurations were examined in the experimental testing:

1) BPS with mobility coordination over dedicated channels.
2) PD without mobility coordination over dedicated channels.
3) PD with mobility coordination over dedicated channels.
4) PD with mobility coordination and frequency re-use.

Further, all configurations were tested in idle networks with no other traffic than the one generated by the multi-STA device, and also with background traffic where two static STAs were used to load each AP with a constant controlled traffic load of 10 Mbit/s uplink and 10 Mbit/s downlink. The traffic load was chosen to reflect a low-medium usage of the network and to observe an impact on the latency without reaching congestive conditions. For configurations with BPS scheme, the gateway on the network-side was configured to steer traffic to the STA that it had last received data from to ensure single-connectivity behavior in both uplink and downlink. In those cases, where the mobility coordinator was not used (the PD case without mobility coordination over dedicated channels),

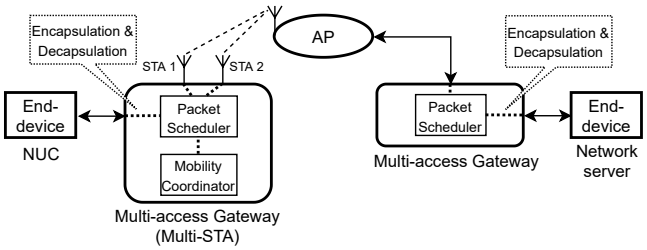

Fig. 6: Test setup using multi-access gateways. Solid connections represent Ethernet connections, while dashed connections represent Wi-Fi.

the STAs were configured to initiate network scans when they reach -85 dBm RSSI, from which they will search for another suitable AP nearby and initiate the handover. This is done to reduce the stickiness of the connection to an AP, as it would otherwise remain connected until the connection is lost (at $\sim$ -93 dBm).

IV. Wi-Fi Performance Results

We first take a look to the performance of the PD scheme with and without mobility coordination. The RSSI measured at each STA is illustrated in Fig. 7 for both schemes. The data confirms that if the STAs are used without mobility coordination, the RSSI will, as expected, be very similar due to the low spatial diversity of the setup. Spatial diversity in the setup could be improved by separating as much as possible the antennas of the different STAs, but that would not be realistic as, in practice, industrial hardware impose restrictive constrains on the communication modules and antenna location. With the mobility coordinator enabled, the RSSI-traces for the different STAs became uncorrelated due to each STA connecting to a different AP. It is, however, also observed that when the signal strength degrades to -85 dBm (corresponding to the interval between 55-85 s) and if no eligible AP can be reached by the secondary STA (this happens when the AMR is only in coverage with AP 3 in Fig. 5), the STA will fully disconnect and either remain in a searching state until a suitable AP is found, or it will ping-pong between reconnecting to the previous AP and disconnecting due to low RSSI. When considering the RTT latency performance, the results illustrate

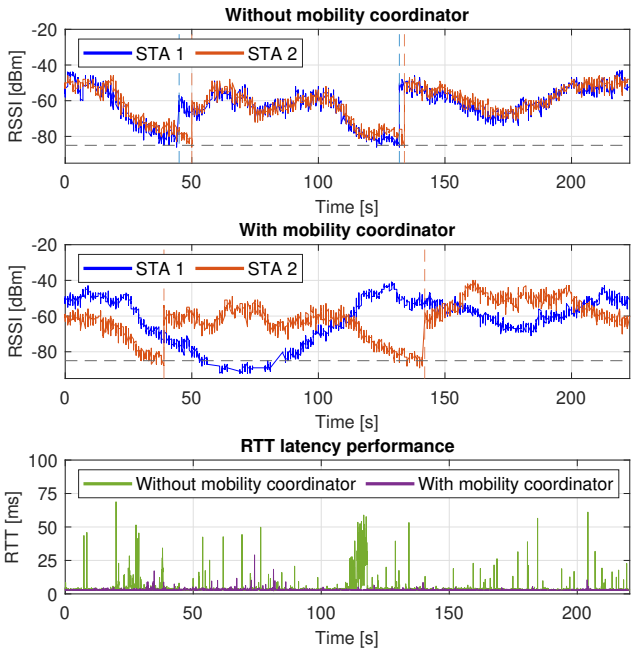

Fig. 7: RSSI traces and overall Wi-Fi RTT latencies for two different PD schemes: with and without mobile coordinator. The -85 dBm threshold is highlighted with the horizontal line, and handovers are highlighted with vertical lines.

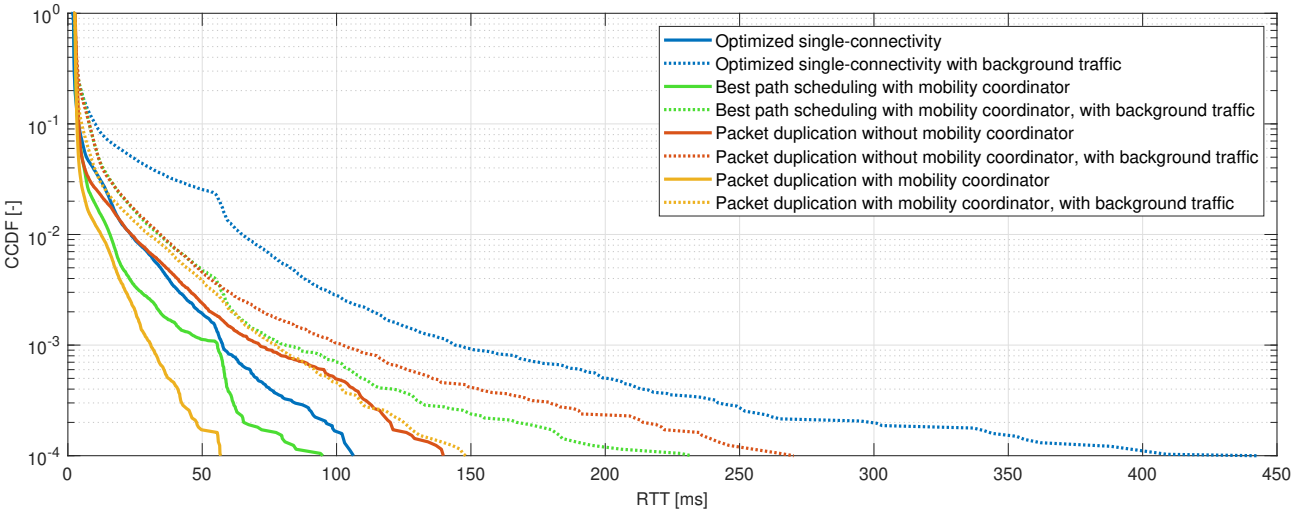

Fig. 8: Empirical CCDFs of the RTT for single- and multi-connectivity configurations operating over dedicated channels.

TABLE I: Summary of RTT latency statistics and PER measurement results for the single- and multi-connectivity configurations operating over dedicated channels.

Test Configuration	Coordination scheme	Channel condition	Min	Avg	99.9%-ile	Jitter	PER
Optimized single-connectivity [1]	Without mobility coordinator	Idle network	1.48 ms	3.23 ms	57.60 ms	1.56 ms	0.044%
		Background traffic	1.29 ms	6.26 ms	143.00 ms	4.78 ms	0.065%
Best path scheduling	With mobility coordinator	Idle network	2.23 ms	3.36 ms	55.60 ms	1.00 ms	0.005%
		Background traffic	2.24 ms	4.80 ms	80.70 ms	2.94 ms	0.028%
Packet duplication	Without mobility coordinator	Idle network	2.32 ms	3.88 ms	71.10 ms	1.42 ms	0.018%
		Background traffic	1.99 ms	4.87 ms	102.00 ms	2.66 ms	0.041%
	With mobility coordinator	Idle network	2.14 ms	3.15 ms	30.80 ms	0.58 ms	0%
		Background traffic	2.06 ms	4.12 ms	77.10 ms	1.75 ms	0.001%

that using the mobility coordinator translates into significantly higher stability (reduced amount of latency spikes, and spikes of shorter duration) than for the uncoordinated configuration.

Fig. 8 displays the empirical complementary cumulative distribution functions (CCDF). Each of the presented results sets were computed from more than 100,000 RTT latency samples of each tested configuration. Key latency statistics and PER are summarized in Table I. As a reference in both Fig. 8 and Table I, results from the optimized single-connectivity mobility performance test presented in [1] are also included. The results indicate that by having two STAs available in a single device, using either BPS or PD with mobility coordinator, significant latency improvements can be achieved for the 95%-ile and above, as compared to the single-connectivity reference. While taking the RSSI of the STAs into account by using the mobility coordinator will result in a more robust connection, the main improvement stems from the multi-STA being able to fully mitigate the latency impact of handovers. Apart from in latency, this has also a positive effect on the PER, that is reduced significantly in BPS, and almost completely in PD. In the cases where background traffic was present, the increase in latency is significantly lower in BPS and PD as compared to the single-connectivity configuration. In the case of PD multi-connectivity without mobility coordination, the performance is very similar to the one for single-connectivity for the idle network case, while a gain is observed for the case with background traffic. At 99.9%-iles, the performance of the BPS and PD schemes with mobility coordination and idle network can be as low as 56 and 31 ms, respectively, as compared to 58 ms in the single-connectivity case (4-46% gains). When background traffic is present, the gains for the multi-connectivity schemes are even larger (43-46%).

All previous results assumed some level of network planning and were obtained using individual dedicated frequency channels for each of the APs. However, the impact of uncoordinated deployments where the APs operate under frequency re-use was also evaluated. This evaluation was done for the PD scheme and the results are shown in Fig. 9. The results indicate that if dedicated frequency channels cannot be guaranteed for each AP to avoid overlapping regions, the performance of the PD multi-connectivity scheme will be slightly degraded in the presence of background traffic.

V. Discussion

Using mobility coordination will result in increased performance when multiple APs on dedicated channels are available, with PD performing better than BPS both in terms of latency and PER. However, as PD causes an increase of the overall load of the network, thus causing interference for other de-

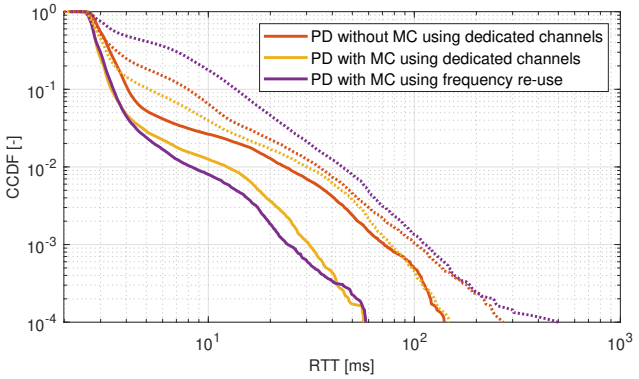

Fig. 9: Empirical CCDFs of the RTT for the packet duplication multi-connectivity scheme with dedicated channels and with frequency re-use.

vices, BPS could be a more appealing scheme to be applied for multi-connectivity. This will depend on both the network setup and amount of active Wi-Fi devices, but the performance may differ if a different type of traffic load is applied, e.g. exponential traffic patterns as compared to a constant load.

While the presented algorithms do improve the latency and PER performance, some areas for potential improvement have been identified. In regions where only one the primary STA remains operational due to impossibility of the secondary STA to find a suitable AP (such as in the PD scheme with mobility coordination presented in Fig. 7), it would be beneficial to force the secondary STA to connect to the same AP as the primary STA has, despite of having correlated RSSI, a performance gain could be achieved, especially in the presence of background traffic. As an alternative, fine tuning of the RSSI thresholds for the mobility coordinator could also result in an improved performance. If end-device location information is available from an external positioning system, the scheduling algorithms and mobility coordinator could be enriched by including this information in their decisions, eliminating the need for scanning for channels and neighbor APs.

VI. Conclusion

In this paper, a novel approach to introducing multi-connectivity using off-the-shelf Wi-Fi hardware configurations is presented. We present a custom transport-layer packet scheduler located at the edges, therefore not requiring any changes to the network itself, as opposed to e.g. proprietary solutions. This approach uses knowledge of the Wi-Fi connection (e.g. connection state and RSSI) to improve both how the traffic is steered and to introduce a mobility coordinator between the multiple STAs in the device, such that they connect to different APs and avoid simultaneous handovers. Two schemes are evaluated: 1) best path scheduling utilizing a primary STA with highest RSSI and seamlessly switching to the secondary STA when the signal degrades, and 2) packet duplication over the two STAs simultaneously.

The experimental performance evaluation showed that both multi-connectivity schemes improve the Wi-Fi performance as compared to the single-connectivity case. Latency improvements of up to 46% were observed at the 99.9-%iles, lowering the latency from 143 ms to 77-81 ms in the presence of background traffic. Using the mobility coordinator ensures correlated links for the different STAs, which translates into non-simultaneous handovers thus, fully mitigating the impact of mobility between different APs in the communication, resulting in a seamless roaming operation. Packet duplication was found to be the best performing scheme, but it comes at the operational cost of having increased load in the system, and thus the operational conditions should be carefully analyzed before prioritizing it over the best path scheduling scheme.

References

[1] A. Fink *et al.*, "Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility," unpublished, submitted to European Wireless 2021.

[2] E. J. Khatib *et al.*, "Multi-Connectivity for Ultra-Reliable Communication in Industrial Scenarios," in *IEEE Vehicular Technology Conference (VTC2019-Spring)*, 2019, pp. 1–6.

[3] M.-T. Suer *et al.*, "Multi-connectivity as an enabler for reliable low latency communications—an overview," *IEEE Communications Surveys Tutorials*, vol. 22, no. 1, pp. 156–169, 2020.

[4] J. J. Nielsen *et al.*, "Reliability and Error Burst Length Analysis of Wireless Multi-Connectivity," in *International Symposium on Wireless Communication Systems (ISWCS)*, 2019, pp. 107–111.

[5] M. Lauridsen *et al.*, "Reducing Handover Outage for Autonomous Vehicles with LTE Hybrid Access," in *IEEE International Conference on Communications (ICC)*, 2018, pp. 1–6.

[6] C. Raiciu *et al.*, "How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP," in *USENIX Conference on Networked Systems Design and Implementation (NSDI)*, 2012.

[7] A. Croitoru *et al.*, "Towards Wifi Mobility without Fast Handover," in *USENIX Conference on Networked Systems Design and Implementation (NSDI)*. USENIX Association, May 2015, pp. 219–234.

[8] K. Nguyen *et al.*, "Improving WiFi networking with concurrent connections and multipath TCP," in *IEEE International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*, 2013, pp. 1–3.

[9] M. R. Palash *et al.*, "MPWiFi: Synergizing MPTCP Based Simultaneous Multipath Access and WiFi Network Performance," *IEEE Transactions on Mobile Computing*, vol. 19, no. 1, pp. 142–158, 2020.

[10] R. S. Mogensen *et al.*, "Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications," in *IEEE Vehicular Technology Conference (VTC2019-Spring)*, 2019, pp. 1–5.

[11] H. Zhang *et al.*, "WiFi and Multiple Interfaces: Adequate for Virtual Reality?" in *IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, 2018, pp. 220–227.

[12] R. S. Mogensen *et al.*, "Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0," in *IEEE Vehicular Technology Conference (VTC2019-Fall)*, 2019, pp. 1–7.

[13] M. Suer *et al.*, "Evaluation of Multi-Connectivity Schemes for URLLC Traffic over WiFi and LTE," in *IEEE Wireless Communications and Networking Conference (WCNC)*, 2020, pp. 1–7.

[14] 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.

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

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

Chapter 3

Worksheets

3.1 IEEE 802.11 Wi-Fi

This section aims to investigate the IEEE 802.11 Wi-Fi standard. Initially, the background of the standard will be presented, followed by a general description of the main features of Wi-Fi. The performance of the standard will then be considered when introducing mobility, such as switching between networks. Finally, selected standards and approaches for optimizing the performance will be presented.

3.1.1 Background

Wi-Fi has for several years remained a popular choice for wireless communication in many scenarios, both in private and industrial use-cases. Because it is an open standard and not proprietary software, equipment for it is widely available and easy to incorporate into devices, thus making it even more widespread.

The IEEE 802.11 Wi-Fi protocol has been used since its release in 1997 and has since then received numerous updates and revisions as technology has become more sophisticated, taking advantage of e.g. Orthogonal Frequency-Division Multiplexing (OFDM), as well as multiple antennas to perform beamforming and MIMO communication [1]. With these revisions, both the transfer speed and overall reliability has improved significantly. An overview of the main changes between the Wi-Fi versions can be found in Table 3.1.

Table 3.1: Brief overview of the difference between Wi-Fi versions [1].

Wi-Fi 1	Initial release, operating on 2.4 GHz frequency, using DSS and CCK modulation schemes for data.
Wi-Fi 2	Introduces OFDM to support high data rates unlike the single-carrier case of Wi-Fi 1. Operates on the 5 GHz spectrum.
Wi-Fi 3	Combines Wi-Fi 1 and 2 to allow Access Points (APs) and Stations (STAs) to operate on both 2.4 GHz and 5 GHz bands and use both versions' modulation schemes.
Wi-Fi 4	Introduces MIMO and initial versions of beamforming. Allows for using both 20 MHz and 40 MHz bandwidth. Uses modulation schemes such as BPSK, QPSK, 16QAM and 64QAM.
Wi-Fi 5	Beamforming added to the MIMO and multi-user MIMO feature is introduced.
Wi-Fi 6	Introduces OFDMA for both uplink and downlink. Also includes MU-MIMO, beamforming, 1024 QAM and more.

3.1.2 Overview

When it was initially implemented, Wi-Fi was intended to replace the cabled Ethernet connections in a seamless manner, with the most common usage being for local area networking and Internet access. There are several configurations available, with the main two being infrastructure mode and peer-to-peer mode.

- The infrastructure mode operates as a star-topology network, with an AP being the central point of communication. Devices can connect to this AP wirelessly, and access other devices in the network by routing traffic through the AP [2]. Devices can also access other networks (e.g. the Internet) through a wired connection from the AP.
- For the peer-to-peer mode, also known as ad-hoc networks and Wi-Fi Direct, devices communicate directly to each other or follow some general protocol [3]. This type of communication can be useful for cases where no centralized networks are available.

Wi-Fi operates in the unlicensed ISM band and is restricted to operate in designated channels and frequency ranges, with the most common being located in the 2.4 GHz and 5 GHz spectrum. Although this significantly reduces the costs of setting up a network (compared to e.g. 4G LTE) as it is not required to purchase a license, a significant concern for utilizing these bands is the presence of other networks or devices operating on the same channel. This can be other Wi-Fi networks, but also other technologies such as Bluetooth [4]. There are likewise rules and guidelines that must be taken into account to

ensure fair use, such as maximum transmission power and channel bandwidths. Because the wavebands have relatively high absorption and work best in line-of-sight conditions, the range of a network can be significantly impacted by obstacles such as walls or doors [5]. It may therefore be necessary to deploy multiple networks to guarantee sufficient coverage. A benefit of the high absorption is however that it can be easier to separate the networks in a spatial manner.

The data transferred over Wi-Fi is encoded using various levels of quadrature amplitude modulation (QAM). Higher levels of QAM results in higher throughput but is more susceptible to noise. This is one of the areas in which Wi-Fi has improved over the years, with Wi-Fi 5 supporting 256 QAM [1].

Because all devices in the network operate on the same frequency, and approaches such as collision detection from Ethernet isn't feasible in wireless communication, the Carrier-Sense Multiple Access with Collision Avoidance (CSMA/CA) protocol is used instead [6]. This protocol can generally be considered to be a "listen-before-talk" approach to medium access control, where each device listens for current transmissions. If the channel is idle, it will start transmitting immediately. If another device is currently transmitting, it will wait for a random period of time before checking again. This randomness is necessary in case two or more devices waiting for the channel to become idle, as they would otherwise both start transmitting simultaneously, resulting in collisions. The access aspect can be further enhanced by utilizing the IEEE 802.11 RTS/CTS exchange, which aims to solve the hidden/exposed node problems.

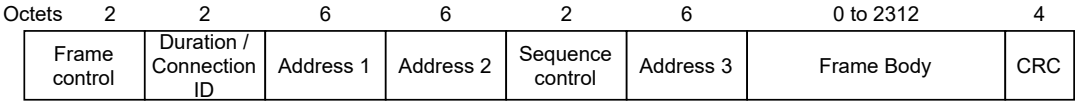

Figure 3.1: Wi-Fi frame format [7].

The link-layer header of the protocol is illustrated in Figure 3.1 [7]. The frame control field is used to indicate the type of frame, such as uplink/downlink, control, data etc. The duration field indicates the time the channel is allocated for a successful transmission, but the field may also be used as an identifier for some control frames. The purpose of the addresses can vary depending on the context, but will in most cases consist of the source/destination MAC address and the address of the network. The sequence control field is used for fragmentation and reassembly as well as to keep track of the number of frames sent between the STA and the AP. Finally, the CRC is used for error checking.

3.1.3 Mobility in Wi-Fi

While static deployments of Wi-Fi STAs can experience changes in the radio conditions, e.g. a shift in signal quality, mobile Wi-Fi STAs experience this in a much more significant manner. How much these changes will depend on the environment, the speed at which the STA roams and distance between STA and AP. This property also extends to other radio communication technologies, but as the focus of the project is specifically on Wi-Fi, this is will not be further investigated. The impact on the latency of static deployments and intra-AP and inter-AP mobility has however been investigated experimentally, where it was found that the latency does indeed increase as a consequence of this. More information can be found in Section 3.4.1.

When the connection between the STA and the AP degrades significantly, or if the connection is lost completely, the STA will need to roam to a new AP. This is known as a handover, and can be divided into 4 main stages: scanning, authentication, association and handshake. These stages are explained below. The impact of each stage has been found experimentally in Section 3.4.2.

Scanning
The device will initially need to determine which APs are available. Each AP will periodically send a beacon frame on the channel at which its network is located. This beacon is meant for other devices to locate it, and contains information such as [8]:

- The timestamp at which the beacon was sent, in order to improve synchronization.
- Its beacon interval (typically 102.4 ms [9]).
- Its capabilities, such as the type of network (ad-hoc or infrastructure) or the encryption type.
- Service Set ID (SSID).
- Supported data rates.

A device can scan for networks through either a passive or active approach. For passive scanning, the device simply listens for beacons from APs. Since the device will only be able to listen on a single channel at a time, it will need to sequentially listen on all available channels, unless a subset of channels has been selected on the device-side. When utilizing active scanning, the device will instead send probe requests on a channel, and then receive responses from APs on that specific channel. This allows for significantly faster scanning, as the time spent on each channel is reduced. Once the scanning process is finished, the device will select an AP to connect to. This is entirely dependent on the setup of the device, but may include factors such as the signal strength, the capabilities of an AP, or by utilizing a prioritized list of SSIDs. The selection algorithm of the devices used for this project is shown in Section 3.2.2.

Authentication
Once an AP has been selected by the network device, it will initiate the connection through authentication. The goal of this step is to establish its identity with the AP [10]. There are two link-level types of authentications defined in the IEEE 802.11 standard: Open System and Shared Key.

The open system authentication consists of two messages, with the first being an authentication request sent from the network device, which typically contains information such as its MAC address. The AP will then respond with either a success or failure message.

A shared key authentication is used in order to limit the network to a subset of authorized users. The shared key, or passphrase, is manually set on the network device and the AP. There are several types of authentication protocols available:

- Wired Equivalent Privacy (WEP)
- Wi-Fi Protected Access (WPA)
- Wi-Fi Protected Access 2 (WPA2)

Since WEP is not recommended for normal use (as its security can easily be bypassed) WPA and WPA2 are the most common. WPA uses IEEE 802.1X authentication and key-exchange, in which dynamic encryption keys are being used. WPA2 works similar to WPA, but is instead considered to be a further security enhancement and occurs during the handshake stage.

Association
Once a device has been successfully authenticated, it can begin its association with the AP. This will register the device to the AP in order to gain access to the network. This is necessary in order for the AP to keep track of which devices are currently connected and avoid e.g. routing packets to an absent device. The device is likewise limited to only be associated with a single AP. This ensures that the device is always 'present' and not e.g. communicating to an AP on a different channel.

The association process starts with the network device sending an association request, which contains the chosen encryption types and other compatible IEEE 802.11 capabilities. The AP will then either grant association or deny it based on vendor-specific implementation.

Handshake
The handshake stage improves the security of the authentication stage through the IEEE 802.11i amendment, implemented as Wi-Fi Protected Access 2 [11]. Here, handshake messages are exchanged to set up and verify the encryption between the AP and STA.

3.1.4 IEEE 802.11 amendments to Improve Roaming

Several amendments have been introduced to the IEEE 802.11 standard with the aim mitigating the impact of handovers and allow for seamless roaming.

IEEE 802.11k

This standard aims to reduce the number of channels which are scanned by creating an optimized list [12]. When the signal strength of the connection to the current AP weakens, the device will start scanning from this list.

This is done through requests sent by the device to nearby APs, from which it compiles a list of candidates to be used when roaming. The request itself is an IEEE 802.11 management frame, known as an action frame. An AP responds to this action frame with another action frame containing a list of neighbor APs on the same WLAN with their Wi-Fi channel numbers. This allows the device to make a decision based on this list instead of performing an active or passive scan of all channels, significantly reducing the time required during a handover.

The neighbor list is created on-demand, and can thus contain different APs depending on which AP the device is connected to [13]. The list can contain APs on the same channel, which is optimal for battery life (as the STA does not have to perform a scan), but can also include APs from other channels. The latter can be more useful to avoid overlapping in terms of frequency for nearby networks.

While not a part of the standard, the Cisco MR36 Access Points used throughout the project provides similar functionality for non-802.11k clients. A prediction neighbor list can be generated for each client, which is based on previous probe requests and RSSI values, and is used to predict which AP the device will most likely roam to. If the device attempts to connect to an AP which is deemed undesirable by the network, the AP can stop the process by sending a deuath-frame. The device will then attempt to connect to the next AP, which may be more reliable based on previous data.

IEEE 802.11r

When switching between networks, the IEEE 802.11r standard enables the use of a feature called Fast Basic Service Set Transition (FT) to make the authentication step faster [14]. This works with both pre-shared keys (PSK) and 802.1X authentication methods.

This form of roaming is made possible by completing the initial handshake with a new

AP before the client needs to connect to it. Because of this, the calculation of the Pairwise Transient Key (PTK) can be done in advance, which are used for communication after the reassociation request or response with the new target AP.

There are two ways for a device to roam networks: Over-the-Air and Over-the-DS, illustrated in Figure 3.2. In the first method, the device communicates directly with the target AP using FT alongside the regular authentication steps. In the latter, the device communicates with the target AP through the AP it is currently connected to, thus communicating through the cabled network between APs instead. This however requires the APs to be connected to the same or compatible controllers.

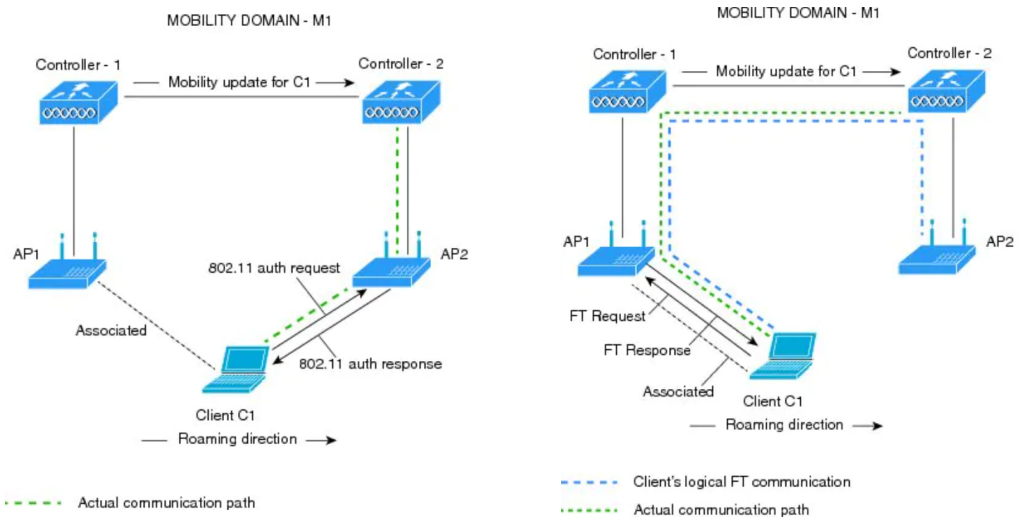

(a) Over-the-Air **(b)** Over-the-DS

Figure 3.2: Client roaming using IEEE 802.11r [15].

IEEE 802.11v

The IEEE 802.11v standard contains enhancements for wireless network management. This includes features regarding roaming, where APs can direct associated clients to other APs. This can be done in a solicited manner, where clients send queries before roaming to obtain a better option of AP to re-associate, similar to the principle in IEEE 802.11k. It can also be done in an unsolicited manner, where the AP may attempt to move the device to another AP to perform load balancing or because of a poor connection quality to the current AP. These are however suggestions and not strict commands, meaning the device can choose to ignore the messages altogether. An AP would otherwise choose to simply disconnect the device and force it to look for other networks, which in turn will harm the performance of applications on the device utilizing the network. This can be due to load balancing, but also in case the device may have a stronger connection if it switches.

3.2 Measurement Software Tool and Single-connectivity Communication Setup

This section will present the software for the data logger used in experiments. This tool allows for gathering a wide range of data, from technology-specific properties (e.g. Received Signal Strength Indicator (RSSI)) to traffic-specific statistics (Round-Trip Time (RTT), Packet Error Rate (PER)). An initial overview of the design of the measurement software is presented which details its purpose and discusses which data should be measured, followed by the implementation related to each data type. This section will likewise address configurations related to the tests, such as the use of an AMR to enable mobility. Finally, this section will present the approach to enabling some of the previously discussed Wi-Fi-specific features and IEEE 802.11 amendments used for tests. A test journal for the tests for Paper 1 can be found in Section 3.4.3.

The platform of choice for this part of the project will be the Intel NUC5i3MYBE detailed in Table 3.2, equipped with the Intel Wi-Fi 6 AX200 network card as opposed to its default Wi-Fi 5 variant. While basic connectivity and setup works out-of-the-box, some additional effort to enable the amendments and solve issues contributing to the perceived latency is required, and is detailed in section 3.2.3.

Table 3.2: Computer hardware and software setup for physical tests.

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

3.2.1 Measurement Software

The measurement software will be implemented using C++ on a Linux-based operating system, specifically Ubuntu 20.04.1 for this project. The purpose of the software is to be compatible on multiple devices. The reasoning behind the choice of operating system is due to availability and because Linux contains software tools to enable interaction with technology-specific interfaces beyond simple communication.

It is initially of interest to gather data for the specific communication technology being utilized. As previously described, this project investigates the performance of Wi-Fi, therefore properties of this technology will be included. Secondly, while the use of cellular technologies (4G LTE, 5G NR) was in the scope of the project initially, but

was later changed to focus solely on Wi-Fi, tools to gather data for these technologies were developed and will likewise be included in this section. Thirdly, capturing the traffic being transmitted and received allows for further processing at a later time to determine traffic patterns and investigate how specific mechanisms of a technology is carried out, such as the handover process of Wi-Fi. Fourth, because of the mobility aspect of the project, keeping track of the location of the STA during operations will allow for determining the physical locations of events such as handovers or areas with high packet loss. Finally, to evaluate the performance of a given solution, the latency between the STA and another device on the network can be utilized.

The measurement software is designed based on the principle of plugin components as illustrated in Figure 3.3, which can be enabled or disabled as necessary, with a realization of each plugin being available for each present interface, such as two separate network cards. One of the main objectives of the measurement software is to allow for gathering of data for different aspects simultaneously on the same device, as to allow for reliable comparison between samples. If measurements are done at e.g. multiple devices at once, and then has to be compared afterwards, additional effort in form of time synchronization will be necessary to ensure sufficient accuracy. Because one of the main focus areas of this project is mobility, this synchronization will become less reliable, as a wireless medium has to be used for synchronization, or otherwise a dedicated solution needs to be used. A centralized trigger will be used to initiate measurements for the majority of the plugins, as to minimize the difference in time between e.g. a Wi-Fi measurement and the position. Since this will trigger the beginning of each sample and each plugin does finish sampling simultaneously, there will however be a slight drift, including those plugins which may require more time than the periodic trigger. For these cases, the next trigger will initiate the generation of a new sample. Two plugins will not be utilizing this trigger: 1) the traffic logger, which operates continuously as data is gathered, and 2) the latency logger which uses an internal trigger mechanism.

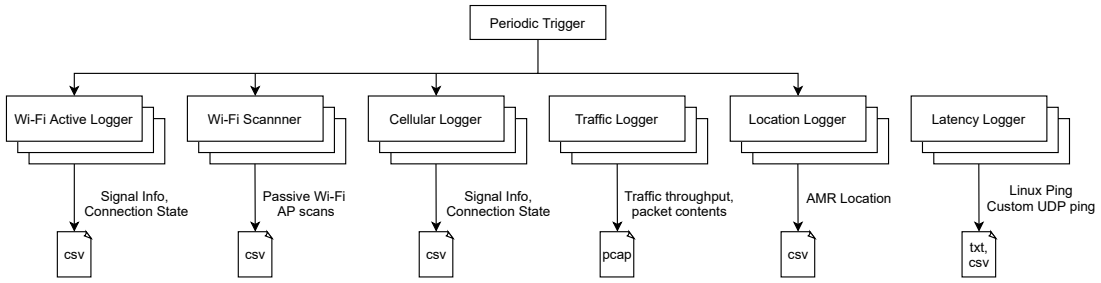

Figure 3.3: Overview of measurement software components.

Wi-Fi Loggers

As Linux is used as platform for the logger, several tools for managing the STA exists and can be utilized. An overview of network-related software and hardware is shown in Table 3.3.

Table 3.3: Overview of selected Wi-Fi-related software on the Linux platform.

Layer	Description
Network Manager	The Network Manager software aims to keep an active network connection available at all times [16], by managing interfaces such as Ethernet and Wi-Fi. Also handles tasks such as IP configuration and using DHCP requests to obtain an address, and handles routes and nameservers. Uses wpa_supplicant to connect to AP's, perform background scans to maintain connectivity etc.
wpa_supplicant	Provides methods for computing encryption keys, handles calls to lower layers to connect to access points and similar aspects [17]. Primarily uses the nl80211 driver to execute commands for the network card, but can also be used with the older wext driver.
802.11-specific drivers	Generic drivers such as nl80211 and cfg80211 to communicate with the network card.
iw	Another command-line interface similar to wpa_supplicant which handles Wi-Fi connections [18].
Intel Driver	Enables communication between the 802.11-specific drivers and the network card.
Intel Wi-Fi6 AX200 Network Card	The physical network card itself.

The wpa_supplicant software was chosen as the target layer, as it provided the functionality of connecting to access points, monitoring status and as well as other aspects. It is also more widely adopted and has more documentation as compared to iw. Because the Network Manager, which was pre-installed on the NUC, contained features such as periodic background scans and would overrule commands made from wpa_supplicant (e.g. sudden roaming events to another AP), it was disabled to maintain consistent behavior.

wpa_supplicant uses an initial configuration when starting to specify e.g. region and information about eligible networks. A minimum example of this configuration is shown in Listing 3.1. Further initial settings can be configured as well, as described in [19].

```
ctrl_interface=DIR=/var/run/wpa_supplicant GROUP=netdev
update_config=1
country=DK

network={
        ssid="SSID"
        psk="PASSKEY"
}
```

Listing 3.1: wpa_supplicant minimal example. The highlighted text is replaced by the corresponding information for the network.

To communicate with the wpa_supplicant from the measurement software, it must attach to a control interface provided by the supplicant. This is done through an external C-library, wpa_ctrl [20]. This allows for issuing commands to the supplicant and control its functionality similar to the command-line interface. Information is returned through either requests or through a callback notifier functionality. Furthermore, the wpa_supplicant software supports logging all events to a file with associated timestamps. While this would normally be used for debugging, this allows for obtaining an insight into the actions of the STA in terms of e.g. scan durations, associations/disassociations and context throughout the measurements.

Active Wi-Fi Logger

The purpose of the active Wi-Fi logger is to initiate roaming events for the STA and to obtain information about the current connection.

The reason handovers need to be triggered by the supplicant is due to the disabled Network Manager, as the supplicant does not actively initiate a handover unless the current connection is lost. This is further investigated in the test journal in Section 3.4.2. The handover is generally initiated through three approaches:

- If the wpa_supplicant detects that the AP cannot be reached, e.g. through a lack of acknowledgements or the beacon signal are no longer detected, it will automatically initiate a scan to roam to a new AP.
- A 'SCAN' command can be issued to the supplicant when an active connection is established, and if another eligible AP (i.e. the supplicant is configured to connect to it) is detected and has higher RSSI than the RSSI of the current APs beacon (with a difference of more than 5 dBm), the STA will automatically roam to this new AP. Note that the RSSI values used for comparison are only from the scan report and does not take current RSSI into account, so the measured RSSI of each AP can vary greatly due to the sample size of 1.
- If the STA has performed a scan recently (through the 'SCAN' command), it has a list of nearby BSSIDs for a 2-minute period by default. During this period, it can immediately start roaming to a specific AP on the list. Listing 3.2 contains an example of the command-line interface input and output when this occurs.

This approach will allow for custom steering of which AP the STA connects to at any given time, but will require other considerations to make this superior to the current "best RSSI" approach. For example, if the location of the STA, its destination and APs along the path is known, the number of handovers can be minimized, and performance increased. This is however not a focus of this project.

While other approaches to initiate the handover exists, approach 1 and 2 three will mainly be used for the project. The third approach will instead be used to testing and development, but is worthy to note as a key feature to be used in more advanced control schemes for the STA. Section 3.2.2 further details the roaming decisions of the wpa_supplicant.

```
> SCAN
OK
<3>CTRL-EVENT-SCAN-STARTED
<3>CTRL-EVENT-SCAN-RESULTS

> SCAN_RESULTS
bssid / frequency / signal level / flags / ssid
b8:19:04:57:2f:35       5500    -51     [WPA2-PSK-CCMP][WPS][ESS]
    NOKIA-2F2A_High

> ROAM b8:19:04:57:2f:35
OK
<3>SME: Trying to authenticate with b8:19:04:57:2f:35 (SSID='NOKIA-2
    F2A_High' freq=5500 MHz)
<3>Trying to associate with b8:19:04:57:2f:35 (SSID='NOKIA-2F2A_High'
    freq=5500 MHz)
<3>Associated with b8:19:04:57:2f:35
<3>CTRL-EVENT-SUBNET-STATUS-UPDATE status=0
<3>WPA: Key negotiation completed with b8:19:04:57:2f:35 [PTK=CCMP GTK=
    CCMP]
<3>CTRL-EVENT-CONNECTED - Connection to b8:19:04:57:2f:35 completed [id=1
     id_str=]
```

Listing 3.2: wpa_ctrl issued roaming event. Lines starting with '>' are commands, lines with '<3>' are callback responses, and lines without an initial identifiers are responses to commands.

To obtain information of the current connection, the commands 'STATUS' and 'SIGNAL_POLL' are issued whenever the periodic trigger is received, with example outputs shown in Listing 3.3 and 3.4, respectively. These commands will provide the key information highlighted in Table 3.4. Of these, the current frequency channel, SSID, BSSID, bitrate and RSSI will be logged, as well as a timestamp with microsecond-precision.

Table 3.4: Overview of Wi-Fi-related software on the Linux platform.

Property	Description
STATUS	
bssid	The Basic Service Set Identifier (BSSID) is a unique MAC address for the AP and is constant regardless of configuration (e.g. current SSID or encryption).
freq	Frequency of current AP in MHz.
ssid	The Service Set Identifier (SSID) is the name of the network, and is used to e.g. identify multiple APs belonging to the same Extended Service Set (ESS).
key_mgmt	Current encryption scheme being used for the connection. The passkey encryption WPA2-PSK is commonly used for many deployments.
ip_address	IP address of interface. Note that while this is reported through the software, this is not configured directly by wpa_supplicant.
SIGNAL_POLL	
RSSI	Received Signal Strength Indicator (RSSI) of the current signal in dBm.
LINKSPEED	Bitrate of the current connection. Because the value of this variable is not well-documented, and did not correlate with the RSSI based on initial measurements, it is logged but not used for further processing.
NOISE	The SNR of current signal. The value 9999 was reported by the supplicant regardless of current connection or conditions, and is used as an error-value as the noise value is not reported by the driver with the given hardware configuration.
FREQUENCY	Identical to freq reported by STATUS.
CENTER_FRQ1	Center frequency of the current channel in MHz.
AVG_RSSI	Average RSSI in dBm over an unknown number of samples. Because no documentation for the number of samples could be found (since this is computed in the driver or in the network card itself), the instantaneous RSSI values will be used instead.
AVG_BEACON_RSSI	The average RSSI in dBm of the beacon signal from the AP over an unknown number of samples.

```
> STATUS
bssid=b8:19:04:57:2f:35
freq=5500
ssid=NOKIA-2F2A_High
id=1
mode=station
wifi_generation=6
pairwise_cipher=CCMP
group_cipher=CCMP
key_mgmt=WPA2-PSK
wpa_state=COMPLETED
ip_address=192.168.18.127
p2p_device_address=14:f6:d8:66:40:46
address=14:f6:d8:66:40:45
uuid=6ce3fea7-e28e-5d80-a2db-eb3e8a5d297b
```

Listing 3.3: Example output of STATUS command. The initial line contains the command.

```
> SIGNAL_POLL
RSSI=-49
LINKSPEED=6
NOISE=9999
FREQUENCY=5500
WIDTH=20 MHz
CENTER_FRQ1=5500
AVG_RSSI=53
AVG_BEACON_RSSI=-52
```

Listing 3.4: Example output of SIGNAL_POLL command. The initial line contains the command.

Passive Wi-Fi Scanner

The purpose of this plugin is to allow for interference-free scanning of nearby APs. Although the amount of interference generated through the beacon requests from the STA is minimal compared to regular traffic patterns, it is desirable to completely remove the interference where possible. As such, it is possible for the wpa_supplicant software to utilize passive scanning, where a device will instead simply listen on a channel. While the exact timing can be dependent on the vendor implementation, the interval between two beacons from an AP most commonly set to a default of 102.4 ms [9], and as such a similar interval can be assumed. If equal time is spent scanning each channel, and ignoring the time required for the radio to switch frequency channels, a total of 3.9 seconds is required for a full scan of 38 channels. It is therefore more desirable to utilize the active scanning, however, as the purpose of this plugin is to obtain an overview of nearby APs, the speed of the scans is not of high priority. Furthermore, as the STA being used for monitoring will not be used for active communication (i.e. it will never attach to an AP), the latency impact of traffic can be ignored.

Based on 100 samples for each setting using the Intel NUC, it was found that the active scanning method requires 0.94 seconds on average to perform the scan, while the passive scanning requires 3.47 seconds. As this is lower than the estimate, it is clear that the

time used for passive scanning may be shorter than anticipated, however, as this is controlled by the network interface itself, this has not been further investigated.

To enable the passive scanning with the wpa_supplicant, the *passive_scan=1* parameter is supplied with the initial configuration. It can however also be enabled during runtime. If only a subset of channels should be scanned, e.g. to increase the sampling rate, these can be specified through the *freq_list* parameter with channels being specified in MHz. The scans are the issued using the 'SCAN' command and results through the 'SCAN_RESULTS' command, as shown in Listing 3.5. The wpa_supplicant uses a callback function to notify when new data is available, such as the scan results. The plugin therefore ignores the external triggers to start samples until results are obtained. The resulting data is then parsed and saved to a file.

```
1  > SCAN
2  OK
3  <3>CTRL-EVENT-SCAN-STARTED
4  <3>CTRL-EVENT-SCAN-RESULTS
5
6  > SCAN_RESULTS
7  bssid / frequency / signal level / flags / ssid
8  9a:18:98:bd:be:a3       5660    -65     [WPA2-PSK-CCMP][WPS][ESS]
       AAU5G_CISCO
9  9a:18:98:bd:bf:0c       5700    -66     [WPA2-PSK-CCMP][WPS][ESS]
       AAU5G_CISCO
10 9a:18:98:bd:be:e4       5680    -79     [WPA2-PSK-CCMP][WPS][ESS]
       AAU5G_CISCO
```

Listing 3.5: Example output of the 'SCAN' command. Lines starting with '>' are commands, lines with '<3>' are callback responses, and lines without an initial identifiers are responses to commands.

Cellular Technology Logger

Similar to the wpa_supplicant, several software solutions exist to handle the communication to and control of cellular modems for 4G and 5G. The inner workings of this software are however out of the scope of this project, and thus the approach to gathering the relevant data is presented.

The software in question for this plugin is the libqmi library [21], specifically the qmicli command-line interface. The library handles communication to WWAN modems and devices using the Qualcomm MSM Interface (QMI) protocol. Information is requested through the commands shown in Listing 3.6, where *-p* specifies that the commands should be run through a proxy (improving the response time for the request), -d specifies the interface to be used for communication (commonly denoted as cdc-wdm0, cdc-wdm1 etc.), and the final argument being the command to be issued to be issued.

```
qmicli -p -d /dev/cdc-wdm1 --nas-get-signal-info
qmicli -p -d /dev/cdc-wdm1 --nas-get-signal-strength
qmicli -p -d /dev/cdc-wdm1 --nas-get-cell-location-info
```

Listing 3.6: qmicli commands used for the cellular plugin.

Table 3.5 presents the information provided in the response to the first two commands (signal-info and signal-strength). These values describe the connection between the modem and its currently serving cell. While signal-info provides more information regarding the signal, it was found that it was not supported for 5G configurations, which can partly be attributed to the equipment being state-of-the-art and do not yet support the same features as its predecessor (4G LTE). Furthermore, the SNR cannot be obtained for 5G as it is not reported by the modem.

Table 3.5: Parameters of –nas-get-signal-info and –nas-get-signal-strength commands [22].

Parameter	**signal info**	**signal strength**	**Description**
RSSI	Yes	Yes	Received Signal Strength Indicator.
ECIO	Yes	No	EC/IO indicates the downlink carrier-to-interference ratio.
IO	Yes	No	Signal quality, similar to ECIO.
SINR	Yes	No	Signal-to-Interference plus Noise Ratio, indicates the throughput capacity of the channel.
RSRQ	Yes	Yes	Reference Signal Received Quality, indicates the quality of the received reference signal.
SNR	Yes	Yes	Signal-to-Noise Ratio.
RSRP	Yes	Yes	Reference Signal Received Power, indicates the power of the Reference Signals spread over the full bandwidth and narrowband.

Listing 3.7 shows an example of the output of the *–nas-get-cell-location-info* command. Here, information regarding which cells are nearby, which cell is currently used, and some signal quality properties are included in the response.

```
Intrafrequency LTE Info
        UE In Idle: 'yes'
        PLMN: '99940'
        Tracking Area Code: '153'
        Global Cell ID: '113920'
        EUTRA Absolute RF Channel Number: '43190' (E-UTRA band 42: TD
    3500)
        Serving Cell ID: '0'
        Cell Reselection Priority: '1'
        S Non Intra Search Threshold: '62'
        Serving Cell Low Threshold: '40'
        S Intra Search Threshold: '62'
```

```
        Cell [0]:
                Physical Cell ID: '0'
                RSRQ: '-11.3' dB
                RSRP: '-77.8' dBm
                RSSI: '-46.5' dBm
                Cell Selection RX Level: '52'
        Cell [1]:
                Physical Cell ID: '1'
                RSRQ: '-11.5' dB
                RSRP: '-80.5' dBm
                RSSI: '-56.4' dBm
                Cell Selection RX Level: '43'
Interfrequency LTE Info
        UE In Idle: 'yes'
LTE Info Neighboring GSM
        UE In Idle: 'yes'
LTE Info Neighboring WCDMA
        UE In Idle: 'yes'
```

Listing 3.7: qmicli –nas-get-cell-location-info result for 4G LTE.

Traffic Logger

By capturing the traffic of on-going communication for later analysis, characteristics such as the throughput, communication flow and for some traffic types the response time can be determined. A commonly used tool to log the data for Linux-based systems is tcpdump [23]. This software allows for both capture and analysis of packets for a given interface, but will for the purposes of this project be used solely for packet capture.

```
tcpdump -tt -Z root -w FILENAME -i INTERFACE -G 30 -W 1 FILTERS
```

Listing 3.8: Tcpdump command.

The software is used in this plugin through the command shown in Listing 3.8, where

- -tt: Print timestamp as seconds since epoch and include fraction of a second.
- -Z: Run as specified user. The user 'root' is used to obtain sufficient privileges.
- -w: Path to file for storing data.
- -G: Time to capture for in seconds.
- -W: Limits number of files to capture to.
- -i: Interface to log data from. The parameter 'any' can be supplied to listen for traffic on all interfaces.

– <filters>: Expression to be evaluated when logging traffic. For example, the logging can be limited to only UDP traffic or all packets containing a specific IP.

A rotating approach is used in terms of saving the data to files. The traffic is logged to a file over a 30-second period, after which a new file is created and logged to. This is done to avoid corrupting the entire file in case an error occurs during logging.

Location Logger

To enable the mobility aspect of the tests in this project, a MiR robot will be used due to its availability and the positioning system. This robot is also able to provide power and a local network for the NUC. The MiR robot uses sensors to determine location of nearby objects, which it then uses to avoid while moving. It can be controlled manually through a web-interface accessible through either its own Wi-Fi AP, or through an external Wi-Fi network it has connected to. It is however also possible to create a pre-determined for it to follow independent of the web interface. It will then utilize a pre-scanned map of the environments to plan a route while also using the on-board sensors to avoid obstacles. The robot moves with a maximum speed of 1.5 m/s [24].

When planning a route from one point to another, the MiR robot can determine its own location using a laser-based positioning system with an accuracy of 5 cm [24]. This position, along with the current orientation, can be obtained through a REST interface with the robot. Representational State Transfer (REST) is a software architecture containing a set of constraints used with web services. Here, systems can access and manipulate data on the service through stateless operations, such as HTTP GET/POST operations. The GET operation can therefore be used to obtain the location data in JSON format as shown in Listing 3.9, with the position highlighted in red. It can also be seen that other data such as current mission status is available.

```
1  HTTP GET mir.com/api/v2.0.0/status
2
3  HTTP RESPONSE:
4  {
5      "allowed_methods": null,
6      "battery_percentage": 47.400001525878906,
7      "battery_time_remaining": 12157,
8      "distance_to_next_target": 32.4244384765625,
9      "errors": [],
10     "footprint":
       "[[0.506,-0.32],[0.506,0.32],[-0.454,0.32],[-0.454,-0.32]]",
11     "joystick_low_speed_mode_enabled": false,
12     "joystick_web_session_id": "",
13     "map_id": "25588a0b-fbdf-11ea-898a-0001299f16e3",
14     "mission_queue_id": 742,
15     "mission_queue_url": "/v2.0.0/mission_queue/742",
16     "mission_text": "Moving to 'front' (32.5 meters to goal)",
```

```
    "mode_id": 7,
    "mode_key_state": "idle",
    "mode_text": "Mission",
    "moved": 50561.28877147002,
    "position": {
        "orientation": 179.6822052001953,
        "x": 35.491058349609375,
        "y": 23.78909683227539
    },
    "robot_model": "MiR200",
    "robot_name": "MiR_202003003",
    "safety_system_muted": false,
    "serial_number": "",
    "session_id": "85cd7f3f-f2b7-11ea-ad20-0001299f16e3",
    "state_id": 5,
    "state_text": "Executing",
    "unloaded_map_changes": false,
    "uptime": 65817,
    "user_prompt": null,
    "velocity": {
        "angular": 2.953000068664551,
        "linear": 0.763059675693512
    }
}
```


Listing 3.9: MiR HTTP GET position

Latency Logger

With the latency performance being a main focus point for the project, it is necessary that sufficient care is taken to ensure this plugin provides reliable results. This latency logger will provide information regarding the latency distribution (i.e. either the Round-Trip Time (RTT) or one-way delay as well as packet drop statistics).

The RTT latency will be used for measurements as compared to the one-way delay, as time synchronization between two devices can lack the required accuracy when cabled connections are not available. The results can therefore instead be compared to the expected time between a request and a response for an application, and because Wi-Fi does not operate in a scheduled manner as e.g., 4G LTE does, the results can furthermore provide some insight into the one-way delay.

Two approaches have been investigated towards measuring the latency and packet drop statistics: the build-in Linux ping functionality, and a custom UDP-based ping.

Linux Ping

Most devices with internet-connection are capable of sending and responding to Internet Control Message Protocol (ICMP) ping requests. While this functionality is commonly used to determine the reachability of a host, it furthermore provides the RTT for a given

ping based on the duration until the response was received. Because this functionality is compatible with a vast number of devices, this requires minimal configuration of the target.

```
> ping -I eth0 -i 0.1 -s 64 -c 5 -D 192.168.0.1
PING 192.168.0.1 (192.168.0.1) from 172.19.243.82 eth0: 64(92) bytes of
    data.
[1621234619.800194] 72 bytes from 192.168.0.1: icmp_seq=1 ttl=63 time
    =2.84 ms
[1621234619.900276] 72 bytes from 192.168.0.1: icmp_seq=2 ttl=63 time
    =2.64 ms
[1621234620.001389] 72 bytes from 192.168.0.1: icmp_seq=3 ttl=63 time
    =2.29 ms
[1621234620.101527] 72 bytes from 192.168.0.1: icmp_seq=4 ttl=63 time
    =2.18 ms
[1621234620.202380] 72 bytes from 192.168.0.1: icmp_seq=5 ttl=63 time
    =2.13 ms

--- 192.168.0.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 403ms
rtt min/avg/max/mdev = 2.125/2.413/2.835/0.276 ms
```

Listing 3.10: Linux Ping example.

The pings can be configured based on a large number of parameters, with selected options shown in the example in Listing 3.10, where

- -I: Interface
- -i: Packet interval (given in seconds between each packet transmission)
- -s: Packet size
- -c: Number of ping requests to be sent before exiting
- -D: Include timestamp in output
- <target>: IP of destination

If multiple interfaces are to be used simultaneously, multiple instances of the ping program have to be called, each with their own internal timer. When pinging at very low intervals (< 10 ms), the tool was found to instead use a lower-bound ping interval of 10 ms regardless of the provided setting. The limit did however not have an impact on the tests performed in this project, where an interval of 50 ms was utilized.

Custom UDP Ping

Instead of relying on the built-in ping functionality, a custom application was created for the project. This application has certain advantages over the built-in functionality: 1) supports packet intervals of down to $\sim$1 ms, 2) can use custom traffic models for packet

departure times (e.g. exponential models) and 3) can synchronize packet transmission across interfaces.

This approach consists of two parts: the transmitter and receiver. Because this approach won't use the widely-available ICMP protocol, and to mimic its behavior in terms of packet processing and minimizing overhead, the User Datagram Protocol (UDP) is used. This will likewise ensure information regarding packet drop statistics are available, which would otherwise be handled by the Transmission Control Protocol (TCP). The transmitter initially opens sockets on two ports, one being used for requests, and the second for responses. The contents of the payload of the request packet are illustrated in Figure 3.4, where the ID designates the interface being used for transmission (being kept track of by the transmitter), the sequence number is used for identifying packets for packet drop statistics, and the timestamp being the time at which the packet was sent from the application layer. The remainder of the packet consists of filler bytes to match a desired frame size.

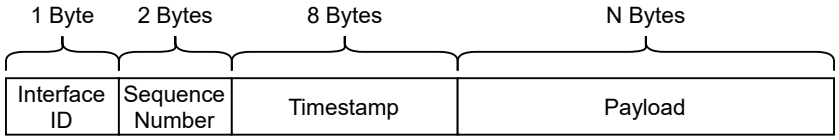

Figure 3.4: Custom UDP ping packet structure.

To transmit the packets, an internal trigger functionality is used, similar to the one used to coordinate the plugins. This trigger is used to synchronize the transmission of packets between the interfaces at the application-level. Although perfect synchronization cannot be guaranteed due to internal thread scheduling of the operating system and packet handling through the different interfaces (e.g. Wi-Fi vs. LTE), this allows for better comparison between them as compared to not using any coordination whatsoever. Although the trigger was used with a constant interval similar to the Linux ping, it can be extended to use any given traffic model. This can e.g. allow for RTT measurements similar to that an AMR would experience in terms of how it contacts a fleet manager.

Because this approach operates in the application layer, an application has to be running at the receiver which can respond to the ping requests. When this application receives a packet on the specified port, it transmits an identical packet back in terms of the packet structure in Figure 3.4, resulting in the transmitter performing the necessary RTT calculations. This packet is sent to the port the transmitter is listening on.

3.2.2 Roaming with wpa_supplicant

Because the majority of roaming events are based on the wpa_supplicant, it is relevant to look into how this is done, and determine which conditions are necessary to trigger a

roam. This section will provide a general overview for this aspect, based on the open-source code available at [17].

wpa_supplicant has an internal function solely for determining if a roam within the same ESS is necessary, where it compares the current BSS to a selected BSS from the scan results. The following parameters are considered for the current and selected BSS:

- $*\_SSID$: Service Set ID
- $*\_level$: Received Signal Strength Indicator (RSSI)
- $*\_snr$: Signal-to-Noise Ratio
- $*\_est$: Estimated throughput (based on SNR)
- to_5ghz: Set to 1 if current BSSID is in 2.4 GHz band and selected BSSID is in 5 GHz band, 0 otherwise.
- $diff$: Difference in RSSI between current and selected BSS ($selected\_level - current\_level$).

The scan results are used for the parameters for the selected BSS, while the information for the current BSS is from polled information from the driver, or likewise from the scan results of the driver did not respond to the poll. If the driver did not respond, the SNR for the current connection is set to a value of 0.

The decision of whether to roam is chosen from the list below (with 'selected' denoting the selected BSS, and 'current' denoting the current BSS):

1. Has a request to reassociate been requested explicitly?
   - If yes, initiate roam.
2. Is the STA currently disconnected?
   - If yes, initiate roam.
3. Is the current SSID unknown?
   - If yes, initiate roam.
4. $current\_SSID \neq selected\_SSID$
   - If yes, initiate roam. Selected network is in a different network block.
5. Has driver performed roam?
   - If yes, skip wpa_supplicant roam.

   Fetch current BSS entry from scan results.
6. Is current BSS located in scan results?
   - If no, initiate roam. Current BSS is not found in scan results.

7. Is the current BSS identical to the selected BSS?
    - If yes, skip roam. Selected and current BSS are the same.
8. Was the current BSS found in previous scan results?
    - If yes, initiate roam.

Same-ESS-specific roaming:

9. Does selected have a preferred BSSID?
    - If yes, initiate roam.
10. $selected\_est > current\_est + 5000$
    - If yes, initiate roam. Selected has better estimated throughput, given a threshold.
11. $current\_level < 0$ and
    $current\_level > selected\_level + to\_5ghz \cdot 2$ and
    $selected\_est < current\_est \cdot 1.2$
    - If yes, skip roam. Current has better signal level. Also prefers switching to 5 GHz band, but will not switch if the estimated throughput is decreased by some threshold.
12. $current\_est > selected\_est + 5000$
    - If yes, skip roam. Current has better throughput, given a threshold.
13. $current\_snr > 25$
    - If yes, skip roam. Current BSS has good SNR (higher than 25). The reasoning for this is based on the rule of thumb that anything above 20 SNR is considered 'good', and that 25 SNR is minimum for 54 Mbps data rate.

    Calculate min_diff based on current signal level, preferring good RSSI:

    - If $current\_level < -85$, $min\_diff = 1$
    - Otherwise if $current\_level < -80$, $min\_diff = 2$
    - Otherwise if $current\_level < -75$, $min\_diff = 3$
    - Otherwise if $current\_level < -70$, $min\_diff = 4$
    - Otherwise if $current\_level < 0$ (is a valid value), $min\_diff = 5$
    - Otherwise $min\_diff = 5$

    Alter min_diff based on estimated throughput values, with magnitude of change depending on difference between current and selected throughput:

    - If $current\_est > selected\_est \cdot 1.5$, increment min_diff by 10
    - Otherwise if $current\_est > selected\_est \cdot 1.2$, increment min_diff by 5
    - Otherwise if $current\_est > selected\_est \cdot 1.1$, increment min_diff by 2

- Otherwise if $current\_est > selected\_est$, increment min_diff by 1
- Otherwise if $selected\_est > current\_est \cdot 1.5$, decrease min_diff by 10
- Otherwise if $selected\_est > current\_est \cdot 1.2$, decrease min_diff by 5
- Otherwise if $selected\_est > current\_est \cdot 1.1$, decrease min_diff by 2
- Otherwise if $selected\_est > current\_est$, decrease min_diff by 1

If to_5ghz is set to 1, decrease min_diff by 2

14. $diff < min\_diff$
    - If yes, skip roam. Difference between signal levels is too small.
    - If no, initiate roam due to difference in signal level.

3.2.3 Enabling Wi-Fi Improvements

As described in Section 3.1.4, there are three main improvements which are to be implemented: client balancing, optimized scanning and IEEE 802.11r Fast BSS Transfer. Client balancing was enabled through the Cisco Meraki dashboard, where a toggle could be found to enable this functionality. The optimized scanning was enabled through the configuration file for wpa_supplicant, where the global parameter *freq_list* is used to specify which frequency channels to be scanned, as described in Section 3.2.1. Enabling the IEEE 802.11r functionality requires configuration on both client and network side. For the Cisco Meraki setup, the FT functionality was enabled using the dashboard, with three modes of operation being available:

- Off: IEEE 802.11r is disabled
- Adaptive: For older Apple devices which only support parts of the IEEE 802.11r standard can be used with the network and obtain some of the benefits. Devices with IEEE 802.11r enabled will not be able to connect at all.
- On: Devices using IEEE 802.11r will be able to utilize the protocol to obtain improved performance.

After changing this setting to On, it was confirmed to have an effect by scanning the network and observing the encryption type, see Listing 3.11. It was however still possible to connect to it, which was found to be part of a fallback mode specific to Cisco Meraki that enables non-FT devices to connect [15].

```
1 Normal network:
2 BSSID / frequency / signal level / flags / SSID
3 9a:18:98:be:b6:ce   5660   -48   [WPA2-PSK-CCMP][ESS]        AAU5G_CISCO
4
5 IEEE 802.11r-enabled network:
6 BSSID / frequency / signal level / flags / SSID
7 9a:18:98:be:b6:ce   5660   -45   [WPA2-PSK+FT/PSK-CCMP][ESS]  AAU5G_CISCO
```

Listing 3.11: Scan results

When connecting to an FT-enabled AP using a non-FT STA configuration, there are as expected no difference in the handover performance, as the STA goes through the normal phases of scanning, authorizing, associating and 4-way handshake, after which it is finally connected. To enable FT on the device, it is necessary to, according to the wpa_supplicant, specify the encryption key as FT-PSK in the configuration [19]. This allows for connecting to the AP using this new encryption type, and the logs from wpa_supplicant shows that new FT keys are being computed in the beginning of the 4-way handshake phase.

To determine the performance of the FT-feature, a handover test was conducted. Here, the robot was configured to drive in a pre-fixed path through the lab, and when the connection from the device to an AP reaches -92 dBm, the wpa_supplicant software will automatically detect the losses of the beacons from the AP, do a scan and establish a connection to another eligible AP. Although the FT-feature was believed to be enabled, no improvements were observed, as shown in Figure 3.5. The Cisco Meraki Dashboard likewise presented the NUC as not being IEEE 802.11r-enabled. To verify if this was correct, an Apple iPhone 12 was connected to the network. This is due to the phone being listed as supporting the IEEE 802.11r feature. This did result in the dashboard showing IEEE 802.11r as being enabled, but due to the platform it was not possible to obtain logs similar to wpa_supplicant, and likewise there were no FT-specific logs available in the dashboard to identify when FT was being used.

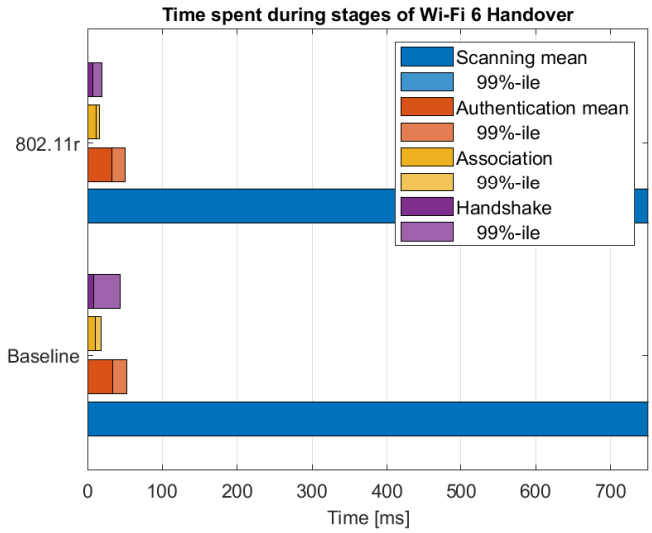

Figure 3.5: Time spent during handover-phases based on wpa_supplicant logs.

Because of the similar performance between FT and non-FT, it was believed that this was a configuration issue. A secondary device was used for network monitoring/packet sniffing to capture the connection process. In the association request packet, it was con-

firmed through Wireshark that the encryption type in the Robust Security Network-field (RSN) was set to FT-PSK, as shown in Figure 3.6. It was also found that a mobility-related field was present in the extended capabilities block, which has been confirmed to be non-present when IEEE 802.11r is disabled on the client. It contains information of which features the device supports and shows that both the Resource Request Protocol (RRP) and over-the-BS was not supported. The RRP field is not supported with IEEE 802.11r, so this behavior is expected [25]. The over-the-BS functionality was also confirmed to be disabled in the wpa_supplicant logs, where the contents of the association packet can be seen being constructed based on the information from the driver. This can be seen in Listing 3.12.

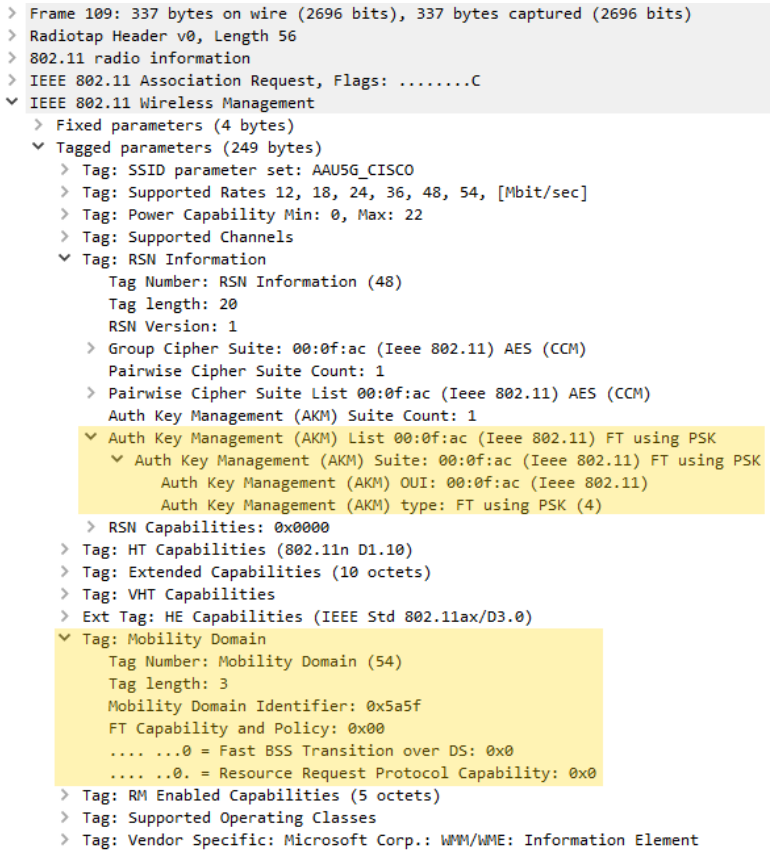

Figure 3.6: Association request packet inspected using Wireshark.

```
1 FT: MDE - hexdump(len=3): 5f 5a 00
2 FT: Mobility domain - hexdump(len=2): 5f 5a
3 FT: Capability and Policy: 0x00
4 FT: Stored MDIE and FTIE from (Re)Association Response - hexdump(len=5):
    36 03 5f 5a 00
5 SME: FT IEs - hexdump(len=179): [Removed]
6 wlp2s0: SME: FT mobility domain 5f5a
```

Listing 3.12: Mobility domain contents from wpa_supplicant log.

Although the association packet stated otherwise, and since the over-the-DS roaming functionality was supported in the wpa_supplicant software, this was tried as well to determine if IEEE 802.11r roaming worked. This feature is initiated by the user since the device has to communicate with the current AP, and will therefore not be applicable in the tests performed for the handovers (where the device would lose connection completely). It was observed that the command could be successfully executed and that from both the wpa_supplicant logs and network monitoring it could be observed that an over-the-DS request was sent to the network. This request was however ignored by the network, as no form of response was transmitted back to the device. There is likewise no indication of this request from the network monitoring tools. It is likely that the request was ignored since the device advertises it cannot support these messages. Another option is that the packet itself was not valid, however this has not been verified.

As described in [26], the IEEE 802.11r feature is mainly targeting the WPA2-EAP encryption type, instead of the current WPA2-PSK. Switching from the pre-shared key to an authentication server will result in a larger authentication/handshake process, as more messages are transmitted before access is granted. It is therefore much more beneficial to use FT, as these may then be skipped, and should be noticeable when inspecting handovers. While the Cisco networks support using a cloud-hosted authentication server, this was not deemed useful in the current setup, as the latency was of interest - offloading traffic to an external site will therefore cause significant and unpredictable delays, where a local server will avoid this issue. A radius authentication server was therefore installed on the local network using freeradius [27], using a standard installation with a single user. After verifying the setup with FT disabled, where a device could successfully connect to the network with the new encryption type, FT was reenabled and the device set to utilize EAP-FT encryption. It was, like previously, clear from the wpa_supplicant logs that the FT-related computations were made on the device, but judging from both the communication between AP and device, as well as from the performance in terms of latency, the desired effects weren't present.

Based on the previous observation with the lack of support for FT features, the installation of the software on the device itself was investigated. wpa_supplicant was tested with both the latest release as well as the latest development version, from which it could also be confirmed that the program was indeed built with IEEE 802.11r features enabled. Likewise, if the software were run without this feature enabled, using the FT-PSK encryption type would not be recognized at all, resulting in the program being unable to process its configuration. Tests similar to those done previously (with regards to handovers) were done, but with similar results. The driver for the network card itself, iwlwifi [28], was updated with the latest release, again with no difference. The general drivers for 802.11 (nl80211, cfg80211) were shipped with wpa_supplicant, so these were not believed to be the cause of the problem.

Although the IEEE 802.11r features are widely supported for many Windows and especially Apple systems, it was clear that the support for Linux-based clients was lacking.

It should however be noted that the wpa_supplicant is also used for Android-based systems, and since there are several people who claims to have successfully used IEEE 802.11r on these [29], it could not be completely ruled out that the system was unsupported. It was likewise found that on the Linux support page for Intel Wireless Devices that the wpa_supplicant software is supported [30]. After asking on an Intel Community Forum, it was however clear that the support for the chosen platform (the Intel NUC5i3MYBE) was not tested to work with Linux, and they likewise did not recommend changing hardware due to compatibility issues [31].

As described in [32], the issue may have been related to how roaming is defined. When mobility is introduced, it becomes relevant to actively monitor the current connection and either act preemptively or during roam-time. When acting preemptively, the device can monitor e.g. the signal strength to determine when it is necessary to perform a roam. For the latter, the device may initiate the roam when e.g. the AP beacon frames are no longer being received and only then start searching for other networks. Since all of the handover tests were done by moving the robot out of the coverage area manually, a test was conducted where the device would be moved from one access point to another without it losing connection, resulting in the measured RSSI from the second AP being much larger than the previous AP. When then performing a scan, the device would disconnect from its current AP and successfully perform a roam to the AP with highest signal strength. The roam was verified to be using IEEE 802.11r, as it was included in the wpa_supplicant log that an FT over-the-air handover would be performed, and from the lack of a handshake-stage. This behavior was likewise verified to be repeatable for all APs.

It is believed that this lack of FT handover could be a software bug in the behavior of the supplicant, as the device was still computing the FT keys, or due to it simply not being considered a roam (by losing the connection entirely). As described in Section 3.2.1, the Network Manager software-initiated scans in 2-minute intervals which were disabled due to causing notable spikes in latency. These scans are believed to be able to initialize the handovers, but the behavior was not identified due to the manager being disabled.

It can be concluded that FT roaming must be initiated either by the user or by a program sending commands to the wpa_supplicant software. In order to mimic the previous behavior of only roaming when necessary, a feature was added to the measurement software, where a scan would be initiated when the signal reached -85 dBm or lower. The signal strength was chosen based on previous gathered data, where it was identified that the time period from -85 dBm until a handover was around 15 seconds with the given movement patterns, which is deemed to be more than sufficient to disconnect from the current AP properly (rather than simply losing connection). While this value could be closer to the -92 dBm point at which the software would automatically disconnect, it is not believed to cause more handovers during the tests, as the route is planned out in a way that the device will move straight from one access point to another, rather than

staying in the edge of one APs coverage area.

Finally, although the device can successfully do FT over-the-air handovers, it remains unknown why over-the-DS handovers are not supported. The Cisco Meraki dashboard likewise still considers IEEE 802.11r to be disabled for the device. This can likely be due to it not supporting these features, but may also be related to Cisco working together with e.g. Apple to support their devices, and may then have other identifiers to consider before considering them to be IEEE 802.11r-capable.

3.3 Multi-connectivity in Wi-Fi

This section will present the design of the radio-aware multi-connectivity layer-4 scheduling mechanism. The hardware and software platform as well as the multi-access gateway software will initially be presented, followed by the design of the layer-4 packet scheduling mechanism and mobility coordinator. A test journal verifying the functionality of the multi-connectivity approach is found in Section 3.4.4 and a test journal for the tests for Paper 2 in Section 3.4.5.

An overview of the solution is illustrated in Figure 3.7. Here, outgoing traffic is routed to a packet scheduler, which encapsulates and transmit the packet on an interface. This scheduler can be configured to use either of the two illustrated interfaces, e.g. Interface 1, or both through packet duplication. Because of the encapsulation and duplication, an eligible receiver needs to be present on the network to decapsulate and forward the traffic to the intended receiver. This is handled using the multi-access gateway software. To introduce the radio-aware aspect, an external mobility coordinator is used to configure the packet scheduler. An example of this could be that the current interface is experiencing poor radio conditions, and that Interface 2 should be used instead.

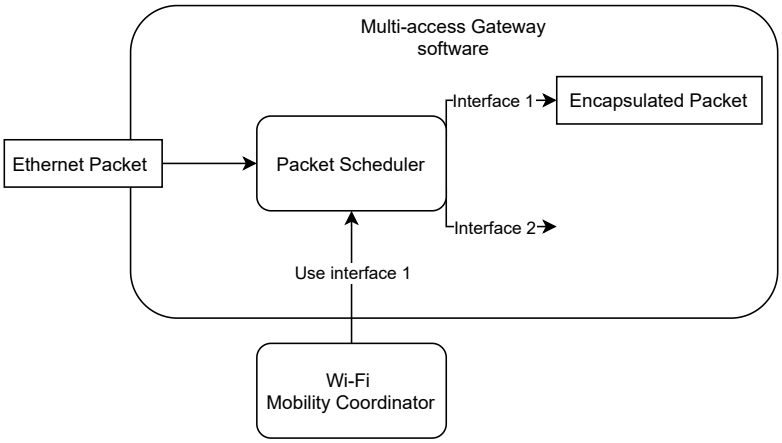

Figure 3.7: Overview of the radio-aware multi-connectivity packet scheduler.

3.3.1 Multi-access Gateway

The Multi-access Gateway provides seamless connectivity between groups of network devices. By encapsulating traffic on the local network (through e.g. Ethernet), traffic is forwarded to other gateways, which then transmits the packet to their local network. This is illustrated in Figure 3.8, where the gateways allow for all end-devices to communicate with one-another. Instead of relying on cabled connections, the gateways can furthermore take advantage of any Linux-compatible wireless communication technology,

such as Wi-Fi, 4G LTE and 5G NR.

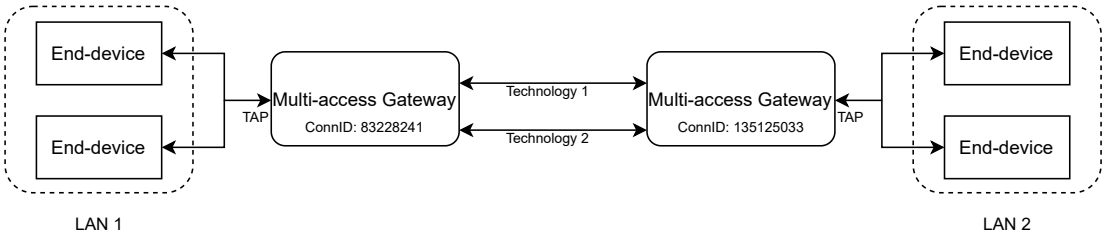

Figure 3.8: Multi-access gateway example setup, allowing two separate LANs to communicate with one-another.

The gateways are initially configured using the following information:

- Ethernet-interface to be used as TAP for incoming and outgoing information. This interface is connected to the end-devices on the LAN, where traffic from end-devices is forwarded to other gateways, or packets from other gateways are transmitted on the TAP.
- A list of interfaces to be used for communication with other gateways. The gateway will duplicate the traffic on all specified interfaces.
- A list of other gateways, including a unique connection ID and IPs of their interfaces.

When a packet is received by the gateway from the TAP, the entire Ethernet frame (including headers) is encapsulated in a new frame, illustrated in Figure 3.9. The original unaltered packet, highlighted in blue, is placed in the payload of the new frame in addition to the connection ID of the transmitting gateway and a sequence number highlighted in green. This number, sequentially increasing for each encapsulated frame, is used for identification of duplicate packets at the receiving gateway. When a gateway receives an encapsulated packet, it first checks if the packet has already been received from the gateway using the connection ID and sequence number, in which case the packet is ignored. This is illustrated in Figure 3.10. Otherwise, the decapsulated packet (in blue) is transmitted to the local network through the TAP. This is to avoid transmitting the same packet to the local network twice, which would then be processed by the end-devices.

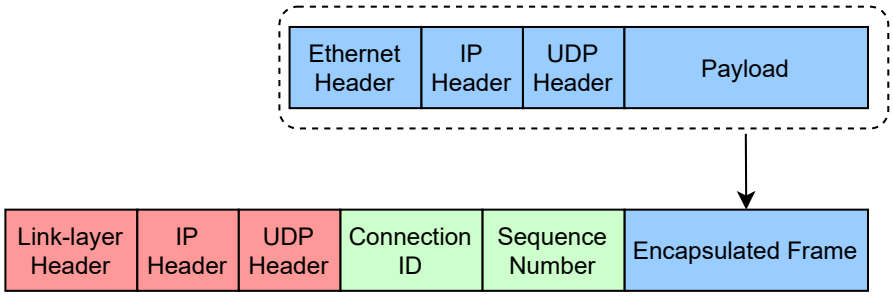

Figure 3.9: Multi-access gateway packet encapsulation. The original frame is highlighted in blue, metadata specific to the gateway in green, and headers from layer 2, 3 and 4 in red.

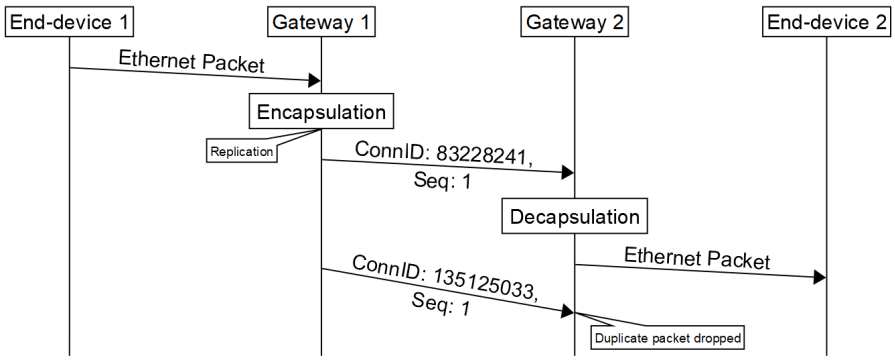

Figure 3.10: Sequence diagram of communication using gateways, highlighting the encapsulation and decapsulation process.

The UDP protocol is used for transmissions between gateways, as connection-oriented traffic will rely on the TCP header in the original packet. Therefore, if a packet does not arrive at the intended end-device, a retransmission request will arrive at the original transmitter end-device, as the effects of the gateways are invisible to the local networks.

The gateways rely on routing protocols such as Address Resolution Protocol (ARP) broadcasting to determine how traffic should be routed. When a broadcast-packet is received, as illustrated in Figure 3.11, it is encapsulated and forwarded to all other gateways. When packets are received at the gateway, it saves the source MAC address and the originating connection ID (both present in the packet), and will thus be able to forward packets to that specific gateway based on the destination MAC address. After receiving the request, the end-device will generate an ARP reply with the corresponding destination MAC address. Because the gateway has already learned the source of this MAC address, it encapsulates and forwards the packet to the corresponding gateway, which likewise remembers the source MAC address and connection ID.

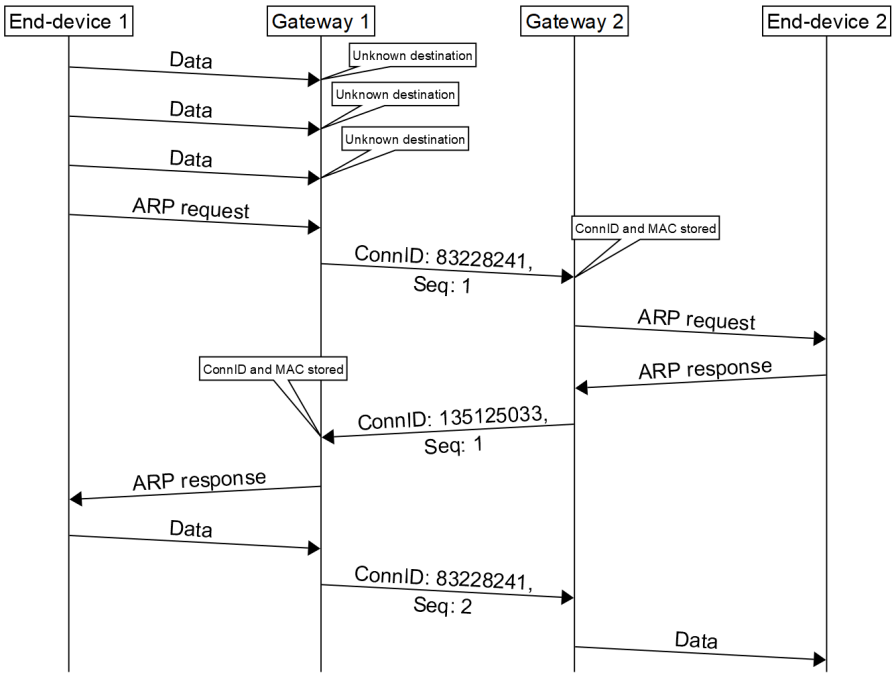

Figure 3.11: Sequence diagram of routing using gateways.

Hardware Platform

For this part of the project, the end-devices will consist of the same Intel NUCs presented in Section 3.2. Because the multi-access gateway will utilize two Wi-Fi 6 modems to enable multi-connectivity, a Gateworks GW6104 detailed in Table 3.6 will be used, as this is compatible with multiple network cards on a hardware level. It will, similar to the Intel NUC, use Ubuntu, and as such the measurement software will be compatible. It was, however, found that the IEEE 802.11r functionality was not compatible with this setup.

Table 3.6: Computer hardware and software setup for physical tests.

HW/SW	**Details**
Device Model	GATEWORKS GW6404
CPU	Cavium OcteonTX™ ARMv8 SoC @ 1.5GHz
RAM	4GB
OS	Ubuntu 20.04.2 LTS
Kernel	5.4.45
Wi-Fi Network Cards	2x Intel Wi-Fi 6 AX200

3.3.2 Packet Scheduler

Having multiple interfaces allows for a new approach to transmitting data. For the purpose of this project, we consider scheduling at layer-4, i.e. the IP layer. As the original traffic is encapsulated, the purpose of the packet scheduler is to ensure it is routed to the corresponding gateway with emphasis on which interface to be used. This section will assume two Wi-Fi interfaces are available, but the theory can easily be extended to involve more.

Two approaches to packet scheduling are considered, illustrated in Figure 3.12: packet duplication and best path scheduling. The scheduler is a central part the multi-access gateway software, and is adapted for the purposes of this project.

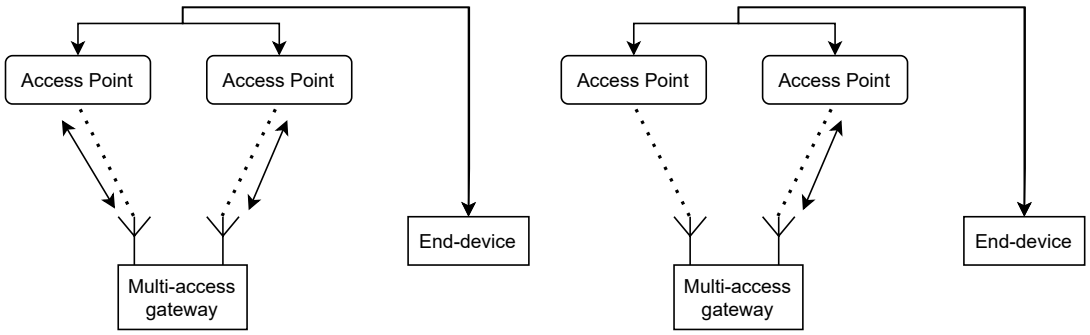

(a) Packet duplication. **(b)** Best path scheduling.

Figure 3.12: The two packet scheduling configurations used for the project.

Packet Duplication

Because the multi-access gateways ensure only a single copy of the packet is forwarded to the LAN, duplicating all traffic over the two available STAs can allow for a significant increase in reliability and decreased latency, but will effectively result in doubling all transmitted data. This will therefore harm other devices on the network due to the Listen-Before-Talk mechanism of Wi-Fi for a goodput of less than 50% when also considering packet headers.

Implementing the packet duplication on the multi-access gateway is done through the example configuration in Section 3.3.1. The gateway is configured to use both interfaces, but with same destination IP. Because the two interfaces operate independently, they will both listen on the medium until it is idle before transmitting data, but because of the random hold-off time, the probability of the packet colliding between the two STAs is considered to be low (comparable to collision between two individual STAs).
On the network-side, the gateway transmits data through the same (Ethernet) interface,

but targeting both of the Wi-Fi-enabled gateways' interfaces. The packets will be sent to the Ethernet-interface sequentially (in the same order), but from this point, scheduling will initially rely on the Ethernet interface, followed by the scheduling mechanisms in the AP.

Best Path Scheduling

The second approach only uses one STA at a time for communication, but still takes advantage of the presence of multiple STAs. To determine which STA to use for communication, the scheduler uses the RSSI of both STAs. Initially, the STA with highest RSSI is used for communication, referred to as primary. If the other STA, referred to as secondary, has an RSSI value of 5 dBm above the primary, the scheduler switches between the two, and the new primary (the one with highest RSSI) is used for communication. Because either STA of the multi-access gateway can be used to transmit packets, this switch is fully seamless and does not introduce any of the interruptions observed with handovers during roaming. This is illustrated in Figure 3.13. The RSSI is checked every 1.5 seconds, as the signal strength is not expected to change drastically.

Although the multi-access gateway can choose between which interface is best suited for communication, it is necessary to do the same on the gateway located at the network-side. While it could send packets to both interfaces, similar to the packet duplication, and rely on either message arriving, this will have the same implications as for the first approach in terms of goodput. Instead, it will be configured to only reply to whichever STA it last received data from, being determined using the IP. This ensures that a single instance of each packet is transmitted (resulting in higher goodput), and since the STAs can still receive traffic while in secondary mode, the packets will not simply be lost when switching between them. This does however rely on the roaming gateway to transmit data frequently to update this. For example, if the STAs are connected to different APs, but the connection for the primary is lost and the gateway switches and does not transmit data for a period of time, the gateway on the network will be unable to reach it.

While this packet scheduler only uses RSSI to determine which STA to use, the algorithm can be extended to include other aspects as well. For example, channel congestion estimation algorithms can enrich the scheduler to only use STAs with low congestion.

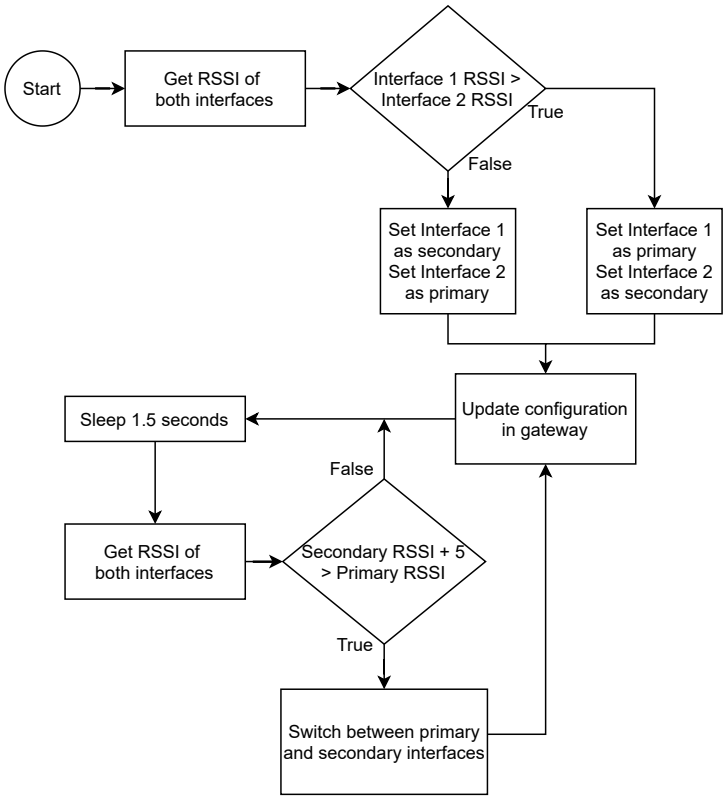

Figure 3.13: Flowchart of best path scheduling interface selection.

3.3.3 Mobility Coordinator

Using the packet scheduler by itself can yield improvements to the latency performance of a STA, but it may still experience some of the issues of a single-connectivity configuration; that is, the two STAs operate independently, and will thus likely be connected to the same AP due to low spatial diversity. Likewise, because of similar RSSI values, they may perform a handover simultaneously, resulting in a period without connectivity on either STA.

To improve upon the packet scheduler, the mobility coordinator will, in general, ensure the following:

1. The STAs are always connected to different APs.
2. The STAs never perform handovers simultaneously.

In network deployments where multiple APs overlap in terms of coverage area, it is possible to take advantage of this to increase the diversity in terms of paths for com-

munication. This will of course depend highly on the given network deployment. The mobility coordinator will operate based on a blacklist-approach using wpa_supplicant. Here, each AP can be listed as an individual network according to the BSSID (i.e. MAC address of the AP). These can then be enabled or disabled during run-time. If a device is connected to a network being blacklisted, it will immediately disconnect and search for another non-blacklisted AP.

The implementation of the mobility coordinator differs slightly between the two packet scheduling mechanisms, as maximum uptime is desirable for packet duplication, while the secondary STA for best path scheduling allows for more flexibility in terms of associating with the best AP.

Packet Duplication

Figure 3.14 contains a flowchart illustrating the algorithm for the packet duplication configuration. Upon startup, the algorithm checks whether the two STAs are connected to the same AP, in which case the current is blacklisted for one of the STAs. After a few seconds, the second AP will either remain disconnected or have connected to a new AP, which is then blacklisted for the first STA. After this, the main loop is entered. A 2.5 second delay is used for each loop to allow for STAs to establish a connection to new APs. Although redundant, the program again checks if both STAs are connected to their own respective AP, in which it blacklists it for STA 2.

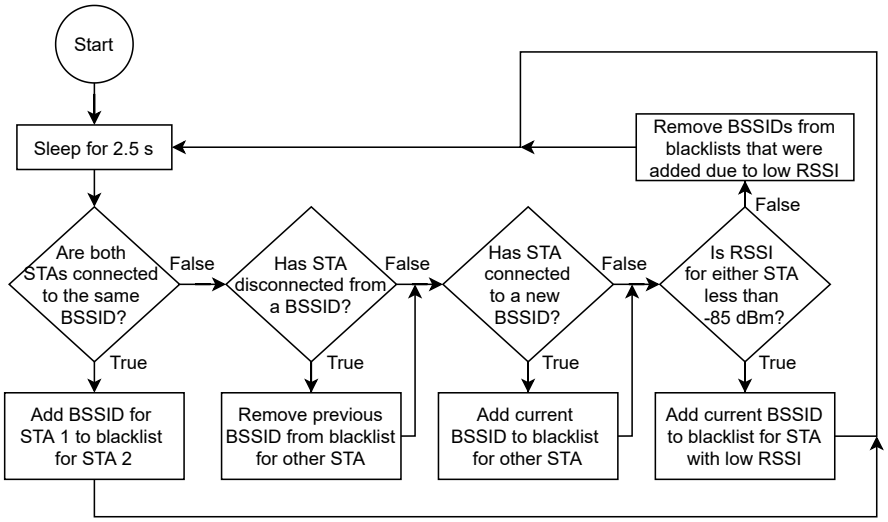

Figure 3.14: Mobility coordinator flowchart for packet duplication.

To keep the blacklists up-to-date, the algorithm keeps track of BSSID events, i.e. disconnections and new connections to APs. When a STA disconnects from an AP, its BSSID

is removed from the blacklist of the other AP. If a STA connects to a new BSSID, its BSSID is added to the blacklist for the other AP.

Finally, similar to the single-connectivity case, a roaming event is initiated when the RSSI of a STA reaches a threshold, in this case -85 dBm. Instead of performing a scan, the current BSSID is blacklisted temporarily until the next loop iteration. This forces the STA to disconnect from the current AP and perform a scan. While this does result in a longer period without connectivity than a regular scan, the other STA remains connected, and as such this disconnect will not have a significant impact on the overall latency. An improved approach would be to actively select which AP to roam to in the mobility coordinator, however this functionality was not supported on the multi-access gateway setup, but was on the Intel NUC. The cause of this was not looked into.

Best Path Scheduling

The main difference between best patch scheduler with and without the mobility coordinator is that effort is put into ensuring the two STAs do not connect to the same AP, and that the secondary STA roams to new APs. To achieve this, the current connected BSSID for the primary STA is blacklisted for the secondary. When the secondary reaches a threshold, selected to -85 dBm, a SCAN command is sent to the STA in an attempt to trigger a roam. If another AP is found which satisfies the requirements for a roam (see Section 3.2.2), the STA will initiate the handover. Otherwise, it remains connected to its current AP, and a 7.5 second timeout is used to avoid excessive scanning (corresponding to 5 loop iterations).

3.4 Test Journals

3.4.1 Static and Intra-AP Mobility Wi-Fi Latency Tests

Purpose of the Test

Introducing mobility to a Wi-Fi connection can result in increased latency. This test aims to determine what kind of impact this is. To do so, static, intra-AP and inter-AP measurements will be done.

Theory

Wi-Fi uses several mechanisms to improve the connection between an AP and STA, such as through pilot signals to estimate the phase offset, but as these are not in the scope of this project and will not be investigated. When the STA is mobile, the connection is less reliable, as these mechanisms will constantly need to adapt due to changes in environment, effects of scattering off nearby surfaces and by the signal strength degrading with an increased distance.

Test Setup

- Intel NUC equipped with Intel AX200 Wi-Fi 6
- 2x Wi-Fi 6-enabled devices for transmitting background data
- 3x Cisco MR36 Wi-Fi 6 access points
- Edge-server connected to the network through Ethernet
- MiR robot for movement and positioning data

The test was conducted in the AAU Smart Production Lab shown in appendix A. To avoid the static measurements being overly dependent on the location, 4 measurements were done at various locations, shown in blue on Figure 3.15. These points were chosen because of the overlap with the route used for intra-AP measurements, shown in red. Mean RSSI values at each location can be seen in Table 3.7.

Table 3.7: RSSI at the static measurement points.

Location	Single-user	With background traffic
Charger	-57 dBm	-52 dBm
Table	-53 dBm	-52 dBm
Door	-47 dBm	-45 dBm
Path	-47 dBm	-47 dBm

Measurements were done for two types of conditions: single-user without any additional traffic on the network, and multi-user where two devices are connected to the network and either sends or receives 10 Mbit/s UDP traffic using iperf3. The frequency channel 100 in the 5 GHz band was used with a bandwidth of 20 MHz.

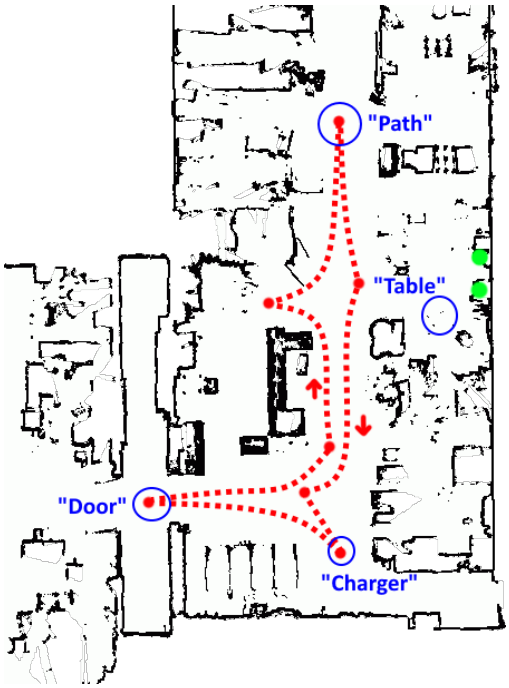

Figure 3.15: Route for intra-AP measurements, shown in red. The location for static measurements are also highlighted in blue. Note that the door separating the two lab environments was closed for these tests.

Results

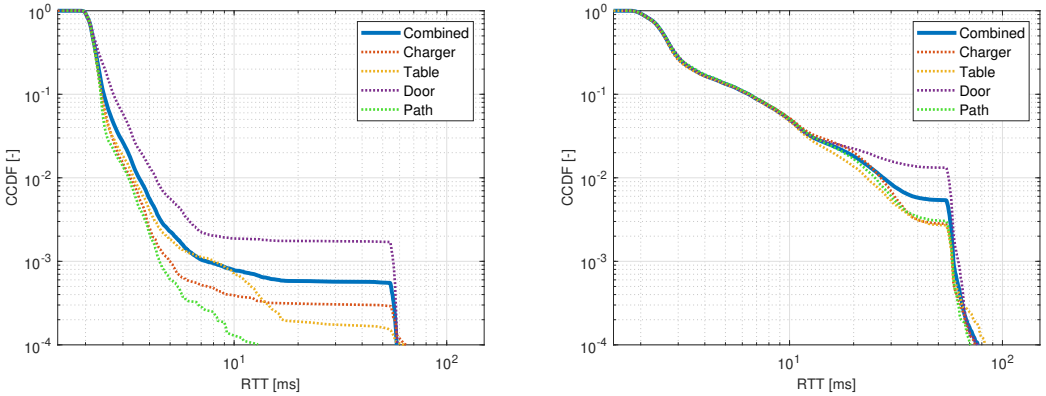

(a) Single-user network.

(b) With background traffic.

Figure 3.16: CCDF for static locations.

Table 3.8: Summary of RTT latency and packet error rate measurement results for the static test setups.

Setup		Min	Avg	99.9%-ile	Jitter	PER
Combined	Idle network	1.5 ms	2.3 ms	7.4 ms	0.3 ms	0%
	With traffic	1.5 ms	3.9 ms	58.8 ms	2.9 ms	0%
Charger	Idle network	1.5 ms	2.2 ms	5.0 ms	0.2 ms	0%
	With traffic	1.7 ms	3.9 ms	57.9 ms	2.8 ms	0%
Table	Idle network	1.7 ms	2.2 ms	8.0 ms	0.2 ms	0%
	With traffic	1.5 ms	3.8 ms	58.1 ms	2.5 ms	0%
Door	Idle network	1.6 ms	2.4 ms	55.9 ms	0.6 ms	0%
	With traffic	1.7 ms	4.2 ms	62.0 ms	3.5 ms	0%
Path	Idle network	1.6 ms	2.2 ms	4.4 ms	0.2 ms	0%
	With traffic	1.7 ms	3.9 ms	58.1 ms	2.7 ms	0%

Figure 3.16a illustrates the latency performance of the four static deployments, with key statistics and PER highlighted in Table 3.7. Three of the configurations have similar performance to $4 \cdot 10^{-3}$. The fourth measurement at the door is however higher than the rest throughout the percentiles, which can be attributed to the location being nearby a reflective surface, but the cause has not been investigated further. All latencies are however contained within 60 ms. The drop just after 50 ms is suspected to be due to a packet batching mechanism at the APs, where in this case two ping-responses are sent in the same downlink transmission, such that one of the responses experience the normal RTT latency plus an additional 50 ms from the ping interval. Since no settings reflecting this aspect has not been found in the Cisco Meraki Dashboard for the APs, an approach to change this property has not been found with the current deployment. When background traffic is present, the latency increases down to the 90%-ile, after which the latency reaches ∼55 ms as previously described.

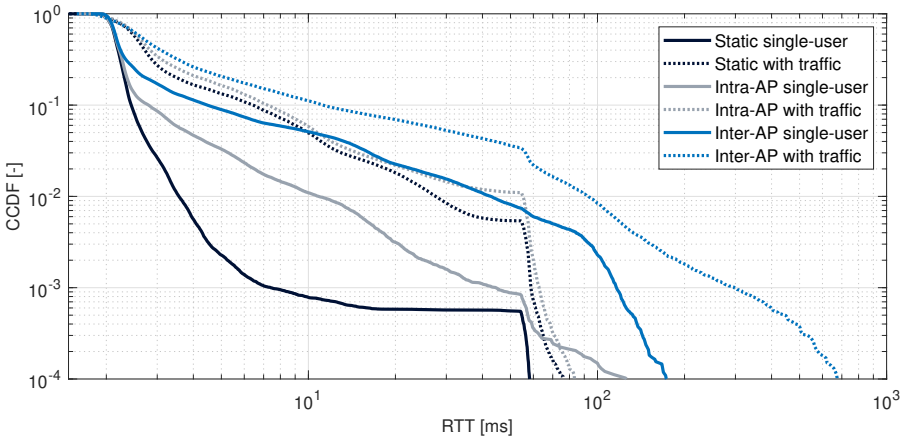

Figure 3.17: CCDF for static, intra-AP and inter-AP conditions.

Table 3.9: Summary of RTT latency and packet error rate measurement results for the static, inter-AP and intra-AP configurations.

Setup		**Min**	**Avg**	**99.9%-ile**	**Jitter**	**PER**
Static	Idle network	1.5 ms	2.3 ms	7.4 ms	0.3 ms	0%
	With traffic	1.5 ms	3.9 ms	58.8 ms	2.9 ms	0%
Intra-AP	Idle network	1.6 ms	2.6 ms	44.4 ms	0.8 ms	0%
	With traffic	1.6 ms	4.4 ms	62.7 ms	3.5 ms	0%
Inter-AP	Idle network	1.8 ms	3.9 ms	116.0 ms	2.4 ms	0.043%
	With traffic	1.6 ms	7.8 ms	297.0 ms	5.9 ms	0.066%

The CCDF shown in Figure 3.17 illustrates the RTT performance for static and dynamic cases, where the intra-AP mobility as presented in the test setup is shown. The inter-AP performance is also shown for reference, from the test journal in Section 3.4.3. The latency of intra-AP increases exponentially from the 90%-ile throughout the percentiles for idle networks, which confirms that the latency does indeed increase when mobility is introduced. This latency is further increased when considering the inter-AP mobility, however due to the presence of handovers and different environments (i.e. the full testing environment is used as opposed to a single part), these two cannot be directly compared. Nonetheless, the similarity between the intra-AP and inter-AP latency distributions remains notable, as this reveals that the act of performing handovers is not the sole reason for worsened performance compared to static conditions.

Conclusion

When a STA is mobile, i.e. non-static, the latency increases as compared to idle cases, especially for the 99.9%-ile. Based on a comparison between intra-AP and inter-AP mobility, it is suspected that the handover itself is not the sole cause of this increase in latency.

3.4.2 Wi-Fi Handover Tests

Purpose of the Test

This test aims to determine where the time is spent during each handover with regards to the information provided by wpa_supplicant. Another goal with the test is to obtain insight as to when handovers occur and what caused them. Finally, the effect of using RSSI-initiated roaming instead of wpa_supplicant-initiated roaming is to be determined.

Theory

When a Wi-Fi device moves out of the coverage area of an access point (AP), it will eventually reach a point where communication becomes impossible. The detection of this can be done through several means, such as when the beacon frames from the AP is not detected, if a large amount of packet loss is observed, or if the RSSI reaches a certain threshold. When this happens, the device will begin to search for new APs to which it can establish a new connection to.

Roaming from one AP to another is known as a handover, as described in Section 3.1.3. Based on experience described in 3.2.3, it is relevant to determine what happens when handovers are initiated from wpa_supplicant instead of relying on a timeout mechanism.

Test Setup

- Intel NUC equipped with Intel AX200 Wi-Fi 6
- 3x Cisco MR36 Wi-Fi 6 access points
- Edge-server connected to the network through Ethernet
- MiR robot for movement

The test was conducted in the AAU Smart Production Lab shown in appendix A. Two scenarios were considered: supplicant-triggered and RSSI-triggered roaming. Likewise, the improvements of scanning fewer channels and using IEEE 802.11r was investigated. To obtain the desired data, the NUC was configured to ping a target device connected to the network using Ethernet, so that we only consider the latency introduced from the Wi-Fi connection between the NUC and the AP. Measurements relating to the handover itself was obtained from the wpa_supplicant log, as this contains information and timestamps for each stage.

Scenario 1 - Supplicant-triggered roaming
Two robots are with identical setups were programmed to drive in a fixed path, shown in red on Figure 3.18. The robots were offset such that they do not perform handovers simultaneously. To allow for the supplicant to properly register the RSSI, the devices were configured to ping another device on the network connected through Ethernet, with 64 B every 50 ms. While this traffic may affect the messages when a device roams to a different network, it is not believed to have any notable impact. This will however also allow for measuring the effect on latency.

Scenario 2 - RSSI-triggered roaming
A single robot was used for this test. It was configured to follow another path, shown in green on Figure 3.18. The network device would at predetermined locations perform a scan, which would then trigger a roaming event as another, more suitable AP is nearby.

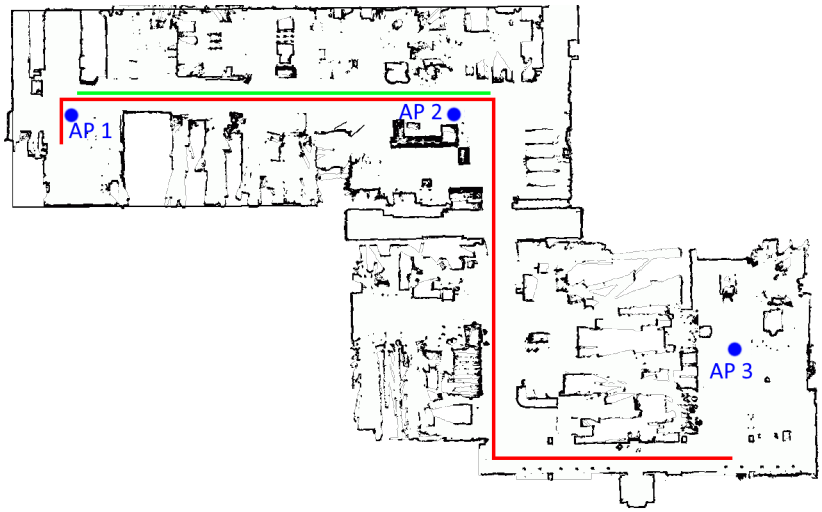

Figure 3.18: Routes for the tests.

Results

Figure 3.19 illustrates the correlation between RSSI and latency. It is here seen that spikes in latency occurs during handovers, with supplicant-triggered handovers resulting in latencies exceeding over 1000 ms. While this only occurs once for the shown data in Figure 3.19a, it was observed from the logs of wpa_supplicant that the device would detect an absence of AP beacon frames, and thus initiate the handover. This resulted in the communication being blocked until the device had scanned and successfully established a connection with a new AP.

For RSSI-triggered handovers, the latency was bound to 1000 ms for the handovers. This was in found to be due to a link-level routing issue, with an additional second of

latency being added, but does not affect the results of this test. When a handover is triggered, a small period of degraded performance is observed as the device performs the scanning for new APs. Here, latencies of ∼100 ms are observed, but as communication is still possible, it is a better alternative to the supplicant-triggered handover.

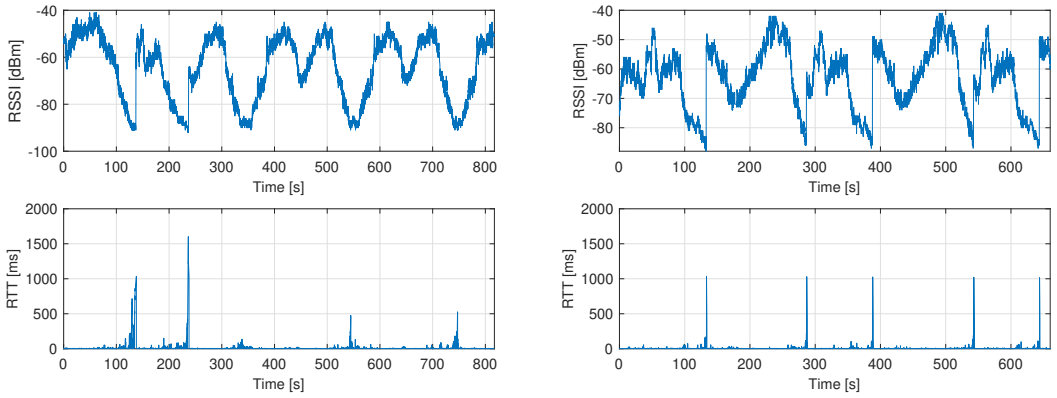

(a) Supplicant-triggered

(b) RSSI-triggered

Figure 3.19: RSSI and latency mobility measurements.

It was found that the device was unable to perform the IEEE 802.11r over-the-air roaming when utilizing supplicant-triggered handovers due to the reasons described in Section 3.2.3. The performance of the RSSI-triggered handover is therefore used going forward.

Draft Figure 3.20 shows that the scanning stage takes a substantial amount of time compared to the other stages. This explains the increased latency during handovers for supplicant-triggered handovers, as this stage would also be without any connection to any AP. It is also seen that the IEEE 802.11r protocol is functioning as expected as the 4-way handshake is absent. Using optimized scanning, i.e. only scanning for 3 channels compared to 38, yields significantly reduced time required for this stage, which indicates that the scanning time is indeed proportional to the number of channels scanned. The final version of the figure is found in Paper 1.

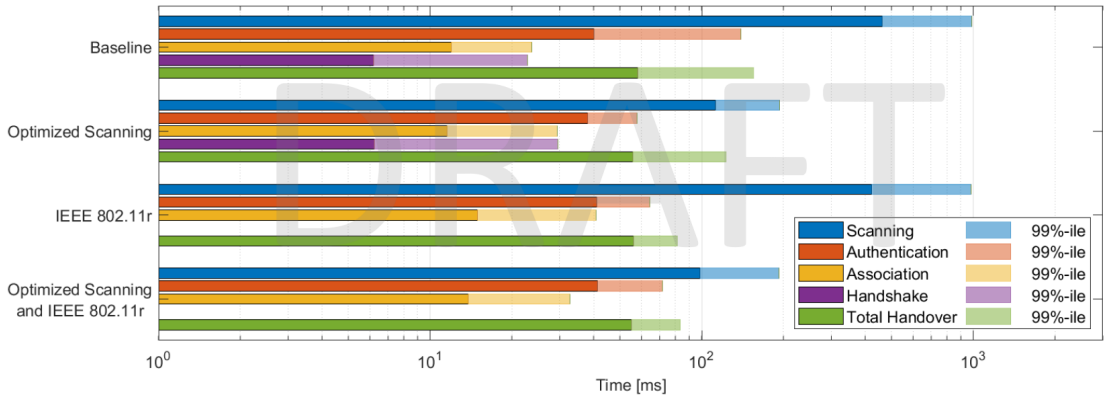

Figure 3.20: Duration of each stage in handovers through four different configurations.

Conclusion

Using supplicant-triggered handovers results in large latencies of up towards 2 seconds, while RSSI-triggered handovers have much lower spikes in latency. This is due to the device losing connection completely for the first form of handover, resulting in the communication interruption also containing the scanning period, which is significantly longer than the handover itself. For RSSI-triggered handovers, this scanning period instead results in degraded performance, but still allows for communication to occur.

For the handover, the authentication and association stage are of similar time duration, with the handshake being slightly shorter. Using optimized scanning greatly reduces the scanning period, and using IEEE 802.11r over-the-air roaming removes the handover aspect completely without any notable impact on the other stages.

3.4.3 Wi-Fi Latency Tests

Purpose of the Test

This test aims to determine the impact on the latency to the network when utilizing optimized scanning and IEEE 802.11r. Furthermore, it is also relevant to determine the impact when introducing additional interference on the same channel.

Theory

When utilizing Wi-Fi with mobility, the latency will change based on the environment, the amount of interference from other devices or machinery, as well as other factors. Moreover, when roaming from one network to another, the impact of the handover between these are of interest. The impact of this is described in Section 3.1.2.

Test Setup

- Intel NUC equipped with Intel AX200 Wi-Fi 6
- 6x Wi-Fi 6-enabled devices for transmitting background data
- Cisco MR36 Wi-Fi 6 access point
- Edge-server connected to the network through Ethernet
- MiR robot for movement and positioning data

The test was conducted in the AAU Smart Production Lab shown in appendix A. The robot was programmed to follow a specific route, as shown in Figure 3.21. The device performs a handover every two minutes on average using this route. The channels 100, 108 and 116 in the 5 GHz band were used for AP 2, 1 and 3 respectively, all at a bandwidth of 20 MHz. The APs were furthermore configured to use 15 dBm transmit power. Finally, the network will utilize "Cisco Client Balancing", which utilizes IEEE 802.11v and additional proprietary mechanisms to steer clients to other APs depending on the network load.

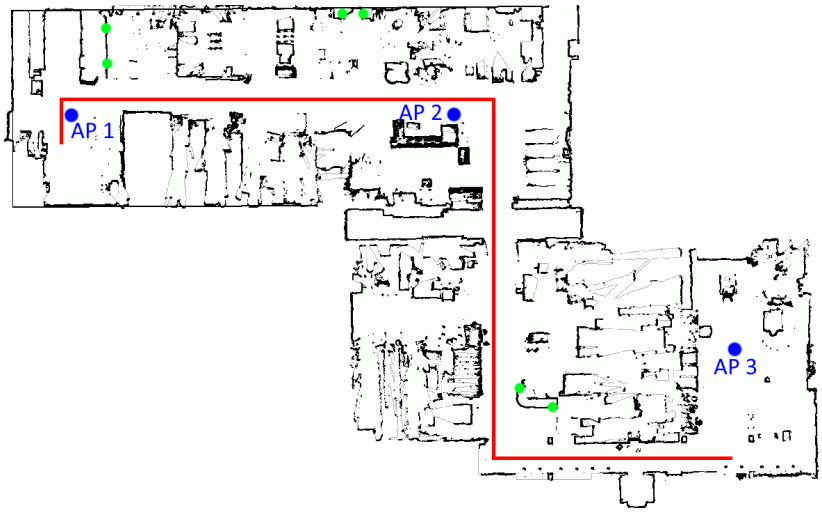

Figure 3.21: Route used during tests. The green markers indicate the location of the devices sending and receiving background traffic.

The device will be configured using the measurement software described in Section 3.2.1. It will thus be able to obtain its location directly from the MiR robot using Ethernet. For Wi-Fi, the device was configured to connect to any of the three Cisco APs. When the device detects an RSSI of -85 or lower, it will trigger a scan which will then initiate the handover. Some other Wi-Fi-related parameters was utilized as described below. Finally, the Linux ping functionality was used to determine the latency, by pinging a computer connected to the network using Ethernet. A ping interval of 50 ms was used, with a packet size of 64 Bytes.

A test was done for baseline, for when optimized scanning is enabled (reducing the number of channels to scan from 38 to 3), when IEEE 802.11r is enabled, and with both optimizations enabled. These four tests were also done when introducing background traffic. Two devices were connected to each AP, with one device transmitting 10 Mbit/s uplink, and the other receiving 10 Mbit/s downlink. Both of these used iperf3 with UDP traffic. Finally, the frequency channels of the APs were configured to be using the same channel, further increasing the interference as the devices will now be on the same channel.

Results

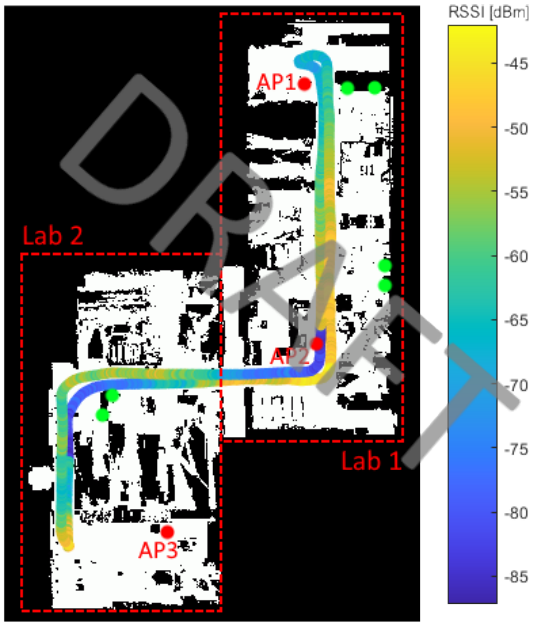

Figure 3.22: RSSI heatmap of a single lap for an idle network. The APs and traffic-generating STAs are shown in red and green, respectively.

A draft of the heatmap of the RSSI in idle conditions can be seen in Figure 3.22, with the final version being found in Paper 1. It is here seen that the STA did not roam to AP 1 since the RSSI was ∼70 dBm. This was however not the case for all of the measurements, where the STA would occasionally roam to it when interference was present. The heatmap further shows that there is a clear overlap in terms of coverage area between AP 2 and 3. Although these regions of handover can be better planned by changing the location of the APs or changing their transmitting power, this was not the focus of the measurements. Nonetheless, the heatmap helps us to determine the handover regions are located, with most handovers found to be in the same area.

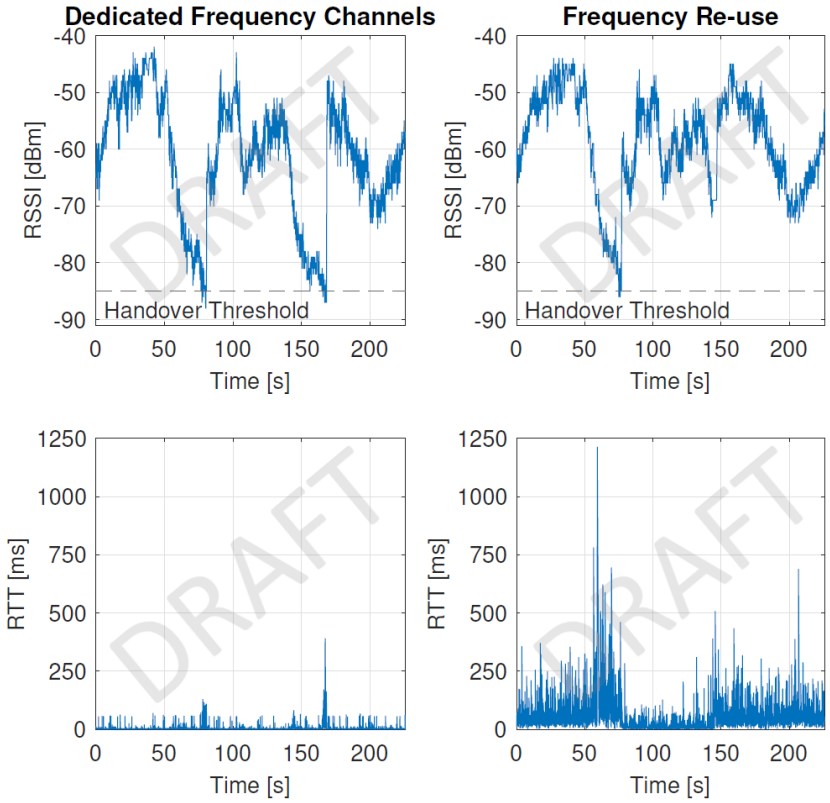

Figure 3.23: RSSI and RTT for a single lap with a loaded network.

The draft in Figure 3.23 illustrates the correlation between RSSI and RTT for a single lap starting from AP 2 and moving towards AP 3. The final version of the figure can be found in Paper 1. It is here seen that the STA will initiate the scanning process once -85 dBm is reached for the multi-channel configuration. During this phase, there is a slight increase in latency due to the scanning, followed by a large spike from the handover itself. When using a single-channel configuration where the three APs overlap, the overall latency is much higher with exception of AP 3. This is due to second region, Lab 2, being separated from Lab 1 by a brick wall, blocking a significant part of the interference. When the STA then moves back to Lab 1, we see a handover occurring earlier than for low-interference conditions at an RSSI of -73 dBm. This is due to the handover being timeout-triggered instead of being triggered by the RSSI. The reason that we do not see a spike in terms of RTT during this handover is due to the impact of the interference being much more significant than the handover itself.

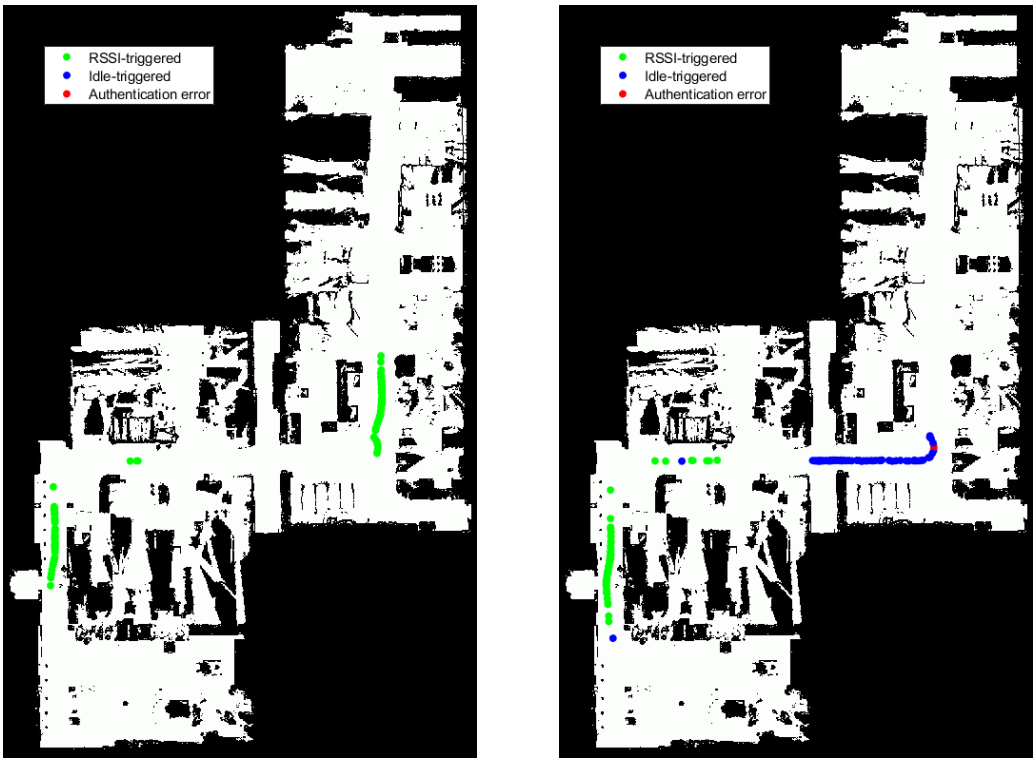

(a) Low interference **(b)** High interference

Figure 3.24: Locations of handover-events with the type of handover and errors associated with the process.

It can from Figure 3.24a be seen that when low interference is present (10 Mbit/s uplink and downlink UDP traffic), all handovers are triggered based on the RSSI. It is further seen that the location of the handovers is generally clustered in the same area, but not at an exact point. When introducing frequency re-use, seen in Figure 3.24b, all of the handovers in one of the handover regions are caused by a timeout/idle-trigger, with one of the handovers failing altogether. This is likely due to the combined traffic of all three networks in the middle area, causing a significant impact on both packet reliability and latency.

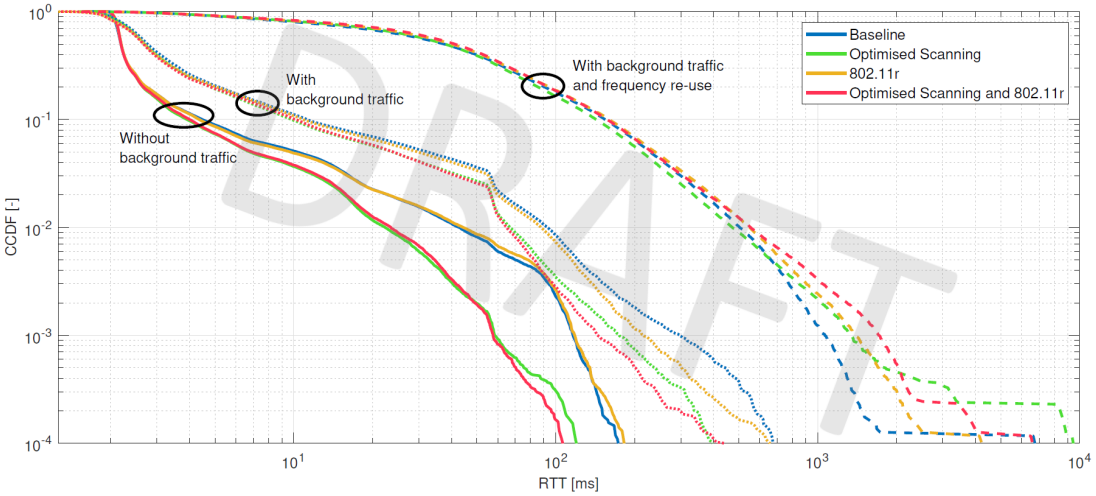

Figure 3.25: RTT measurements for Wi-Fi with mobility in various conditions.

Table 3.10: Summary of RTT and PER measurement results for the different Wi-Fi schemes and network configuration setups.

Test Setups		Min	Avg	99.9%-ile	Jitter	PER
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%

It can be seen in the draft Figure 3.25 and Table 3.10 that the overall latency distribution is highly affected by the amount of traffic present in the network due to the Listen-Before-Talk mechanism. It is however also seen that improving aspects related to the handover will result in improved latency after the 90%-ile. This is especially evident by optimizing 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 load on the network. 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 802.1X. The final versions of the figure and table can be found in Paper 1.

If frequency re-use is utilized for all APs, it is evident that the performance is severely affected. It can likewise be seen that the aforementioned improvements to the handover

will not have any effect on the latency. The source of the increased latency comes from general interference, but also due to the roaming being triggered by a timeout-mechanism as seen in Figure 3.23.

The (PER) was not found to have been affected by the handover improvements, but is instead proportional to the amount of traffic present on the network.

Conclusion

The latency for Wi-Fi at can be decreased by applying radio tuning and IEEE 802.11 amendments targeting the handover events. By optimizing the number of channels to scan to match that of the network environment, corresponding to a decrease from 36 to 3 for this test, the scanning period can be decreased from 0.5-1 to 0.1-0.3 seconds. Using IEEE 802.11r reduces the handover itself by skipping the handshake stage, of which the impact is most notable when traffic is present on the network.

3.4.4 Multi-connectivity Functionality Tests

Purpose of the Test

This test aims to verify the behavior of the two packet scheduling schemes: packet duplication and best path scheduling. Furthermore, the functionality of the mobility coordinator can be validated through this test.

Theory

The design and implementation of the packet schedulers and mobility coordinator is presented in Section 3.3.

Using the packet duplication scheme, both STAs should transmit and receive data independently. With all traffic being duplicated equally, i.e. all packets are duplicated, it is expected that similar throughput is achieved. When a STA disconnects from an AP and remains in a disconnected state, the throughput should drop to 0 until a connection is reestablished.

The best path scheduling scheme utilizes the STA with highest RSSI, and switches between the two STAs when the secondary exceeds the primary by a margin of 5 dB. The traffic is only routed through the primary STA, but both STAs can still receive data if they are connected to an AP. The gateway on the network-side only transmits data to the STA that it had last received data from to ensure the data is not duplicated in either uplink or downlink.

The mobility coordinator aims to ensure the STAs connect to different APs at all times through a blacklisting mechanism using the Basic Service Set IDs (BSSIDs) of the APs. Furthermore, it initiates the roaming events for the STAs and avoids simultaneous handovers on both STAs. For the best path scheduling, it will only initiate roaming events for the secondary (inactive) STA.

Test Setup

- Gateworks PC equipped with 2x Intel AX200 Wi-Fi 6
- 3x Cisco MR36 Wi-Fi 6 access points
- Edge-server connected to the network through Ethernet
- MiR robot for movement

The test was conducted in the AAU Smart Production Lab shown in appendix A. To enable mobility in the setup, an Intel NUC and the Gateworks PC was placed on top of an AMR. This AMR was programmed to follow a specific route throughout the test environment, shown in Figure 3.26. The figure likewise illustrates the coverage areas of each AP in the test environment. The channels 132, 136 and 140 in the 5 GHz band were used for AP 2, 1 and 3 respectively, all at a bandwidth of 20 MHz. The APs were furthermore configured to use 15 dBm transmit power.

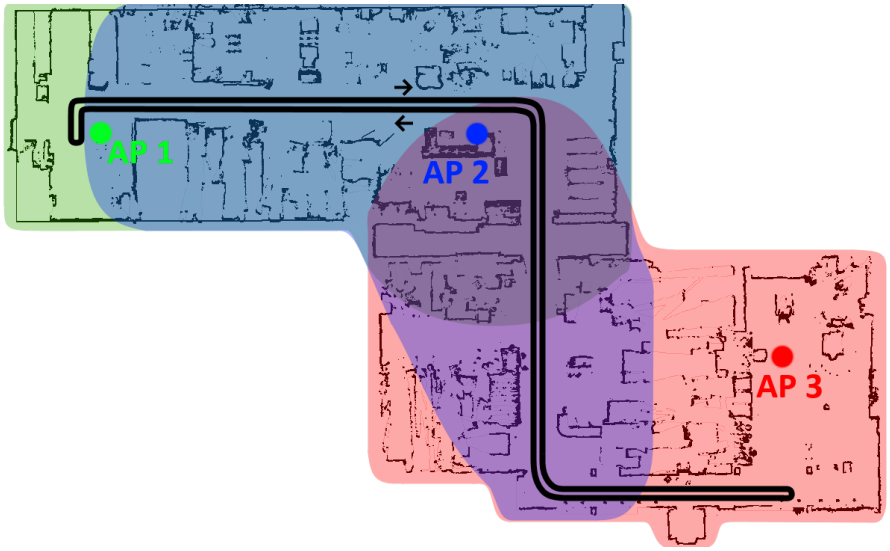

Figure 3.26: Floor plan of the test environment, including AP locations (green, blue, and red dots) and coverage areas as well as the AMR measurement route.

The Gateworks PC was running the measurement software detailed in Section 3.2.1, logging Wi-Fi statistics, positioning (from the robot through Ethernet) and traffic. To generate traffic, a ping with a packet size of 64 B and 50 ms (20 packets per second) was generated from the NUC on the AMR, transmitted to another device on the network-side through the gateway. The two packet scheduling schemes were tested with the gateway, packet scheduler and mobility coordinator configured accordingly.

Results

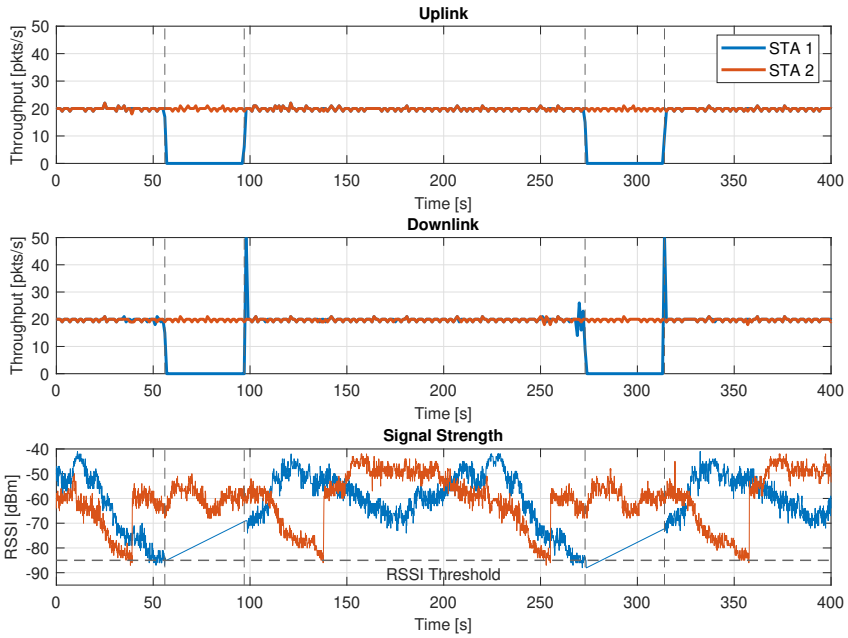

Figure 3.27: Throughput and RSSI of the packet duplication scheme using STA 1 and STA 2 highlighted in blue and orange, respectively. The vertical lines indicate the handovers for STA 1.

Figure 3.27 illustrates the throughput of each STA, which confirms that both STAs are being used for communication during the packet duplication scheme. It can however also be observed that STA 1 remains disconnected for a considerable amount of time, during which only STA 2 is used. When STA 1 reconnects, a significant spike in downlink traffic occurs, which stem from the gateway on the network still transmitting data to the disconnected STA, which gets queued at the APs.

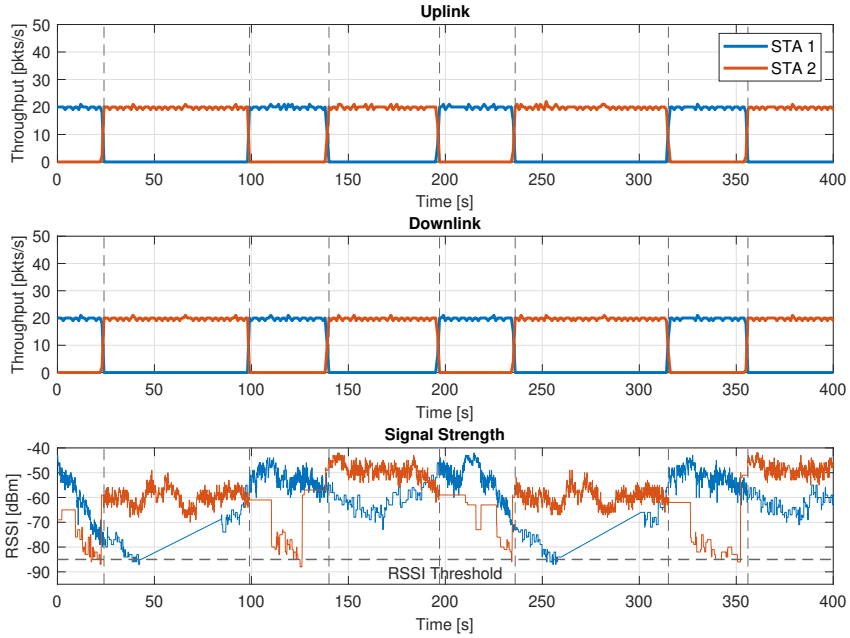

Figure 3.28: Throughput and RSSI of the best path scheduling scheme using STA 1 and STA 2 highlighted in blue and orange, respectively.

The measurements illustrated in Figure 3.28 likewise confirms that the best path scheduling scheme does indeed switch between the two STAs depending on the RSSI, and likewise confirms that the network-side gateway does not duplicate traffic and adapts correctly to the received traffic. Furthermore, the RSSI measurements indicate that the mobility coordinator does indeed initiate scans on the secondary STA, and leaves the primary connected. Note that the low resolution of RSSI for the secondary stems from it not receiving any traffic, and thus the RSSI is not updated as regularly.

Conclusion

Both the packet scheduler and mobility coordinator functions correctly, being able to both steer the traffic depending on the network conditions and being able to introduce coordination between the two STAs.

3.4.5 Multi-connectivity Latency Tests

Purpose of the Test

This test aims to determine the latency and packet reliability impact of using the radio-aware multi-connectivity configurations. Due to similar test setups, the results will be comparable to the results from the single-connectivity test in Section 3.4.3

Theory

Using the radio-aware multi-connectivity layer-4 packet scheduler, detailed in section 3.3, is likely to show improvements in terms of latency and packet reliability over a STA using single-connectivity. Using packet duplication will increase the probability of a packet arriving, but may be subject to the same restrictions in terms of medium access because of the Listen-Before-Talk mechanism. Using best path scheduling will mitigate the latency impact of handovers, thus reducing the latency at lower percentiles.

Test Setup

- Gateworks PC equipped with 2x Intel AX200 Wi-Fi 6
- 6x Wi-Fi 6-enabled devices for transmitting background data
- 3x Cisco MR36 Wi-Fi 6 access points
- Edge-server connected to the network through Ethernet
- MiR robot for movement and positioning data

The test was conducted in the AAU Smart Production Lab shown in appendix A. To enable mobility in the setup, an Intel NUC and the Gateworks PC was placed on top of an AMR, shown in the draft Figure 3.29. The robot was programmed to follow a specific route, shown in Figure 3.30. If the mobility coordinator was not used, the STA would perform a handover every 2 minutes on average using this route (for more details see Section 3.4.3). The channels 132, 136 and 140 in the 5 GHz band were used for AP 2, 1 and 3 respectively, all at a bandwidth of 20 MHz. The APs were furthermore configured to use 15 dBm transmit power. Finally, the network will utilize "Cisco Client Balancing", which utilizes IEEE802.11v and additional proprietary mechanisms to steer clients to other APs depending on the network load.

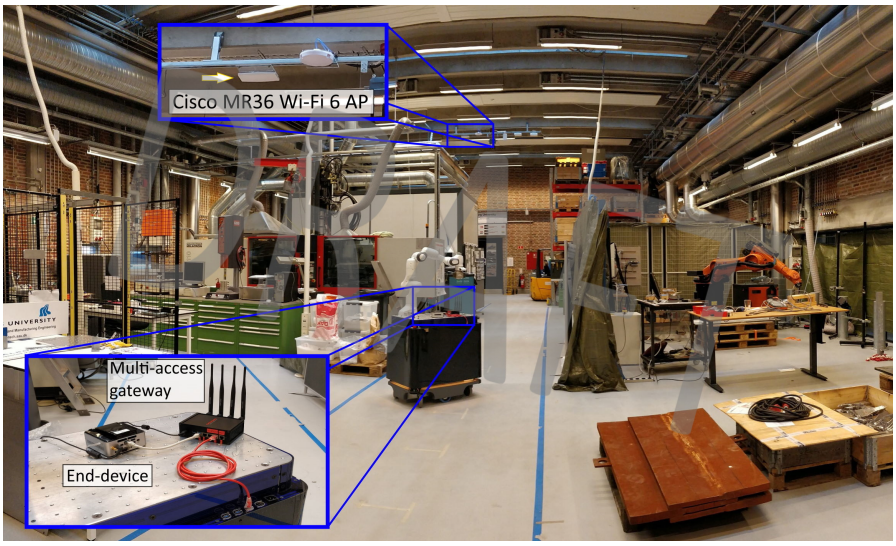

Figure 3.29: Test setup with the Intel NUC and Gateworks set up on the MiR robot. Final figure can be found in Paper 1.

Figure 3.30: Route used during tests. The green markers indicate the location of the devices sending and receiving background traffic.

The Gateworks PC was running the measurement software detailed in Section 3.2.1, logging Wi-Fi statistics and positioning (from the robot through Ethernet). It will moreover use the gateway software described in Section 3.3.1 using both Wi-Fi STAs (depending on the multi-connectivity configuration). An Intel NUC will likewise run the gateway software, but will only transmit data through a single Ethernet cable connected to the APs. The multi-connectivity configurations are presented in Section 3.3. The Gateworks PC will finally be configured to only scan the three utilized channels, as to

compare with the previously measured best-case single-connectivity Wi-Fi. To measure the latency and packet error rate, a NUC connected to either gateway pinged through the gateway software at an interval of 50 ms and packet size of 64 Bytes.

The following configurations were investigated: 1) best path scheduling with coordination, 2) packet duplication without coordination, 3) packet duplication with coordination. All three configurations were tested with and without background traffic. Configurations 2 and 3 were likewise tested with background traffic and using the same frequency channel across all APs, referred to as frequency re-use. The background traffic consisted of 10 Mbit/s uplink and downlink on each AP using the two Intel NUCs (i.e. 2 NUCs per AP), with one receiving and the other transmitting iperf3 UDP traffic.

Results

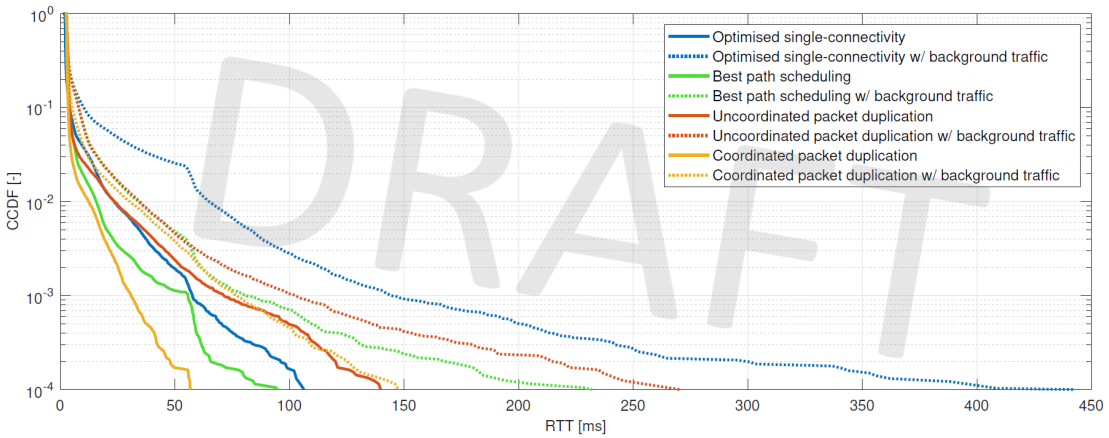

Figure 3.31: RTT Empirical CCDFs for single- and multi-connectivity configurations.

Table 3.11: Summary of RTT and PER measurement results for the single- and multi-connectivity configurations.

Test			Min	Avg	99.9%-ile	Jitter	PER
Configuration	**Coordination scheme**	**Channel condition**					
Optimised single-connectivity	-	Idle network	1.48 ms	3.23 ms	57.60 ms	1.56 ms	0.044%
		Background traffic	1.29 ms	6.26 ms	143.00 ms	4.78 ms	0.065%
Best path scheduling	Coordinated	Idle network	2.23 ms	3.36 ms	55.60 ms	1.00 ms	0.005%
		Background traffic	2.24 ms	4.80 ms	80.70 ms	2.94 ms	0.028%
Packet duplication	Uncoordinated	Idle network	2.32 ms	3.88 ms	71.10 ms	1.42 ms	0.018%
		Background traffic	1.99 ms	4.87 ms	102.00 ms	2.66 ms	0.041%
	Coordinated	Idle network	2.14 ms	3.15 ms	30.80 ms	0.58 ms	0%
		Background traffic	2.06 ms	4.12 ms	77.10 ms	1.75 ms	0.001%

The data from draft Figure 3.31 and Table 3.11 suggests that by having two STAs available for a device and using best path scheduling, significant latency improvements

can be observed for the 95%-ile to the 99.9%-ile. Although taking the RSSI of the STAs into consideration to maintain a more robust connection, the main improvements stem from the multi-STA being able to utilize the two radios and fully mitigate the latency impact of handovers. The PER has likewise been decreased by 0.04 percent point. When traffic is introduced to the network channel, the increase in latency is significantly lower than that of the single-STA configuration.

When using packet duplication without coordination between the STAs, the latency performance is similar to that of single-connectivity Wi-Fi up until the 99.9%-ile. The latency is however improved compared to single-connectivity when traffic is introduced to the network, however without considerable improvements to the PER. However, by duplicating the traffic over two separate APs, the latency performance is substantially improved, with a reduction in the 99.9%ile latency of 27 ms for idle networks, and 66 ms for loaded networks. Packet losses are also minimized, with only a single packet drop out of 100.000 measured for the loaded network. The final version of the figure can be found in Paper 2.

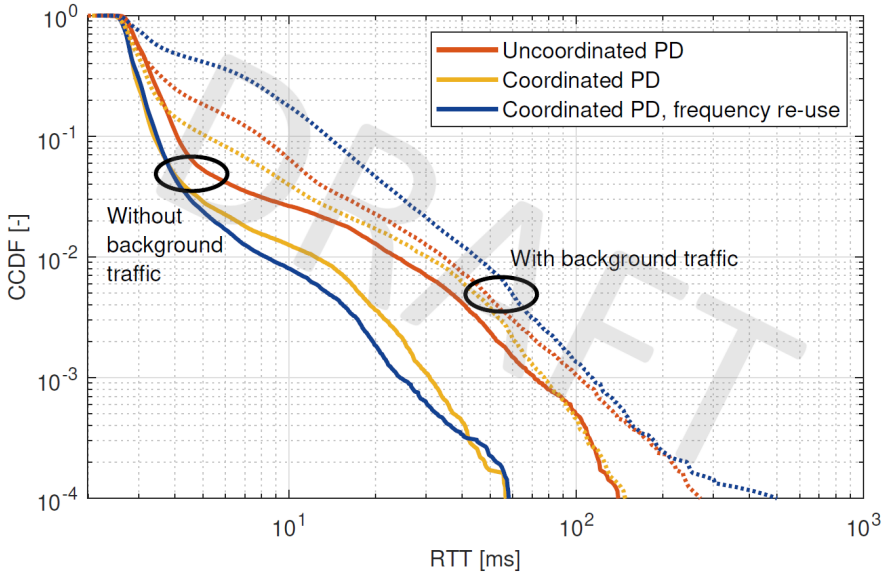

Figure 3.32: Packet duplication (PD) RTT performance with dedicated channels and using frequency re-use.

The performance of the packet duplication configuration for channel conditions with and without dedicated frequency channels is illustrated in draft Figure 3.32. If it is not possible to have dedicated frequency channels for each AP and avoid overlap for APs with same channel, the performance for the packet duplication scheme will be slightly degraded if background traffic is present. The final version of the figure can be found in Paper 2.

Conclusion

Using multi-connectivity further increases the performance of Wi-Fi. Using the single path scheduling mechanism increases the latency performance and decreases the PER by using the AP with highest measured RSSI and by mitigating the performance impact of handover events. Packet duplication likewise improves the performance, with a reduction of 99.9%-ile RTT from 58 ms to 31 ms from single-STA configurations. However, purely duplicating traffic without a mobility coordinator, which is commonly done in state-of-the-art, will increase the latency at lower percentiles compared to single-STA. When traffic is present on the network, both packet schedulers with mobility coordination yielded similar performance down to the 99.9%-ile, although packet duplication resulted in close-to-no packet losses.

Bibliography

[1] RF Wireless World, "Difference between wifi 6 and wifi 5, wifi 4, wifi 3, wifi 2, wifi 1," 2021. [Online]. Available: https://wiki.debian.org/NetworkManager

[2] X. Chu and Y. Yan, "Performance evaluation of ieee 802.11 infrastructure mode with intra-cell udp traffic," in *2007 Second International Conference on Communications and Networking in China*, 2007, pp. 893–898.

[3] J. H. Lee, M.-S. Park, and S. C. Shah, "Wi-fi direct based mobile ad hoc network," in *2017 2nd International Conference on Computer and Communication Systems (ICCCS)*, 2017, pp. 116–120.

[4] G. Naik, J. Liu, and J.-M. J. Park, "Coexistence of Wireless Technologies in the 5 GHz Bands: A Survey of Existing Solutions and a Roadmap for Future Research," *IEEE Communications Surveys Tutorials*, vol. 20, no. 3, pp. 1777–1798, 2018.

[5] P. Dhere, P. Chilveri, R. Vatti, V. Iyer, and K. Jagdale, "Wireless Signal Strength Analysis in a Home Network," in *2018 International Conference on Current Trends towards Converging Technologies (ICCTCT)*, 2018, pp. 1–5.

[6] A. Bhattacharya and A. Kumar, "Analytical modeling of ieee 802.11-type csma/ca networks with short term unfairness," *IEEE/ACM Transactions on Networking*, vol. 25, no. 6, pp. 3455–3472, 2017.

[7] Fraida Fund, "Understanding the 802.11 Wireless LAN MAC frame format," 2017. [Online]. Available: https://witestlab.poly.edu/blog/802-11-wireless-lan-2/

[8] J. Hintersteiner, "Wi-Fi Beacon Frames Simplified," 2021. [Online]. Available: https://www.engeniustech.com/wi-fi-beacon-frames-simplified/

[9] "802.11 Beacon Intervals - The Real Story." [Online]. Available: https://www.cwnp.com/cwnp-wifi-blog/80211-beacon-intervals/

[10] Intel, "Understanding IEEE 802.11 Authentication and Association," 2021. [Online]. Available: https://www.intel.com/content/www/us/en/support/articles/000006508/network-and-i-o/wireless.html

[11] J. Chen, M. Jiang, and Y. Liu, "Wireless LAN security and IEEE 802.11i," *IEEE Wireless Communications*, vol. 12, no. 1, pp. 27–36, 2005.

[12] Cisco, "802.11k Assisted Roaming," 2020. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-1/Enterprise-Mobility-8-1-Design-Guide/Enterprise_Mobility_8-1_Deployment_Guide/Chapter-11.html

[13] 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.

[14] A. A. Tabassam, H. Trsek, S. Heiss, and J. Jasperneite, "Fast and seamless handover for secure mobile industrial applications with 802.11r," in *2009 IEEE 34th Conference on Local Computer Networks*, Oct. 2009, pp. 750–757, iSSN: 0742-1303.

[15] Cisco, "802.11r BSS Fast Transition Deployment Guide," 2016. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/80211r-ft/b-80211r-dg.html

[16] Debian Wiki, "NetworkManager," 2021. [Online]. Available: https://wiki.debian.org/NetworkManager

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

[18] "iw Documentation." [Online]. Available: https://wireless.wiki.kernel.org/en/users/documentation/iw

[19] J. Malinen *et al.*, "wpa_supplicant sample configuration," 2019. [Online]. Available: https://w1.fi/cgit/hostap/plain/wpa_supplicant/wpa_supplicant.conf

[20] "wpa_supplicant Control Interface." [Online]. Available: https://w1.fi/wpa_supplicant/devel/ctrl_iface_page.html

[21] "Freedesktop Wiki - libqmi." [Online]. Available: https://www.freedesktop.org/wiki/Software/libqmi/

[22] "Mobile Signal Strength Recommendations." [Online]. Available: https://wiki.teltonika-networks.com/view/Mobile_Signal_Strength_Recommendations

[23] "tcpdump Official Webpage." [Online]. Available: https://www.tcpdump.org/

[24] MiR, "Mir100 specification," 2021. [Online]. Available: https://www.mobile-industrial-robots.com/en/solutions/robots/mir100/

[25] Cisco, "Enterprise Mobility 8.1 Design Guide," 2020. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-1/Enterprise-Mobility-8-1-Design-Guide/Enterprise_Mobility_8-1_Deployment_Guide/Chapter-11.html

[26] J. Hintersteiner, "WiFi Fast Roaming, Simplified," 2016. [Online]. Available: https://www.networkcomputing.com/wireless-infrastructure/wifi-fast-roaming-simplified

[27] "Freeradius." [Online]. Available: https://freeradius.org/

[28] Intel, “iwlwifi,” 2020. [Online]. Available: https://wireless.wiki.kernel.org/en/users/drivers/iwlwifi

[29] “Finally got 802.11r roaming working,” 2019. [Online]. Available: https://www.reddit.com/r/openwrt/comments/515oea/finally_got_80211r_roaming_working/

[30] Intel, “Linux Wireless Documentation,” 2020. [Online]. Available: https://wireless.wiki.kernel.org/en/users/documentation

[31] A. E. Fink, “Support post on Intel Community Forum,” 2020. [Online]. Available: https://community.intel.com/t5/Wireless/Problem-using-802-11r-FT-on-Ubuntu-AX200/m-p/1233573

[32] 7signal, “Mysteries of Wi-Fi Roaming Revealed,” 2020. [Online]. Available: https://go.7signal.com/download-wifi-roaming-whitepaper

Appendix A

AAU 5G Smart Production Lab

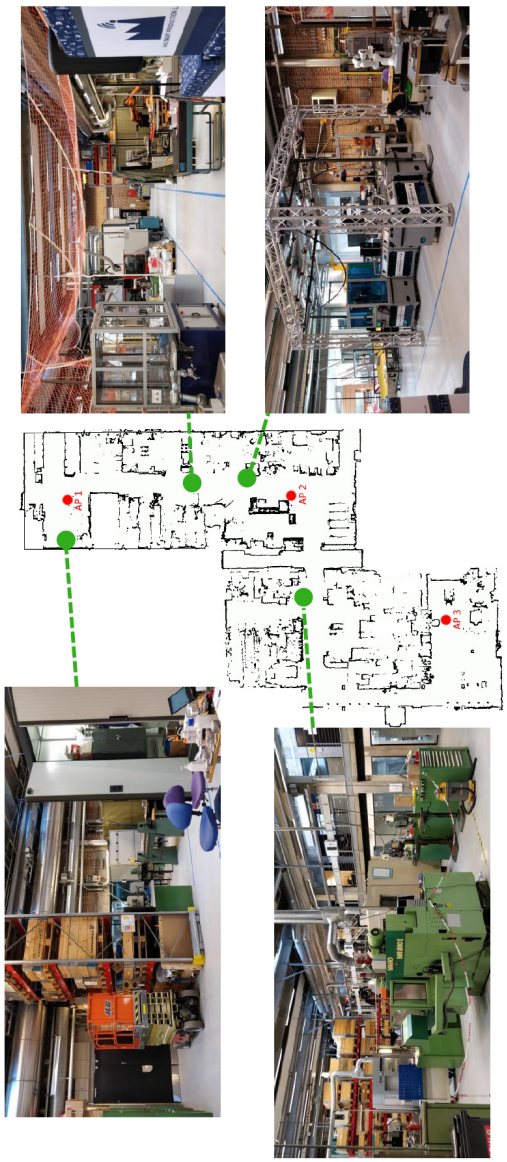

Figure A.1: Testing environment used throughout the project. The Cisco MR36 Access Points (AP) are highlighted in red.

Appendix B

Feedback during the Thesis

This appendix contains the feedback and/or questions received as part of presentations and paper reviews.

B.1 Online Presentation for Nokia Bell Labs

On the 12th of May 2021, an online presentation was held for Nokia Bell Labs where ∼36 people attended. Selected questions raised in the session is shown below. Paraphrased responses to the questions is likewise shown.

Q1: *Is the effect of scanning included in the CCDF results?*
Yes. For the cases where the STA would scan while connected, some data transmission still occurs, albeit with significantly higher latency which is also reflected in the CCDF. If the STA disconnects, some of the effect can be seen in the CCDF but that is not the case for the dropped packets.

Q2: *Were there ping-pong effects between the 2 APs (resulting in increased delay)? Are they captured in the CCDF?*
Because of the setup of the Access Points in the testing environment, this was not the case, as the handover would normally be initiated once the RSSI had degraded to the -85 dBm threshold, at which the STA would be fairly close to the next AP. It therefore didn't for example stay in a region between two APs coverage areas. However, if this was the case, it would result in more handovers occurring, which would have a negative impact on the latency performance.

Q3: *The results show RTT increases during handover, but what is the effect to packet loss?*
The packet loss was observed to be increased, with 0.043% observed. However, there isn't any notable difference between the tested configurations, which could mean that the optimizations don't sufficiently affect the part causing the packet loss. We did however notice an increase in PER as background traffic was introduced.

Q4: *It says that there is a full disruption of traffic for ∼55 ms. What happens to packets generated during that window? Are UL packets buffered in STA until link*

is available again, or dropped? DL packets are all simply lost because they arrive at the source AP and there is no forwarding to target STA?
While this was not investigated as part of the study, some guesses can be made. For uplink, the STA will queue the packets until connection is reestablished, up until some threshold. This is dependent on the implementation of the STA. For downlink, it could be that the APs can hold back the traffic if it has noticed that the STA has disconnected and is initiating a new connection to another AP on the network. While this requires some coordination between the APs, it is not unreasonable with the Cisco Meraki setup. It is however just a guess.

Q5: *How were static measurements performed? Standing at the same point?*
Measured from 4 different points at one of the APs to have some variation in terms of RSSI.

Q6: *Sudden drop at 50-60 ms in CCDF in the static case?*
So this has been investigated a bit and was found to be related solely to downlink. The ping responses are received in the same message from the AP, so it might be due to some MAC aggregation mechanism. Why it doesn't reply immediately is however a good question which I don't have an answer for.

Q7: *At which layer is packet loss considered? Only transport? Retransmission mechanisms at higher layers or only low layers of Wi-Fi*
For these measurements, we operate at the transport layer. So the packet losses here are those which the application did not receive any reply from, and because the Linux ping functionality is used, there are no higher-layer retransmission mechanisms.

Q8: *Impact of speed of the robot/UE?*
The robot moved with a maximum speed of 1.5 m/s, but the general impact of the speed of the STA has not been investigated.

Q9: *How to represent latency for lost packets? Terminator to the curves to indicate the packet error rate in probabilities?*
It is true that the packet losses are not reflected in the CCDFs. It would make sense to maybe cut off the tails if we don't have enough confidence.

Q10: *Comparison to 3GPP technologies (4G/5G, MulteFire, NR-U)?*
Wi-Fi has great performance for mean or lower percentiles, but is easily affected by interference, especially compared e.g. 4G and 5G. Because of the STA doing the roaming by itself, it has to discover nearby nodes and won't know e.g. current usage at each, so there are some downsides. But as we see from the data, it can be improved to perform much better when considering the reliable latencies. As for MulteFire and NR-U I can't say since i don't know the specifics of those technologies.

B.2 Reviews for Paper 1 Submission

Paper 1 was submitted to and accepted for presentation at the European Wireless 2021 conference. The reviews received with the notification of acceptance are shown below, along with initial comments.

Review 1

Relevance and timeliness: Excellent (5)
Technical content and scientific rigour: Solid work of notable importance (4)
Novelty and originality: Some interesting ideas and results on a subject well investigated (3)
Quality of presentation: Excellent (5)

Strong aspects
The paper discusses a well-designed and convincing study of factors contributing to the performance of WiFi in environments characterized by AP handovers. The variations considered are relevant, and the paper is clearly written with excellent graphics.

Weak aspects
As well-written as the paper is, its findings do not seem especially novel or important. The conclusion that mobility performance is improved by the use of dedicated clean channels seems intuitively obvious. The paper serves a valuable purpose by confirming this intuition and quantifying the effect, and it should be accepted and published as such. But I can't rank it highly for innovation.

Recommended changes
No suggestions. The paper is fine as written; a more groundbreaking topic would result in a different paper.

While the findings may not be groundbreaking, as current deployments are measured to obtain some form of reference, it is a fair statement. Regarding using dedicated channels being obvious, I agree but as the reviewer also claims it is nice that it has been confirmed experimentally.

Review 2

Relevance and timeliness: Good (4)
Technical content and scientific rigour: Valid work but limited contribution (3)
Novelty and originality: Some interesting ideas and results on a subject well investigated (3)
Quality of presentation: Well written (4)

Strong aspects
It is a very clearly written and easy to follow paper. The experimental procedure carried out is sufficiently rigorous and the indications obtained from the experiments are interesting.

Weak aspects
It is a paper that shows an empirical experiment and compares the performance of already known algorithms. Therefore, there is not much innovation in the contents for this reason. The results found refer to a specific scenario and therefore it is not possible to judge their generalizability.

Recommended changes
The paper under review basically consists of an experiment that serves to compare, from the point of view of latency measurement, different mobility management approaches in environments with WiFi coverage and applications with stringent requirements in terms of delay. The takeaway message in output to the experimental campaign illustrated is that much remains to be done to ensure Wi-Fi solutions optimized for latency to meet the requirements of IIoT applications on the move.
The lack of variability of the environments in which the measurements were carried out and of the provision of averaged values on a number of scenarios different from each other (with the same load conditions and approach to channel management, of course) slightly reduces the value of the experimentation. The authors should at least provide some more words of comment on the generalizability of the results in different IIoT scenarios.

Again, as the paper targets the performance of readily-available hardware and software, the argument for lack of "new" concepts is fair. Regarding the lack of variability in the environments, it is true that being able to e.g. measure in multiple setups could improve the quality of the work and help relate the data to more settings, however with the environment closely resembling a real factory, the propagation conditions and finally the mobility conditions (through the use of an actual AMR) are realistic, and should thus be relatable to other setups. The background traffic was however indeed artificial, but with the actual load at a factory being highly dependent on current equipment, this constant traffic of 10 Mbit/s should strike a middle ground. As for clarifying about the generalizability of the results in different IIoT scenarios, this is reasonable as it can help the reader relate to the findings in the paper and compare it to other settings.

Review 3

Relevance and timeliness: Good (4)
Technical content and scientific rigour: Valid work but limited contribution (3)
Novelty and originality: Significant original work and novel results (4)
Quality of presentation: Excellent (5)

Strong aspects
This technical paper focuses on measurement of latency in the context of WLAN for industrial scenarios. Specifically, it focuses on handover aspects.

Weak aspects
The paper should have provided more general analytical considerations to increase the value and impact of its results.

Recommended changes
The article is well-written and its organisation is clear. No significant changes are needed.

This is similar to the feedback from reviewer #2 for the generalizability of the results.