

Experimental Evaluation and Enhancements of Wireless Communication Technologies for the Factories of the Future

Mogensen, Rasmus Suhr

DOI (link to publication from Publisher):
[10.54337/aau535063853](https://doi.org/10.54337/aau535063853)

Publication date:
2023

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Mogensen, R. S. (2023). *Experimental Evaluation and Enhancements of Wireless Communication Technologies for the Factories of the Future*. Aalborg Universitetsforlag. <https://doi.org/10.54337/aau535063853>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

**EXPERIMENTAL EVALUATION AND
ENHANCEMENTS OF WIRELESS
COMMUNICATION TECHNOLOGIES
FOR THE FACTORIES OF THE FUTURE**

**BY
RASMUS SUHR MOGENSEN**

DISSERTATION SUBMITTED 2023



AALBORG UNIVERSITY
DENMARK

Experimental Evaluation and Enhancements of Wireless Communication Technologies for the Factories of the Future

Ph.D. Dissertation
Rasmus Suhr Mogensen

Dissertation submitted January, 2023

Dissertation submitted: January, 2023

PhD supervisor: Assoc. Prof. Gilberto Berardinelli
Aalborg University

Assistant PhD supervisors: Asst. Prof. Ignacio Rodriguez Larrad
Aalborg University
Senior Research Specialist Troels Kolding
Nokia Bell Labs

PhD committee: Professor Cedomir Stefanovic (chairman)
Aalborg University, Denmark
Professor Martin Wollschlaeger
Technische Universität Dresden, Germany
Principal Scientist Ognjen Dobrijevic
ABB Sweden, Sweden

PhD Series: Technical Faculty of IT and Design, Aalborg University

Department: Department of Electronic Systems

ISSN (online): 2446-1628
ISBN (online): 978-87-7573-767-3

Published by:
Aalborg University Press
Kroghstræde 3
DK – 9220 Aalborg Ø
Phone: +45 99407140
aauf@forlag.aau.dk
forlag.aau.dk

© Copyright: Rasmus Suhr Mogensen

Printed in Denmark by Stibo Complete, 2023

Curriculum Vitae

Rasmus Suhr Mogensen



Rasmus Suhr Mogensen received his B.Sc. degree in Electronics and IT and M.Sc. in Networks and Distributed Systems from Aalborg University in 2013 and 2018 respectively. He is currently pursuing a Ph.D. degree at Aalborg University in collaboration with Nokia Bell Labs. His research interests include Industry 4.0, industrial wireless technologies, transport layer protocols and autonomous mobile robots.

Abstract

The 4th Industrial revolution, or Industry 4.0 (I4.0), envisions a transition from linear production to highly flexible equipment that offers production customizability with high reliability. The increased flexibility and agile nature of the production line is expected to increase not only production output but also enable adaptability to easier support changing market needs and product lines.

One of the key technological enablers for this vision is reliable, low latency wireless communication, which presents a disruptive evolution in factory network. It will enable flexible communication paths between all entities of the factory allowing for not only quick reconfiguration of equipment but also support of entirely new use cases., e.g. autonomous mobile robots. Wireless communication already sees adoption in factories today in the form of ruggedized Wi-Fi equipment for use-cases that are tolerant to latency and packet losses. For more advanced use-cases there exists proprietary extensions which further improves performance, at the cost of flexibility and more complex co-existence. Therefore, more open technologies such as Wi-Fi 6 and LTE/5G NR are starting to be used in such contexts as they advertise features which promise better support of the diverse sets of communication requirements and future use-cases.

This thesis aims at experimentally investigating the potential of modern wireless technologies in an I4.0 context as well as proposing, exploring and evaluating enhancements that further strengthen their applicability for challenging use-cases. The main research questions we aim at answering are the following: What are the characteristics of communication networks and data traffic in real world factories? To what extent can existing off-the-shelf technologies be used for static and mobile elements on the factory floor? How can enhancements via open building blocks such as transport layer multi-connectivity be used to improve the performance? As the manufacturing industry is a new “vertical” in the wireless space, we first propose an experimental framework for integration of modern wireless solutions in factories. We then show how the information gathered in real factories can be used to design an appropriate wireless solution for a wireless Manufacturing

Execution System.

Secondly, we investigate how off-the-shelf technologies such as Wi-Fi, LTE and 5G NR Rel. 15 perform not only in controlled test environments but also in real factories. This includes both stationary equipment and autonomous mobile robots. Thirdly, we show how transport layer multi-connectivity can be used to greatly improve network availability and reliability by combining potentially disjoint networks and technologies. Our research shows that, for static use-cases with low load Wi-Fi is capable of delivering latencies in the order of 9ms at 99.99 %-tile but degrades to >1s if the load is increased or in the presence of interference, while a private LTE or 5G NR Rel. 15 is delivering latencies in the order of 30/15 ms respectively. Under mobility conditions, a significantly larger degradation is observed in Wi-Fi if not optimized for the specific deployment. We also show that deploying radio aware transport layer multi-connectivity schemes can completely mitigate the effect of the Wi-Fi handover for mobile use cases, bringing the latency down to around 120 ms. This can further be improved by exploiting selective packet duplications and per traffic flow QoS in more complex scenarios.

Resumé

Den 4. industrielle revolution, eller Industri 4.0 (I4.0), omhandler visionen om at skifte fra lineære produktionslinjer til fleksibel produktion der tilbyder større tilpasning med pålidelighed end tidligere. Den øgede fleksibilitet og agile karakter af produktionen forventes ikke kun at øge produktionsoutput, men også muliggøre lettere tilpasning af skiftende markedsbehov og produkter.

En af de vigtigste teknologier for realiseringen af denne vision er pålidelig, trådløs kommunikation med lav latenstid da disse giver enestående muligheder for at revolutionere fabriks netværk. Introduktionen af trådløs kommunikation vil muliggøre fleksible kommunikationsveje mellem alle enheder på fabrikken, hvilket giver mulighed for både hurtig omkonfiguration af udstyr samt understøttelse af helt nye brugsscenarier som f.eks. autonome mobile robotter. Trådløs kommunikation ser allerede taget i brug på fabrikker i dag i form af robust Wi-Fi udstyr til brug i situationer, der er tolerante over for højere latenstid og pakke tab. For mere krævende situationer findes der proprietære løsninger, som forbedrer ydeevnen yderligere på bekostning af fleksibilitet og mere kompleks sameksistens med andre teknologier. Derfor begynder mere åbne teknologier såsom Wi-Fi 6 LTE og 5G NR. at blive brugt i sådanne sammenhænge, da de reklamerer for funktioner, som lover bedre support til de forskellige sæt af kommunikationskrav og fremtidige brugs scenarier.

Denne afhandling har til formål at eksperimentelt undersøge potentialet i moderne trådløse teknologier i en I4.0 sammenhæng. Dette inkluderer at foreslå, udforske og evaluere forbedringer, der yderligere styrker deres anvendelighed til udfordrende brugsscenarier. De vigtigste forskningsspørgsmål, vi ønsker at besvare, er følgende: Hvad er kendetegnene for kommunikationsnetværk og data trafik på fabrikker i den virkelige verden? I hvilket omfang kan eksisterende teknologier bruges til statiske og mobile elementer på i produktionen? Hvordan kan forbedringer via åbne byggeklodser såsom transportlags multi-connectivity bruges til at forbedre ydeevnen? Da produktionsindustrien er en ny "vertikal" sektor inden for det trådløse marked, foreslår vi først en eksperimentel ramme for integration af moderne trådløse

løsninger på fabrikker. Vi viser derefter hvordan en sådan ramme bruge til at indsamle formation i virkelighedens fabrikker samt hvordan dette kan bruges til at designe en passende trådløs løsning til et Manufacturing Execution System.

For det andet undersøger vi, hvordan allerede tilgængelige teknologier såsom Wi-Fi, LTE og 5G NR. Rel. 15 performer i kontrollerede testmiljøer samt rigtige fabrikker. Dette omfatter både stationært udstyr og autonome mobile robotter. For det tredje viser vi, hvordan transportlags multi-connectivity kan bruges til at forbedre netværkets tilgængelighed og pålidelighed ved at kombinere potentielt uafhængige netværk og teknologier. Vores forskning viser, at for statiske anvendelsestilfælde med lav belastning er Wi-Fi i stand til at levere en latenstid i størrelsesordenen 9 ms ved 99,99 %-tile, men forringes til >1 s, hvis belastningen øges eller ved tilstedeværelse af interferens, mens et privat LTE eller 5G NR Rel. 15 netværk leverer latenstider i størrelsesordenen 30/15 ms. Under mobilitetsforhold observeres en væsentlig større forringelse i Wi-Fi, hvis den ikke er optimeret til den specifikke implementering. Vi viser også, at implementering af radiobevidste transportlags multi-connectivity fuldstændigt kan fjerne effekten af Wi-Fi handover i brugsscenarier med mobilitet, hvilket bringer latenstiden ned til omkring 120 ms. Dette kan forbedres yderligere ved at udnytte selektive pakkeduplikationer og QoS pr. trafikflow i mere komplekse scenarier.

Glossary

3GPP 3rd Generation Partnership Project

5G NR 5G New Radio

AAU Aalborg University

ACK Acknowledgement

AGV Automated Guided Vehicle

AMR Autonomous Mobile Robot

AP Access Point

BPS Best Path Scheduling

CSMA/CA Carrier Sense Multiple Access with Collision Avoidance

DetNet Deterministic Networking

DL Downlink

DT Delay Tolerant

HTTP Hyper Text Transport Protocol

I4.0 Industry 4.0

IEEE Institute of Electrical and Electronics Engineers

IETF Internet Engineering Task Force

IIoT Industrial Internet of Things

IoT Internet of Things

IP Internet Protocol

IRT Isochronous Real Time

Glossary

IT Internet Technology

JSON JavaScript Object Notation

LBT Listen-Before-Talk

LTE Long Term Evolution

MAC Medium Access Control

MADE Manufacturing Academy of Denmark

MES Manufacturing Execution System

mmWave Millimeter Wave

MP-QUIC Multipath Quick User Datagram Internet Connection

MP-TCP Multipath Transport Control Protocol

NR-U New Radio-Unlicensed

NRT Non-Real Time

OFDMA Orthogonal Frequency Domain Multiple Access

OT Operational Technology

PDCP Packet Data Convergence Protocol

QoS Quality of Service

QUIC Quick User Datagram Internet Connection

RAN Radio Access Network

Rel. Release

RSSI Received Signal Strength Indicator

RT Real Time

RTT Round Trip Time

SACA Stream Aware Congestion Algorithm

SD Selective Duplication

SDQoS Selective Duplication with Quality of Service

SR-MPQUIC Selective Redundant Multipath Quick User Datagram Internet Connection

Glossary

STA Station

SW Software

TCP Transport Control Protocol

TSN Time Sensitive Networking

UE User Equipment

UL Uplink

URLLC Ultra-Reliable Low Latency Communication

UWB Ultra Wide Band

Glossary

Thesis Details

Thesis Title: Experimental Evaluation and Enhancements of Wireless Communication Technologies for Factories of the Future
PhD Student: Rasmus Suhr Mogensen
Supervisors: Assoc. Prof. Gilberto Berardinelli, Aalborg University
Asst. Prof. Ignacio Rodriguez Larrad, Aalborg University
Senior Research Specialist Troels Kolding, Nokia Bell Labs

This PhD thesis is the result of three years of research at the Wireless Communication Networks (WCN) section (Department of Electronic Systems, Aalborg University, Denmark) in collaboration with Nokia Bell Labs (Aalborg, Denmark). The work was carried out in parallel with mandatory courses required to obtain the PhD degree.

The main body of the thesis consists of the following articles:

- Paper A:** I. Rodriguez, **R. S. Mogensen**, A. Fink, T. Raunholt, S. Markussen, P. H. Christensen, G. Berardinelli, P. Mogensen, C. Schou, and O. Madsen, "An Experimental Framework for 5G Wireless System Integration into Industry 4.0 Applications", *Energies*, vol. 14, no. 15, 4444, Jul. 2021.
- Paper B:** **R. S. Mogensen**, I. Rodriguez, C. Schou, S. Mortensen, and M. Sørensen, "Evaluation of the Impact of Wireless Communication in Production via Factory Digital Twins", *Manufacturing Letters*, vol. 28, pp. 1-5, Apr. 2021.
- Paper C:** **R. S. Mogensen**, I. Rodriguez, G. Berardinelli, G. Pocovi and T. Kolding, "Empirical IIoT Data Traffic Analysis and Comparison to 3GPP 5G Models", 2021 IEEE 94th Vehicular Technology Conference (VTC2021-Fall), pp. 1-7, 2021,
- Paper D:** **R. S. Mogensen**, I. Rodriguez, G. Berardinelli, A. Fink, R. Marcker, S. Markussen, T. Raunholt, T. Kolding, G. Pocovi, and S. Barbera,

“Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0”, 2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall), pp. 1-7, 2019.

Paper E: A. Fink, **R. S. Mogensen**, I. Rodriguez, T. Kolding, A. Karstensen, and G. Pocovi, “Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility”, 26th European Wireless (EW 2021), pp. 125-132, 2021.

Paper F: I. Rodriguez, **R. S. Mogensen**, A. Schjørring, M. Razzaghpour, R. Maldonado, G. Berardinelli, R. Adeogun, P. H. Christensen, P. Mogensen, O. Madsen, C. Møller, G. Pocovi, T. Kolding, C. Rosa, B. Jørgensen, and S. Barbera, “5G Swarm Production: Advanced Industrial Manufacturing Concepts Enabled by Wireless Automation”, IEEE Communications Magazine, vol. 59, no. 1, pp. 48-54, Jan. 2021.

Paper G: **R. S. Mogensen**, C. Markmøller, T. K. Madsen, T. Kolding, G. Pocovi, and M. Lauridsen, “Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications”, 2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring), pp. 1-5, 2019.

Paper H: A. Fink, **R. S. Mogensen**, I. Rodriguez, T. Kolding, A. Karstensen, and G. Pocovi, “Radio-Aware Multi-Connectivity Solutions based on Layer-4 Scheduling for Wi-Fi in IIoT Scenarios”, IEEE Wireless Communications and Networking Conference (WCNC 2022), pp. 1821-1826, 2022.

Paper I: **R. S. Mogensen**, S. B. Damsgaard, I. Rodriguez, G. Berardinelli, A. Fink, T. Kolding, and G. Pocovi, “A Novel QoS-Aware Multi-Connectivity Scheme for Wireless IIoT”, IEEE Access, vol. 10, pp. 104123-104134, Oct. 2022.

This thesis has been submitted for assessment in partial fulfilment of the PhD Degree. The thesis is based on the submitted or published papers that are listed above. Parts of the papers are used directly or indirectly in the extended summary of the thesis. As part of the assessment, co-author statements have been made available to the assessment committee and also available at the Faculty.

Acknowledgements

I would like to express my sincere thanks to my supervisors Gilberto Berardinelli, Ignacio Larrad, Troels Kolding, in addition to Guillermo Pocovi (who although not in the role of an official supervisor) has provided support, encouragement, and guidance throughout this Ph.D. project even in difficult times. Furthermore, I would like to thank my fellow Ph.D. colleagues, especially, Sebastian Damsgaard who has provided assistance in debugging the numerous hardware and software challenges. Additionally, I would like to thank the master students I have had the pleasure of supervising specifically Andreas Fiink from whom I have gotten invaluable assistance with the plethora of measurements conducted for the work presented in this thesis as well as being output for my rantings.

Lastly, I would like to thank my family who kept asking me "Are you done yet?". To which I can finally "YES!, now stop asking me".

Rasmus Suhr Mogensen
Aalborg University, January 12, 2023

Acknowledgements

Contents

Curriculum Vitae	iii
Abstract	v
Resumé	vii
Glossary	ix
Thesis Details	xiii
Acknowledgements	xv
I Thesis Summary	1
1 Introduction	3
1.1 Current state of industrial Ethernet solutions	4
1.2 Wireless communication technologies in Industry 4.0	6
1.2.1 Solutions for unlicensed spectrum bands	6
1.2.2 Solutions for licensed spectrum bands	7
1.2.3 Recent research on wireless technologies for industrial applications	8
1.3 Thesis objectives and scope	9
1.4 Methodology	11
1.5 Research outcomes	13
1.5.1 List of scientific papers	13
1.5.2 Other research contributions	15
1.6 Thesis outline	19
References	21
2 Approaching industrial end users	27
2.1 Motivation	27
2.2 Objectives	27

Contents

2.3	Summary of contributions	28
2.4	Main findings	29
2.5	Research impact and follow-up studies	30
	References	31
3	Performance of commercially available wireless technologies	33
3.1	Motivation	33
3.2	Objectives	33
3.3	Summary of contributions	34
3.4	Main findings	35
3.5	Research impact and follow-up studies	36
	References	37
4	Transport-layer multi-connectivity for existing wireless networks	39
4.1	Motivation	39
4.2	Objectives	40
4.3	Summary of contributions	40
4.4	Main findings	41
4.5	Research impact and follow-up studies	43
	References	43
5	Conclusion & Future work	45
5.1	Recommendations	48
5.2	Future work	49
	References	50
II	Annex	51
A	MiR fleet operations: scalability analysis	53
A.1	Short conclusion	59
	References	59
B	Overview of the out-of-the-box performance of wireless technologies in operational conditions at the AAU 5G Smart Production Lab	61
B.1	Short conclusion	63
	References	65
C	SR-MPQUIC: real-world tests and analysis	67
C.1	Short conclusion	68
	References	69

III	Papers	71
A	An Experimental Framework for 5G Wireless System Integration into Industry 4.0 Applications	73
A.1	Introduction	75
A.2	Learnings from Industry and Related Framework Design Considerations	77
A.3	Experimental Framework for 5G Integration	79
A.3.1	Industrial Automation Operational Flow	79
A.3.2	Industrial Research Lab Facilities	83
A.3.3	5G Hardware and Software Prototyping Tools	85
A.4	Framework Applicability Example: 5G Autonomous Mobile Robots	97
A.4.1	5G-Connected Autonomous Mobile Robots	97
A.4.2	5G Mobile Edge-Cloud Planner for Autonomous Mobile Robots	102
A.5	Considerations for Future Framework Versions	105
A.6	Conclusions	106
	References	108
B	Evaluation of the Impact of Wireless Communication in Production via Factory Digital Twins	113
B.1	Introduction	115
B.2	Integrating Wireless in Factory Digital Twins	116
B.3	Digital Twin Simulation Settings	119
B.4	Production performance results & Discussion	120
B.5	Conclusions	123
	References	123
C	Empirical IIoT Data Traffic Analysis and Comparison to 3GPP 5G Models	125
C.1	Introduction	127
C.2	Modeling of IIoT traffic in 5G	129
C.2.1	IIoT traffic models	129
C.2.2	5G QoS model	129
C.3	Industrial Use Cases	131
C.3.1	Unit Test Cell (UTC)	131
C.3.2	Visual Inspection Cell (VIC)	132
C.3.3	Autonomous Mobile Robot (AMR)	132
C.4	Results & Discussion	135
C.4.1	Individual Use Case Results	135
C.4.2	Summary of traffic statistics	140
C.4.3	Comparison with current 3GPP 5G IIoT models	140

C.5 Conclusion	142
References	143
D Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0	145
D.1 Introduction	147
D.2 Integration of Wireless Industrial Ethernet	149
D.3 Multi-access Wireless Gateway	151
D.4 Test setup	154
D.5 Performance Results	156
D.6 Conclusions	158
References	159
E Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility	161
E.1 Introduction	163
E.2 Handover Improvements in Wi-Fi	164
E.3 Test Setup	167
E.4 Test Results	170
E.5 Conclusion	178
References	179
F 5G Swarm Production: Advanced Industrial Manufacturing Concepts enabled by Wireless Automation	181
F.1 Introduction	183
F.2 Manufacturing Industry Goals and Wireless Automation Evolution	185
F.3 The AAU 5G Smart Production Lab	190
F.4 Performance of Industrial Wireless Use Cases	192
F.4.1 Wireless Control of Industrial Production	192
F.4.2 Control of Industrial Mobile Robots	194
F.5 Conclusion	196
References	196
G Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications	199
G.1 Introduction	201
G.2 Selective Redundant MP-QUIC	203
G.2.1 Stream aware congestion algorithm (SACA)	204
G.3 Methodology and Test Setup	205
G.4 Results	207
G.4.1 Static Delay, Variable Packet Loss	208
G.4.2 Wired Domestic IEEE Internet model	210

G.4.3	Wireless model - LTE highway drive test	210
G.5	Conclusions	212
	References	213
H	Radio-Aware Multi-Connectivity Solutions based on Layer-4 Scheduling for Wi-Fi in IIoT Scenarios	215
H.1	Introduction	217
H.2	Radio-aware Scheduling Schemes for Wi-Fi	219
H.2.1	Wi-Fi Packet Scheduler	219
H.2.2	Wi-Fi Mobility Coordinator	220
H.2.3	Wi-Fi Network Planning	222
H.3	Experimental Test Setup	222
H.4	Wi-Fi Performance Results	225
H.5	Discussion	229
H.6	Conclusion	230
	References	231
I	A Novel QoS-Aware Multi-Connectivity Scheme for Wireless IIoT	233
I.1	Introduction	235
I.2	QoS-Aware Link Selection Scheme	238
I.2.1	QoS-Aware Multi-Connectivity Scheme Implementation	241
I.3	Experimental Validation	245
I.3.1	Test description	245
I.3.2	Setup and Measurement Procedures	247
I.3.3	Parameter optimization	250
I.4	Results and Discussion	252
I.5	Conclusions and future work	256
	References	257

Contents

Part I

Thesis Summary

Chapter 1

Introduction

With the advent of Industry 4.0 (I4.0), a transitioning phase in manufacturing is currently taking place. The vision is to move from specialized production lines and traditional conveyor belts towards a flexible and agile configuration, enabling product customization, and maximization of production output. Several technological enablers and concepts are being discussed and formulated to realize this vision. These enablers range from big data, and artificial intelligence to autonomous mobile robots (AMR), matrix production, and digital twins [1–4]. Furthermore, one could argue that wireless communication is a prerequisite for increased flexibility and enablers such as AMRs. An industrial production line consists of several entities with different communication requirements, e.g., some require low latency and high reliability, while others do not. Depending on functionality and type, these entities can be segmented into different layers in the automation pyramid, shown in Figure 1.1 along with an indication of the tolerated latencies.

Typically, communication at the pyramid’s base has strict real-time (RT) network requirements. At the same time, upper layers are less affected by more significant latency variations (delay tolerant (DT)). Examples of such communication requirements can be found in [5–7]. To facilitate this communication, the recent trend has been to move from proprietary Fieldbus technologies towards IEEE 802.3 (Ethernet) based solutions as they offer a good mixture of network flexibility and deterministic performance [8]. Besides the wired elements, there is an emergence of mobile elements in factories in the form of autonomous guided vehicles (AGV) and AMR [9]. The communication requirements can vary greatly depending on the control granularity needed [10]. For example, complex movement instructions/decisions based on RT sensor inputs will impose stricter requirements as compared to simply providing the robot’s destination in the form of a set of coordinates. However, even for delay-tolerant control through a central fleet manager, the reliabil-

ity and availability of the wireless communication infrastructure may affect performance.




Fig. 1.1: Automation pyramid in factory automation and associated reference latencies [11].

Therefore, several challenges concerning the transition to a wireless infrastructure need to be accommodated before the envisioned flexibility and agile nature of I4.0 can be realized. This form of network access is very different from the current Ethernet-based network in which users only share networking resources in trunk connections. Given its nature of a shared medium, wireless brings extra challenges with respect to wired connections. First, mutual interference between users in the same network exists, severely impacting performance unless appropriately managed. Also, interference can come from external sources operating over the same frequency band, which in some cases can be challenging to avoid. Furthermore, the wireless channel is subject to path loss, shadowing, and temporary signal fades, leading to a stochastic behavior. This can be even more challenging if the environment is highly dynamic or the device itself is mobile. On top of these challenges, one must consider whether a wireless system is deployed in entirely new networks where applications and protocols can be tuned for wireless access (greenfield) or as a part of an already existing network infrastructure with interfaces to conventional Ethernet-based equipment (brownfield). The latter potentially imposes even stricter requirements as feature and functional parity needs to be kept with both existing networks and legacy equipment. On the other hand, it lessens the burden on the verticals and, hopefully, helps increase the short-term adoption rate.

1.1 Current state of industrial Ethernet solutions

Industrial Ethernet solutions such as PROFINET, Ethernet/IP etc. [12, 13], segment the medium access in network-cycles decoupled from the applications running on top. Using these cycles as a resource block, it is then possible to assign access to the application traffic flows in such a manner that their respective QoS requirements are fulfilled. However, this approach requires

1.1. Current state of industrial Ethernet solutions

that flows between these devices are pre-planned throughout the network, and any change requires a reconfiguration and validation of flows. The approach of PROFINET is exemplified in Figure 1.2. PROFINET is a proprietary full-stack solution that has the ability of multiplexing traffic application streams based on the specific QoS parameters for a given flow; it supports both TCP/IP and non-TCP/IP type traffic.




Fig. 1.2: An example of how PROFINET maps different traffic types to the network, based on [14].

- The Isochronous Real-Time (IRT) channel can facilitate cyclic traffic that is strictly scheduled according to an integer number of network cycles, and it must adhere to strict deadlines.
- The Real-Time (RT) channel can facilitate real-time traffic that can be either asynchronous or periodic RT. It does not necessarily adhere to the network cycle nor a common schedule or application cycle. The timeliness of the traffic can be specified in terms of latency or deadline, depending on the application requirement.
- The Non-Real Time (NRT) channel can facilitate sporadic or periodic traffic, but the traffic of this class type is served according to a best-effort scheme.

However, having proprietary solutions such as PROFINET is not necessarily beneficial as some of them use custom layer 1/2, which can lead to incompatibility between industrial devices and regular Ethernet devices. To this extent, IEEE and IETF have introduced two working groups; IEEE 802.1 Time-Sensitive Networking workgroup (TSN) and Deterministic Networking (DetNet) [15] that aim at providing a standardized solution for achieving low latency deterministic layer 2/3 performance. Therefore, a candidate wireless technology should support these communication types natively, allowing it to provide a one-to-one replacement of an existing wired connection.

Nevertheless, designing a wireless communication system that meets traditional use cases' strict performance and functional requirements (drilling, printing, etc.), including new use cases (mobile robots), is a substantial task. While wired communication such as Ethernet does have a shared medium, the communication is typically happening through controlled ingress and

egress point per device pair. This behavior is much harder to achieve in a wireless setting, increasing the probability of interference and performance degradation. This, in turn, imposes a more significant probabilistic constraint on the reliability and availability that can be achieved, thereby affecting the level of QoS that can be provided. Additionally, mobile devices must maintain reliability during mobility, including handover, and can act as potential mobile interferers.

1.2 Wireless communication technologies in Industry 4.0

While multiple wireless communication technologies have been considered for use in industrial networks [16, 17], this thesis focuses on the ones that have gained the most traction in recent years. These are the already widely deployed Wi-Fi (IEEE 802.11) and its 3GPP equivalent MulteFire and 5G NR-unlicensed (NR-U), which operate in unlicensed spectra. Cellular technologies such as LTE and 5G NR, operating in the licensed spectrum, are also studied.

1.2.1 Solutions for unlicensed spectrum bands

Solutions that operate in unlicensed bands need to share the wireless spectrum with other systems. Therefore, they are all subject to some form of medium access regulation mechanism to mitigate resource contention. This imposes challenges in ensuring low latency. Common for all the technologies within this category is that Listen-Before-Talk (LBT) mechanisms are required by regulation [18].

IEEE 802.11 is a communication standard that is widely deployed as a part of Wi-Fi, where only IEEE 802.11n [19] and onwards is considered in this study. For medium access, Wi-Fi relies on carrier sense multiple access with collision avoidance (CSMA/CA) as the primary interference mitigation mechanism. Therefore, by default, no coordination between different stations (STAs) connected to the same AP exists. Moreover, the default handover mechanism between APs is an uncoordinated break-before-make sequence dictated solely by a STA. Wi-Fi has been examined extensively for use in industrial applications; however, until IEEE 802.11ax [20], there has not been any support for RT applications as indicated by [21]. Therefore, multiple solutions have been proposed by various authors that aim to improve usability in RT industrial use-cases. Siemens, for example, has a proprietary solution, based on IEEE 802.11n, Industrial WLAN [22] that uses a centrally coordinated scheme to ensure that devices do not contend for resources. It uses multiple links to provide a redundant connection for guided vehicle applications. This solution does, however, not scale well due to its use of centralized TDMA

with only single user access [23]; however a new IEEE 802.11ax version is in the roadmap [24], which should perform better. Other solutions such as RT-WIFI [25] propose a new time-synchronized medium access control layer for IEEE 802.11 to replace the normal CSMA/CA mechanism. This solution significantly decreases latency compared to regular Wi-Fi in different test scenarios but lacks reliability when interference is introduced. IEEE 802.11ax aims to provide better support for heterogeneous device requirements and determinism by introducing a network-controlled OFDMA framework, which allows for allocation of time-frequency resource blocks to devices. Using this framework, it is also possible to pre-assign resources in time and frequency or even implement LTE like time-based scheduling. However, the use of this framework does require a pure IEEE 802.11ax network [26].

MulteFire [27] and NR-U [28] are new wireless technologies that aim at providing equivalent functionalities of LTE and 5G NR in the unlicensed spectrum, by introducing the notion of cellular networking and centralized control. Whereas the current focus of MulteFire has been on mobile broadband and VoIP and not time-critical industrial applications, NR-U accounts for such applications [29]. While they should offer better performance for use-cases that are static and involve mobility as compared to Wi-Fi, they are still subject to the challenges in shared spectrum environments and as such they require deployment in clean channels for optimal performance [30].

In summary, all the above technologies are required to have an LBT mechanism. To minimize this regulatory requirement's impact, one needs to have enforced spectrum control and a well-planned network. Such tight spectrum control would be feasible to be enforced in indoor factory environments, but not in public environments such as harbors.

1.2.2 Solutions for licensed spectrum bands

LTE [31] and 5G NR [32] operate in licensed bands where an operator or company can purchase exclusive utilization rights for given geographical areas, thereby ensuring that only devices associated with the licensee of the band can legally use it.

LTE is a wireless technology designed for wide-area networking and used primarily for mobile communication as well as mobile broadband. It leverages the principle of cellular networks to meet coverage and bandwidth demands. It can operate in two different transmission modes: Frequency Domain Duplex (FDD) and Time Domain Duplex (TDD), each with different capabilities [33]. As a device moves between cells, it will experience a handover to stay connected. The handover procedure in LTE is network-coordinated following the Break-Before-Make paradigm, meaning that a target cell is provided before the handover is initiated. Even though handover is coordinated, unlike that of Wi-Fi, it may still affect the reliability when this occurs [34]. Uplink

transmissions are performed based on grants to use a specific set of resources, such that devices do not have to contend for resources and there is no intra-cell interference.

5G NR is the latest generation of cellular communication. Unlike LTE, 5G has been designed to include support for IoT, mobile broadband communication, and low latency applications, including I4.0 use-cases. To this extent, several features observed in IEEE 802.1 are included in the Release 15 [35] specification to accommodate the real-time requirements of vertical use-cases. Furthermore, Release 16 adds features to support TSN applications within 5G [7]. This includes fine-grained resource reservation, as well as time synchronization capabilities. However, at the time of the writing of this thesis, such features have yet to become available in commercial modems and networks. Furthermore, reliability improvement concepts, such as multi-connectivity, are included in the standard. These improvements enable duplicated data transmissions on multiple links and are implemented at layer 2/1 of the protocol stack, where the main goal is to harden the channel and mitigate handover and interference impact. Different solutions exist depending on whether duplication is applied at Packet Data Convergence Protocol (PDCP) or Medium Access Control (MAC) layer [36].

1.2.3 Recent research on wireless technologies for industrial applications

The usage of Wi-Fi, LTE, and 5G NR for industrial applications has been considered prior to this study, but has become more prevalent in the last three years.

Wi-Fi 5 [37] has been studied to optimize its usage for industrial applications. By taking advantage of the IEEE 802.11e amendment, the work presented in [38, 39] illustrated MAC-layer tuning to accommodate soft-real time use cases. Furthermore, the authors in [25] have investigated potential improvements to the IEEE 802.11 MAC layer for applicability in real-time industrial use-cases. To this extent, one would expect that the improved flexibility of IEEE 802.11ax makes Wi-Fi 6 [20] suitable for industrial usage, as reported in [40, 41]. Mobility optimizations have been presented in [42], where telemetry information of the mobile elements (AGVs in this case) is used to lower handover outage times. Similarly, the work described in [43] proposed using multiple Wi-Fi modems in conjunction with the redundant MP-TCP [44] to have continuous connectivity in an area covered by two APs.

The idea of using LTE in industrial networks have previously been examined by [45], where the authors evaluate its applicability based on ping end-to-end latency estimations to different segments of a German public LTE network. They obtain a minimum roundtrip time (RTT) latency of 27 ms and an average of 35 ms, although they measure a significant variation in

latency depending on the time of day. Public LTE however, does not fully reflect the potential of LTE. In [46], the authors showcase how optimized LTE can provide sub 10 ms latencies, when optimizing the scheduling and the UE sleeping patterns. Such optimizations will be applicable in a private LTE network. To decrease delays related to the resource allocation procedure, concepts such as semi-persistent scheduling, where resource allocations are scheduled according to a pre-configured scheme or grant-free scheduling [47] can be used. To improve reliability at the cell edge, one can leverage multiple co-located base stations that provide diversity gain for the device (Coordinated Multi-Point (CoMP) [48]). Also, in [49] the authors showcase that LTE falls short of supporting a real-time control loop of a Cartesian robot. Here, 5G is presumed to be suitable.

5G has been investigated in [50], where the authors attempt to translate network latencies into the quality of experience from the operator's point of view. Furthermore, a comparison study of Wi-Fi 6 and 5G NR Release 15 for indoor usage has been presented in [51]. To this extent, the authors of [52] present an in-depth view of the various latency contributors in a 5G NR campus network. With respect to more futuristic features of 5G NR, the authors of [53] present a prototype and performance evaluation of 5G NR URLLC for transporting PROFINET real-time traffic.

To this extent, there is a large body of experimental research, including the one presented in this study, that addresses the reliability levels can be achieved by commercial solutions. However, we see a lack of experimental evaluation pertaining to real applications running in factories, including investigations of how one should successfully integrate existing technologies with e.g., legacy equipment. Additionally, investigations involving mobile robots are limited to simulation studies or simple experiments not including challenges related to handover.

1.3 Thesis objectives and scope

The vision of 5G and other wireless technologies targeting industrial applications is to provide the same performance and reliability as wired solutions while enabling the extra flexibility and agile nature of using a wireless system, such as support of device mobility. However, product optimization and feature implementation are often based on most use-cases in the consumer segment. Therefore, a gap between standardized features/performance and what can be achieved by available product implementations is expected. Organizations such as 5G-ACIA [54] have been created to act as an interface between industrial OTs and end-users and wireless technology providers to showcase the performance of 5G in real industrial applications. The motivation behind this approach is that simply referring to standardization or showing PowerPoint

slides with future performance goals is not sufficient to convince industrial OTs about the benefits of deploying industrial 5G. This Ph.D. project aims at investigating the potential of existing wireless technologies in addressing industrial communication needs, and at proposing novel fast-to-market enhancements.

In this respect, our studies are motivated by similar goals as the ones by 5G-ACIA, with the primary focus being on the performance of currently available wireless technologies such as LTE and 5G NR Release (Rel.) 15 in the licensed spectrum and Wi-Fi 5 and 6 in the unlicensed spectrum, for industrial communication. To this extent, we also see considerable potential in using hybrid access techniques to mitigate the impact of handovers and harden the communication for, e.g., AMR and AGV usage. As it is an enhancement on top of the radio access technology, it provides fast-time-to-market, especially in existing network deployments, with the only requirement of multi-modem support at the device side and software capable of handling the multiple flows. Furthermore, such solutions can easily be transferred to different radio technologies depending on the implementation.

Specifically, the study aims at contributing to the state-of-the-art via the following aspects:

1. Investigating whether industrial verticals possess the necessary know-how to understand the impact and implications of potential wireless solutions.
2. Designing real state-of-the-art hardware that facilitates the integration of wireless communications with legacy industrial equipment in a fully transparent manner.
3. Analyzing and characterizing in-depth the data traffic patterns of real-world industrial use cases.
4. Understanding the actual performance of current wireless communication technologies in real industrial setups, especially the inclusion of mobility scenarios.
5. Designing of novel methods for improving the performance of the existing network using hybrid access techniques.

Based on these aspects, multiple hypotheses (H) and associated research questions (Q) are formulated, which will serve as a basis for the study.

H1: A new methodological framework is required for a smooth integration of modern wireless technologies in existing factories.

Q1: What are the major barriers hindering the adoption of wireless technologies by verticals?

Q2: How to ensure a gradual and efficient adoption of wireless technologies in factories?

H2: Commercially available wireless technologies can be valid solutions for a large set of industrial applications.

Q3: What are the main characteristics of the data traffic for the relevant industrial applications?

Q4: To what extent can existing off-the-shelf wireless technologies address the industrial communication needs?

H3: Intelligent scheduling and path selection is required for transport layer multi-connectivity to support complex multi-flow QoS applications in already deployed networks, especially under mobility conditions.

Q5: What are the limitations of current transport-layer protocols in supporting challenging communication requirements?

Q6: To what extent can radio access technology metrics be used for enhancing multi-connectivity performance with respect to pure transport layer approaches?

1.4 Methodology

While the usage of extensive Monte Carlo simulations is the most common approach for the system-level evaluation of wireless technologies, in this thesis we resort to an experimental evaluation in real industrial laboratories or facilities. This choice is due to the factors described below.

1. **Higher ‘trust’ in results, but with reduced generalization.** In an experimental evaluation, results are obtained from real applications and using real hardware, and as such one cannot dispute the realism of the test. This comes at the expense of lower generalizability of the test results, which has to be properly addressed when disseminated.
2. **Easy conversion from experimental setup to prototype/services, subject to practical limitations.** Contrary to simulation studies, any experimental setup can easily be converted to a prototype implementation or service, subject to the practical limitations of the used hardware/software and the environment.
3. **Improved learnings due to realistic conditions, but limited complexity and scope of the tests.** Simulation models may disregard relevant real-world aspects and lead to inaccurate results. In such cases, experimental results can enrich the findings of simulations or even contribute to new

modeling approaches or technological enhancements. Nevertheless, performing extensive experiments with multiple entities and measurement configurations is very labor-intensive; additionally, the time acceleration factor of simulations is not present.

The first two aspects are expected to provide value to the vertical domain. If the experiments are based on an application in a factory, they can translate it directly to potential gains of introducing/leveraging a wireless infrastructure. Furthermore, such experiments would also provide valuable knowledge about deploying different wireless technologies in factories on performance and IT tasks. Also, experimental setups may be commercialized via prototypes or services by improving production setup or OT equipment. The last aspect is expected to enrich the research as the measurements can be used to verify simulation assumptions and provide input to wireless systems design.




Fig. 1.3: Methodology flow followed in the experimental research described in this study.

The methodology flow used in the experimental research is depicted in Figure 1.3. In brief, the different phases are described as follows:

1. **Problem analysis:** this phase may involve a dialogue with verticals for understanding their level of knowledge on wireless communication and the actual needs, as well as technical analysis. Depending on the specific scope of the research activity, the latter can include an analysis of the data traffic characteristic of the specific industrial application/service to be wirelessly supported, and – in case of a brownfield solution- an analysis of the actual characteristics and limitation of the network infrastructure where the new solution is to be integrated.
2. **System design:** this phase aims to use the knowledge from either the technical analysis to select an appropriate technology, as well as to investigate the proper configuration of networks and technology parameters. Additionally, it includes the design of an appropriate system architecture capable of supporting the required network functionalities if native network support is not present.
3. **Prototype implementation:** a prototype is developed based on the design of the previous phase. This phase also includes initial testing to ensure that intended functionalities are supported.

4. Experiment design: this phase includes defining key performance indicators and test-beds for the experiments. Furthermore, if possible, multiple test scenarios are defined to increase the generalization of the results.
5. Performance analysis: this phase includes an analysis in terms of the chosen key performance indicators, as well as a critical assessment for addressing shortcomings and further optimization options, given the initial goals and performance expectations defined in the system design phase.

As mentioned earlier, the generalization of the findings obtained from experimental studies might be challenging. This is because such findings might be strongly linked to the specific experimental environment and operational conditions. Generalization possibilities will be then addressed in each of the research contributions.

1.5 Research outcomes

This section summarizes the various research outcomes of this Ph.D. study. It includes the list of scientific papers that constitute the main body of the thesis, a description of relevant side contributions, such as the 5G Smart Production Lab and the research tools developed to carry out the experimental work, and also a summary of relevant external collaborations.

1.5.1 List of scientific papers

This Ph.D. thesis is organized as a collection of papers, which are listed below. Note that the publications are not listed chronologically but in a logical order following the statements from Section 1.3. This order will be kept for the technical summary of the work in the following chapters.

Paper A: I. Rodriguez, **R. S. Mogensen**, A. Fink, T. Raunholt, S. Markussen, P. H. Christensen, G. Berardinelli, P. Mogensen, C. Schou, and O. Madsen, “An Experimental Framework for 5G Wireless System Integration into Industry 4.0 Applications”, *Energies*, vol. 14, no. 15, 4444, Jul. 2021.

Paper B: **R. S. Mogensen**, I. Rodriguez, C. Schou, S. Mortensen, and M. Sørensen, “Evaluation of the Impact of Wireless Communication in Production via Factory Digital Twins”, *Manufacturing Letters*, vol. 28, pp. 1-5, Apr. 2021.

Paper C: **R. S. Mogensen**, I. Rodriguez, G. Berardinelli, G. Pocovi and T. Kolding, “Empirical IIoT Data Traffic Analysis and Comparison to 3GPP 5G Models”, 2021 IEEE 94th Vehicular Technology Conference (VTC2021-Fall), pp. 1-7, 2021,

- Paper D:** **R. S. Mogensen**, I. Rodriguez, G. Berardinelli, A. Fink, R. Marcker, S. Markussen, T. Raunholt, T. Kolding, G. Pocovi, and S. Barbera, "Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0", 2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall), pp. 1-7, 2019.
- Paper E:** A. Fink, **R. S. Mogensen**, I. Rodriguez, T. Kolding, A. Karstensen, and G. Pocovi, "Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility", 26th European Wireless (EW 2021), pp. 125-132, 2021.
- Paper F:** I. Rodriguez, **R. S. Mogensen**, A. Schjørring, M. Razzaghpour, R. Maldonado, G. Berardinelli, R. Adeogun, P. H. Christensen, P. Mogensen, O. Madsen, C. Møller, G. Pocovi, T. Kolding, C. Rosa, B. Jørgensen, and S. Barbera, "5G Swarm Production: Advanced Industrial Manufacturing Concepts Enabled by Wireless Automation", IEEE Communications Magazine, vol. 59, no. 1, pp. 48-54, Jan. 2021.
- Paper G:** **R. S. Mogensen**, C. Markmøller, T. K. Madsen, T. Kolding, G. Pocovi and M. Lauridsen, "Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications", 2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring), pp. 1-5, 2019.
- Paper H:** A. Fink, **R. S. Mogensen**, I. Rodriguez, T. Kolding, A. Karstensen, and G. Pocovi, "Radio-Aware Multi-Connectivity Solutions based on Layer-4 Scheduling for Wi-Fi in IIoT Scenarios", IEEE Wireless Communications and Networking Conference (WCNC 2022), pp. 1821-1826, 2022.
- Paper I:** **R. S. Mogensen**, S. B. Damsgaard, I. Rodriguez, G. Berardinelli, A. Fink, T. Kolding, and G. Pocovi, "A Novel QoS-Aware Multi-Connectivity Scheme for Wireless IIoT", IEEE Access, vol. 10, pp. 104123-104134, Oct. 2022.

It should be noted that, apart from the previous papers, which constitute the main body of this thesis, contributions have been made to other scientific papers. However, these were considered as side activities and are therefore not included as a main part of the thesis work.

- **Paper J:** M. Razzaghpour, R. Adeogun, I. Rodriguez, G. Berardinelli, **R. S. Mogensen**, T. Pedersen, P. Mogensen, and T. B. Sørensen, "Short-range UWB wireless channel measurement in industrial environments", 15th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), pp. 1-6, 2019.

****Paper K:** I. Rodriguez, **R. S. Mogensen**, E. J. Khatib, G. Berardinelli, P. Mogensen, O. Madsen, C. Møller, "On the Design of a Wireless MES Solution for the Factories of the Future", Global IoT Summit (GloTS), pp. 1-6, 2019.

1.5.2 Other research contributions

Besides scientific papers, this Ph.D. project has brought other relevant research contributions, which are here concisely presented.

5G Smart Production Lab

This Ph.D. study is one of the first projects related to experimental research of commercial wireless technologies within the manufacturing domain, in a framework of cross-disciplinary collaboration between the Department of Electronic Systems and the Department of Materials and Production at AAU [55]. The project had a significant contribution to the establishment of the 5G Smart Production Lab, an industrial lab designed for experimental analysis of wireless technologies for industry 4.0, which is presented in detail in paper A and F. Most of the activities presented in this thesis have indeed been possible thanks to the effort spent in setting up this experimental facility, which has also been the basis for multiple master dissertations addressing the performance of static and mobile use cases in industrial settings using 5G and Wi-Fi 6; and two ongoing Ph.D. projects looking into the "cloudification" of industrial systems based on 5G and edge-cloud computing. Moreover, it allowed the establishment of numerous external collaborations and 2 national projects related to 5G communication in the factories of the future: "Real Time Industrial IoT with 5G" (2020-2022) [56] and "5G-ROBOT, 5G-Enabled Autonomous Mobile Robotic Systems" (2022-2025) [57]. Furthermore, based on the developments of this study, the lab is now hosting an official 5G-ACIA endorsed testbed: "5G-based Zero Touch Production" (2022-2024) [58].

Experimental research tools developed

As a part of the thesis work, a set of tools has been developed to support the experimental activities. These tools are described in the following:

- **Network Sniffer:** in order to obtain the various data traffic traces used throughout the study, a SW network tap was developed based on a Raspberry Pi 4 running Linux and a USB Ethernet adapter as illustrated in Figure 1.4. It runs software that allows it to be used for network sniffing and/or network emulation (either using statistical models or

playback of latency measurements from empirical analysis). It is plug-and-play and can be installed by technical personnel on the factory floor without in-depth knowledge. More details are given in paper A.




Fig. 1.4: Picture of one of the developed network sniffers for industrial data traffic analysis and characterization.

- **Multi-access gateway:** the multi-access gateway has been developed as a tool that allows for connectivity of legacy Ethernet-based manufacturing equipment over multiple different wireless technologies (e.g., 5G and Wi-Fi 6) using either single or multi-connectivity. The main design philosophy depicted in Figure 1.5 is that each gateway should act as a port in a layer-2 switch and collectively provide the capabilities of a learning switch. Moreover, this functionality should be present regardless of the wireless access technology in a "plug-and-play" fashion by supporting both static pre-configured and ad-hoc deployment of both end devices and gateways.

1.5. Research outcomes




Fig. 1.5: Overview of layer-2 switch concept developed in the multi-access gateway.

The gateway concept has significantly evolved during the study. The first Go-lang-based prototype only supported traffic at a rate of a few 100 Kbps and had to be redesigned and re-implemented in C++, capable of providing around 300 Mbps on a 1 Gbps link. The latest version moved to Rust-lang [59] and enables the support of differentiated service-based flows based on IEEE 801.1Q header, QUIC congestion information and traffic steering using data from external plugins such as the wireless link quality. The various versions of this tool have been used in the demonstrators and deployments in collaboration with industry verticals and various AAU projects adjacent to this study. For further reference, details are given in papers A, D, H and I.

- **Measurement tools:** a generic umbrella SW measurement tool has been developed for time-synchronized state and connectivity tracking of multiple wireless interfaces with commercial cellular or Wi-Fi modems, using libqmi [60] and wpa_supplicant [61], respectively. Additionally one can enable concurrent ping RTT or trace-based performance measurements with optional location context if a source such as an AMR is present.

External collaborations

As highlighted in Section 1.3, engaging with industry verticals and demonstrating to them the potential of wireless technologies is a key aspect to the experimental research presented in this Ph.D. thesis. Therefore, it is important to highlight the main external collaborations established during the project, which are briefly described in the following:

- Some of the research activities have been disseminated to the Manufacturing Academy of Denmark (MADE), which includes members from the Danish manufacturing industry that have expressed interest in the future of manufacturing. One of the main objectives of MADE is to foster innovation in the Danish manufacturing space by facilitating and sponsoring joint research activities between companies and Danish universities. Specifically, the initial research and demonstration of Wi-Fi and LTE for non-realtime industrial usage, as well as the integration of wireless communication in brownfield deployments, have been disseminated. This initial demonstrator and knowledge dissemination have been crucial in facilitating all the subsequent collaborations in the project and served as a basis for all the in-house tools and methodology designed and developed throughout this study.
- A significant part of the work done during the study has been centered around wireless communication for AMRs facilitated by the collaboration with Mobile Industrial Robots, which is a Danish manufacturer of AMRs. Differently from AGVs, AMRs can self-navigate throughout a factory without having to rely on physical markers or lines on the floor. Instead, they are equipped with numerous sensors that enable simultaneous localization and mapping, allowing them to navigate autonomously between virtual checkpoints. In larger deployments, their ‘missions’ are orchestrated by a fleet manager, which allocates the AMRs as necessary. However, a big challenge is that the robot’s performance is subject to the available wireless infrastructure at the customer site. Here, assistance has been provided in formulating guidelines for their customer’s Wi-Fi deployment, exploring cellular connectivity, and providing input to protocol optimizations that significantly improve the scalability. A detailed description can be found in Annex A.
- One of the significant indicators of the value of the work in the project not only from an academic perspective but also from an immediate real-world applicability perspective is our involvement in the first trials of private 5G NR Release 15 in Denmark at Grundfos Bjerringbro. Here, the methodological framework served as the platform for creating a shared understanding for the partners involved (Aalborg University, Ericsson/TDC and Grundfos) on how to exploit the capabilities of 5G

NR. This included mapping/identifying traffic flows serviceable by 5G NR Release 15 and subsequently replacing the existing Ethernet links. This exercise included running parts of their production line over 5G for an extended period of time while characterizing its impact thereof. As a part of the trial, comparisons of 5G NR and Wi-Fi were also performed for scenarios involving the fleet control of AMRs as well as using AR/VR for remote support [62].

- One of the more peculiar collaborations of the study is with TV2, a Danish broadcasting company. Using robotic elements within the media production setup would allow greater flexibility from a production set perspective and a content perspective. Communication-wise, however, many use-cases overlap with traditional manufacturing in that they would like to have time-synchronization and the ability to perform isochronous tasks for e.g., coordination of lighting and robotic cameras. One area where they differ is the amount of upstream data produced; this ranges from a few kbs for light control to hundreds of Mbs for video streaming. Here, we have trialed private 5G NR Release 15 for broadcast cameras in static positions as well as under mobility. The primary learning of this trial is that hybrid deployment of mmWave and sub 6 GHz would be ideal. Furthermore, current regulation forced 1/4 or 3/7 frame structure in band n78 (at least in Denmark); this is not well suited for their use-cases as many are uplink heavy.
- Finally, the findings, tools, and methodologies developed during the study have brought Aalborg University to be accepted as a research institute partner and testbed provider in 5G-ACIA [54], a global association that brings together OT, IT and educational parties related to Industrial 5G.

1.6 Thesis outline

This Ph.D. thesis consists of five chapters, where chapters 2, 3, and 4 contain the main body of the work performed in this study. Each chapter presents a motivation of the work, objectives, main findings of the research, and finally a comment on the impact and potential follow-up studies.

- **Chapter 1: Introduction.** The current chapter provides the background for the study and the main motivations for the work presented, including outlines of the objectives, methodology, and contributions.
- **Chapter 2: Approaching industrial end users.** This chapter addresses research questions Q1 and Q2 and is related to papers A and B. Here

we present the new methodologies developed based on collaboration with external verticals. The circular framework presented in paper B, is for mapping a vertical's readiness and ideas for wireless communication in their factories/products into actionable milestones. This includes assessments of digitization maturity and current knowledge levels with regards to wireless communications, defining a suitable wireless network setup (potentially based on empirical traffic measurements), and deploying/showcasing the solution either in the factory or at AAU's premise. Furthermore, we illustrate how a digital twin-based simulation can be used to evaluate the impact of various wireless communication technologies as a cable replacement in terms of production line output.

- **Chapter 3: Performance of commercially available wireless technologies.** This chapter addresses research questions Q3 and Q4 and is related to papers C, D, E, and F. We present an analysis of selected traffic traces measured in Danish factories to understand the characteristics of the traffic. These characteristics are used to define features that need to be supported for a wireless technology used in a factory setup. Furthermore, we compare them to the modeling assumptions currently used. We then elaborate on the suitability of technologies such as Wi-Fi 5, Wi-Fi 6, LTE and 5G NR Release 15 in industrial use-cases by measuring the reliability in various network and load conditions both for static setups and under mobility.
- **Chapter 4: Transport-layer multi-connectivity enhancements for existing wireless networks.** This chapter addresses research questions Q5 and Q6 and is composed of papers G, H, I. Here the focus is mainly on how transport layer multi-connectivity can significantly improve the performance of already deployed networks such as LTE and Wi-Fi. First, we investigate the current multi-connectivity protocols and why improvements are necessary to make them suitable for industrial usage. Furthermore, we illustrate how one can provide seamless handover in Wi-Fi by using two modems in a coordinated fashion exploiting radio metrics. Finally, we show how combined radio/transport scheduling is necessary to ensure the quality of service.
- **Chapter 5: Conclusion & Future work.** This chapter presents the study's overall conclusion by addressing the hypotheses and research questions formulated in the thesis outline. Furthermore, it includes recommendations and an outlook on potential future work.

References

- [1] B. De Beelde, D. Plets, and W. Joseph, "Wireless Sensor Networks for Enabling Smart Production Lines in Industry 4.0," *Applied Sciences*, vol. 11, no. 23, p. 11248, Jan. 2021, number: 23 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/2076-3417/11/23/11248>
- [2] M. Ghobakhloo, "Industry 4.0, digitization, and opportunities for sustainability," *Journal of Cleaner Production*, vol. 252, p. 119869, Apr. 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0959652619347390>
- [3] Y. Lu, "Industry 4.0: A survey on technologies, applications and open research issues," *Journal of Industrial Information Integration*, vol. 6, pp. 1–10, Jun. 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2452414X17300043>
- [4] L. S. Dalenogare, G. B. Benitez, N. F. Ayala, and A. G. Frank, "The expected contribution of Industry 4.0 technologies for industrial performance," *International Journal of Production Economics*, vol. 204, pp. 383–394, Oct. 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0925527318303372>
- [5] Astrit Ademaj, "IEEE/IEC P60802: Traffic type characteristics," IEEE, Tech. Rep., Mar. 2019.
- [6] M. Aleksy, R. Blunk, A. Cohen, F. Dai, M. Düngen, C. Fischer, M. Gundall, R. Kirsten, M. Kus, C. Markwart, T. Neugebauer, J. Schneider, D. Schulz, and M.-P. Stanica, "TACNET 4.0 Use Cases & Requirements Specification," BMBF-PROJECT: TACNET 4.0, Tech. Rep., Apr. 2018.
- [7] 3GPP, "3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Study on NR Industrial Internet of Things (IIoT); Release 16," 3rd Generation Partnership Program, Tech. Rep. 3GPP TR 38.825 V0.0.0, Sep. 2018.
- [8] T. Carlsson, "Industrial Ethernet is now bigger than fieldbuses." [Online]. Available: <https://www.anybus.com/about-us/news/2018/02/16/industrial-ethernet-is-now-bigger-than-fieldbuses>
- [9] "The difference between AGV and AMR | MiR." [Online]. Available: <https://www.mobile-industrial-robots.com/insights/get-started-with-amrs/agv-vs-amr-whats-the-difference/>
- [10] T. Raunholt, I. Rodriguez, P. Mogensen, and M. Larsen, "Towards a 5g mobile edge cloud planner for autonomous mobile robots," in *2021 IEEE 94th Vehicular Technology Conference (VTC2021-Fall)*, 2021, pp. 01–05.
- [11] R. S. Mogensen, I. Rodriguez, G. Berardinelli, A. Fink, R. Marcker, S. Markussen, T. Raunholt, T. Kolding, G. Pocovi, and S. Barbera, "Implementation and trial evaluation of a wireless manufacturing execution system for industry 4.0," in *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*, 2019, pp. 1–7.
- [12] M. Felser, "Real Time Ethernet: Standardization and implementations," in *2010 IEEE International Symposium on Industrial Electronics*, Jul. 2010, pp. 3766–3771.
- [13] P. Danielis, J. Skodzik, V. Altmann, E. B. Schweissguth, F. Golasowski, D. Timmermann, and J. Schacht, "Survey on real-time communication via

References

- ethernet in industrial automation environments,” in *Proceedings of the 2014 IEEE Emerging Technology and Factory Automation (ETFA)*. Barcelona, Spain: IEEE, Sep. 2014, pp. 1–8. [Online]. Available: <http://ieeexplore.ieee.org/document/7005074/>
- [14] “Real Time Classes - PROFINET University.” [Online]. Available: <https://profinetuniversity.com/profinet-basics/profinet-real-time-classes/>
- [15] N. Finn, J.-Y. L. Boudec, E. Mohammadpour, B. Varga, and J. Farkas, “DetNet Bounded Latency,” Internet Engineering Task Force, Internet-Draft draft-finn-detnet-bounded-latency-01, Jul. 2018. [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-finn-detnet-bounded-latency-01>
- [16] V. K. L. Huang, Z. Pang, C. A. Chen, and K. F. Tsang, “New Trends in the Practical Deployment of Industrial Wireless: From Noncritical to Critical Use Cases,” *IEEE Industrial Electronics Magazine*, vol. 12, no. 2, pp. 50–58, Jun. 2018.
- [17] M. Wollschlaeger, T. Sauter, and J. Jasperneite, “The Future of Industrial Communication: Automation Networks in the Era of the Internet of Things and Industry 4.0,” *IEEE Industrial Electronics Magazine*, vol. 11, no. 1, pp. 17–27, Mar. 2017.
- [18] ECC, “ERC Recommendation Relating to the use of Short Range Devices (SRD,” Tech. Rep. 70-03, Jun. 2019.
- [19] “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 5: Enhancements for higher throughput,” *IEEE Std 802.11n-2009 (Amendment to IEEE Std 802.11-2007 as amended by IEEE Std 802.11k-2008, IEEE Std 802.11r-2008, IEEE Std 802.11y-2008, and IEEE Std 802.11w-2009)*, pp. 1–565, 2009.
- [20] “Ieee standard for information technology–telecommunications and information exchange between systems local and metropolitan area networks–specific requirements part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications amendment 1: Enhancements for high-efficiency wlan,” *IEEE Std 802.11ax-2021 (Amendment to IEEE Std 802.11-2020)*, pp. 1–767, 2021.
- [21] A. Varghese and D. Tandur, “Wireless requirements and challenges in Industry 4.0,” in *2014 International Conference on Contemporary Computing and Informatics (IC3I)*, Nov. 2014, pp. 634–638.
- [22] “Industrial Wireless LAN | Industrial communication | Siemens.” [Online]. Available: <https://new.siemens.com/global/en/products/automation/industrial-communication/industrial-wireless-lan.html>
- [23] D. K. Lam, K. Yamaguchi, Y. Shinozaki, S. Morita, Y. Nagao, M. Kurosaki, and H. Ochi, “A fast industrial WLAN protocol and its MAC implementation for factory communication systems,” in *2015 IEEE 20th Conference on Emerging Technologies Factory Automation (ETFA)*, Sep. 2015, pp. 1–8.
- [24] “SIEMENS - wi-Fi 6 for industry.” [Online]. Available: <https://new.siemens.com/global/en/products/automation/industrial-communication/industrial-wireless-lan/wifi6.html>
- [25] Y.-H. Wei, Q. Leng, S. Han, A. K. Mok, W. Zhang, and M. Tomizuka, “RT-WiFi: Real-Time High-Speed Communication Protocol for Wireless Cyber-Physical

References

- Control Applications,” in *2013 IEEE 34th Real-Time Systems Symposium*. Vancouver, BC, Canada: IEEE, Dec. 2013, pp. 140–149. [Online]. Available: <http://ieeexplore.ieee.org/document/6728869/>
- [26] Cisco, “IEEE 802.11ax: The Sixth Generation of Wi-Fi,” White Paper.
- [27] “MulteFire Release 1.0 Technical Paper A New Way to Wireless,” MulteFire Alliance, White Paper. [Online]. Available: <https://www.multefire.org/>
- [28] 3gpp. Study on nr-based access to unlicensed spectrum; technical report; # 38.889; release 16). [Online]. Available: https://www.3gpp.org/ftp//Specs/archive/38_series/38.889/38889-g00.zip
- [29] R. Maldonado, C. Rosa, and K. I. Pedersen, “Multi-link techniques for new radio-unlicensed urlhc in hostile environments,” in *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*, 2021, pp. 1–6.
- [30] —, “Analysis of high-reliable and low-latency communication enablers for new radio unlicensed,” in *2020 IEEE Wireless Communications and Networking Conference (WCNC)*, 2020, pp. 1–6.
- [31] 3gpp. Evolved universal terrestrial radio access (e-utra) and evolved universal terrestrial radio access network (e-utran); overall description; stage 2, release 14. [Online]. Available: https://www.3gpp.org/ftp//Specs/archive/36_series/36.300/36300-ed0.zip
- [32] —. System architecture for the 5g system (5gs) technical specification # 23.501. [Online]. Available: https://www.3gpp.org/ftp/Specs/archive/23_series/23.501/23501-h60.zip
- [33] “TDD FDD Duplex Schemes :: Radio-Electronics.Com.” [Online]. Available: https://www.radio-electronics.com/info/cellular/telecomms/cellular_concepts/tdd-fdd-time-frequency-division-duplex.php
- [34] M. Lauridsen, T. Kolding, G. Pocovi, and P. Mogensen, “Reducing Handover Outage for Autonomous Vehicles with LTE Hybrid Access,” in *2018 IEEE ICC*, pp. 1–6.
- [35] 3gpp. Nr; nr and ng-ran overall description; technical specification # 38.300; stage 2, release 15. [Online]. Available: https://www.3gpp.org/ftp/Specs/archive/38_series/38.300/38300-fd0.zip
- [36] A. Ravanshid, P. Rost, D. S. Michalopoulos, V. V. Phan, H. Bakker, D. Aziz, S. Tayade, H. D. Schotten, S. Wong, and O. Holland, “Multi-connectivity functional architectures in 5g,” in *IEEE ICC Workshop*, May 2016, pp. 187–192.
- [37] “Ieee standard for information technology– telecommunications and information exchange between systemslocal and metropolitan area networks– specific requirements–part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications–amendment 4: Enhancements for very high throughput for operation in bands below 6 ghz.” *IEEE Std 802.11ac-2013 (Amendment to IEEE Std 802.11-2012, as amended by IEEE Std 802.11ae-2012, IEEE Std 802.11aa-2012, and IEEE Std 802.11ad-2012)*, pp. 1–425, 2013.
- [38] P. A. Ribeiro, L. Duoba, R. Prior, S. Crisostomo, and L. Almeida, “Real-Time Wireless Data Plane for Real-Time-Enabled SDN,” in *2019 15th IEEE International Workshop on Factory Communication Systems (WFCS)*, May 2019, pp. 1–4.

References

- [39] L. Seno, G. Cena, S. Scanzio, A. Valenzano, and C. Zunino, "Enhancing Communication Determinism in Wi-Fi Networks for Soft Real-Time Industrial Applications," *IEEE Transactions on Industrial Informatics*, vol. 13, no. 2, pp. 866–876, Apr. 2017, conference Name: IEEE Transactions on Industrial Informatics.
- [40] A. Traßl, L. Scheuvens, T. Hößler, N. Franchi, and G. P. Fettweis, "On dependability metrics for wireless industrial communications - applied to ieee 802.11ax," in *2019 IEEE 2nd 5G World Forum (5GWF)*, 2019, pp. 286–291.
- [41] E. Avdotin, D. Bankov, E. Khorov, and A. Lyakhov, "Ofdma resource allocation for real-time applications in ieee 802.11ax networks," in *2019 IEEE International Black Sea Conference on Communications and Networking (BlackSeaCom)*, 2019, pp. 1–3.
- [42] 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 *2012 National Conference on Computing and Communication Systems*, Nov. 2012, pp. 1–5.
- [43] M. C. Lucas-Estañ, B. Coll-Perales, and J. Gozalvez, "Redundancy and Diversity in Wireless Networks to Support Mobile Industrial Applications in Industry 4.0," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 1, pp. 311–320, Jan. 2021, conference Name: IEEE Transactions on Industrial Informatics.
- [44] A. Frommgen, T. Erbschäuffer, A. Buchmann, T. Zimmermann, and K. Wehrle, "ReMP TCP: Low latency multipath TCP," in *2016 IEEE ICC*, pp. 1–7.
- [45] P. Schulz, M. Matthe, H. Klessig, M. Simsek, G. Fettweis, J. Ansari, S. A. Ashraf, B. Almeroth, J. Voigt, I. Riedel, A. Puschmann, A. Mitschele-Thiel, M. Muller, T. Elste, and M. Windisch, "Latency Critical IoT Applications in 5g: Perspective on the Design of Radio Interface and Network Architecture," *IEEE Communications Magazine*, vol. 55, no. 2, pp. 70–78, Feb. 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7842415/>
- [46] Guillermo Pocovi, Ilaria Thibault, Troels Kolding, Mads Lauridsen, Rame Canolli, Nick Edwards, and David Lister, "On the Suitability of LTE Air Interface for Reliable Low-Latency Applications," *WCNC 2019*.
- [47] T. Jacobsen, R. Abreu, G. Berardinelli, K. Pedersen, P. Mogensen, I. Z. Kovacs, and T. K. Madsen, "System Level Analysis of Uplink Grant-Free Transmission for URLLC," in *2017 IEEE Globecom Workshops (GC Wkshps)*. Singapore, Singapore: IEEE, Dec. 2017, pp. 1–6. [Online]. Available: <http://ieeexplore.ieee.org/document/8269137/>
- [48] "LTE-Advanced." [Online]. Available: <http://www.3gpp.org/technologies/keywords-acronyms/97-lte-advanced>
- [49] F. Polunin, D. C. Melgarejo, T. Lindh, A. Pinömaa, P. H. J. Nardelli, and O. Pyrhonen, "Demonstrating the Impact of LTE Communication Latency for Industrial Applications," in *2019 IEEE 17th International Conference on Industrial Informatics (INDIN)*, vol. 1, Jul. 2019, pp. 977–982, ISSN: 2378-363X.
- [50] W. Tärneberg, O. Hamsis, J. Hedlund, K. Brunnström, E. Fitzgerald, A. Johnsson, V. Berggren, M. Kihl, A. Rao, R. Steinert, and C. Kilinc, "Towards Intelligent Industry 4.0 5G Networks: A First Throughput and QoE Measurement Campaign," in

References

- 2020 *International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, Sep. 2020, pp. 1–6, iSSN: 1847-358X.
- [51] M. Hoppari, M. Uitto, J. Mäkelä, I. Harjula, and S. Rantala, “Performance of the 5th Generation Indoor Wireless Technologies-Empirical Study,” *Future Internet*, vol. 13, no. 7, p. 180, Jul. 2021, number: 7 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/1999-5903/13/7/180>
- [52] J. Rischke, P. Sossalla, S. Itting, F. H. P. Fitzek, and M. Reisslein, “5g campus networks: A first measurement study,” *IEEE Access*, vol. 9, pp. 121 786–121 803, 2021.
- [53] J. Ansari, C. Andersson, P. de Bruin, J. Farkas, L. Grosjean, J. Sachs, J. Torsner, B. Varga, D. Harutyunyan, N. König, and R. H. Schmitt, “Performance of 5G Trials for Industrial Automation,” *Electronics*, vol. 11, no. 3, p. 412, Jan. 2022, number: 3 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/2079-9292/11/3/412>
- [54] “Members – 5G-ACIA.” [Online]. Available: <https://5g-acia.org/membership/members/>
- [55] “AAU 5G Smart Production - About the Lab.” [Online]. Available: <https://www.5gsmartproduction.aau.dk/about/>
- [56] “Industriel IoT realtid med 5G - Industriens FondIndustriens Fond | Viden former fremtiden.” [Online]. Available: <https://industriensfond.dk/projekt/industriel-iot-realtid-med-5g/>
- [57] “5G-ROBOT.” [Online]. Available: <https://www.5gsmartproduction.aau.dk/5G-ROBOT/>
- [58] “5G-based Zero Touch Production – 5G-ACIA.” [Online]. Available: <https://5g-acia.org/testbeds/5g-based-zero-touch-production/>
- [59] “Rust Programming Language.” [Online]. Available: <https://www.rust-lang.org/>
- [60] “Mobile broadband connectivity / libqmi · GitLab.” [Online]. Available: <https://gitlab.freedesktop.org/mobile-broadband/libqmi>
- [61] “Linux WPA Supplicant (IEEE 802.1X, WPA, WPA2, RSN, IEEE 802.11i).” [Online]. Available: https://w1.fi/wpa_supplicant/
- [62] “5G industrial pilot from Grundfos, TDC NET and Ericsson delivers valuable insights on future possibilities of 5G.” [Online]. Available: <https://www.grundfos.com/about-us/news-and-media/news/5g-industrial-pilot-from-grundfos--tdc-net-and-ericsson-delivers>

References

Chapter 2

Approaching industrial end users

This part of the thesis focuses on the proposed approaches for integrating wireless technologies in factories and the challenges observed during collaboration with verticals.

2.1 Motivation

As this thesis revolves around experimental research of the performance of off-the-shelf wireless technologies and their applicability in Industry 4.0 (I4.0) applications, it is beneficial to include collaboration with the different verticals involved in this cross-disciplinary field. These verticals include information technology (Internet Technology (IT)) such as mobile operators, network vendors, and User Equipment (UE) vendors, as well as operational technology (Operational Technology (OT)) that supplies the manufacturing elements, e.g., mobile robots, and lastly, the end-users, i.e., the factory owner. However, one of the observations made early in the study is the existence of a significant knowledge gap regarding the capabilities of the currently commercially available versions of the different technologies and how they can be introduced in factories/manufacturing elements.

2.2 Objectives

The specific objectives related to this part of the thesis are the following:

- Assessing the current knowledge level of verticals with respect to the capabilities of current wireless technologies for their business.

- Defining an appropriate methodology that assists in smoothly integrating wireless in industries and provide knowledge and insights relevant to both IT and OT parties.
- Identifying solutions for translating the communication performance and opportunities into relevant key performance indicators from the manufacturing perspective.

2.3 Summary of contributions

The contributions in this chapter consist of two peer-reviewed papers and an annex that presents an additional study evaluating and verifying different wireless technologies for AMR communication.

Paper A: An Experimental Framework for 5G Wireless System Integration into Industry 4.0 Applications

This paper argues the need for a common understanding between communication and manufacturing sectors. We present a 6 step methodological framework that aims to provide a common starting point for the manufacturing verticals in creating/moving towards a wireless network infrastructure. The different steps aim to cover: Assessment of digitization maturity, derivation of communication requirements (if possible based on empirical measurements), selection of appropriate technology (if commercially available), deployment of a scaled test (either on-premise or in facilities), performance analysis, and optimization. Additionally, we present an overview of the various tools developed at AAU to carry out the framework's practical aspects, including performance characteristics. Furthermore, we showcase an example of how the framework and tools presented are used in exploring potential evolutionary steps in an Autonomous Mobile Robot (AMR) vendors product by transitioning to a cloud-centric software architecture enabled by 5G New Radio (5G NR).

Paper B: Evaluation of the Impact of Wireless Communication in Production via Factory Digital Twins As production lines are complex systems whose efficiency is also bounded by, e.g., physical limitations, translating the impact of wireless technology performance into production outcome is not straightforward. Furthermore, there are some feasibility aspects that need to be considered, e.g. duration of the measurement, implementation and deployment of the wireless system. Therefore, we present how one can use a digital twin of a factory with empirical measurements and simulation results to extrapolate the long-term productivity effect when replacing wired MES infrastructure with a wireless one. We use a digital twin of the AAU FESTO CP Factory production line, where we model the communication

2.4. Main findings

exchanges based on real traffic traces. To evaluate the impact of different wireless communication technologies, we sample, for each message exchange, either empirical latency distributions for Wi-Fi and Long Term Evolution (LTE) or use a simulation-based latency distribution for 5G NR. Moreover, to characterize the implication of the physical aspects of the production line, this evaluation is performed for multiple line configurations.

Annex A: MiR fleet operations: scalability study This annex presents a detailed view of the scalability study conducted in collaboration with Mobile Industrial Robots. It demonstrates how the methodological framework presented in Paper A directly impacted one of the verticals products. Specifically, we investigate intra-AP scalability by incrementally adding up to 20 robots in a warehouse setting while monitoring the performance of fleet-manager operations. To characterize the scalability, we use traffic traces of the HTTP-based fleet communication and analyze them in terms of bitrate, packets per second, and the loop latency from the fleet manager’s point of view. Based on the analysis, we propose protocol optimization steps to increase the protocol’s efficiency based on the findings. After that, we repeat the study with the new protocol implementation, comparing the scalability of Wi-Fi and LTE for single AP/base station operations, i.e. without mobility.

2.4 Main findings

As indicated in **Paper A**, throughout the numerous collaborations with verticals, it was observed that there is a significant gap when it comes to knowledge of the capabilities of commercially available wireless technologies and how to integrate them properly in existing manufacturing equipment. This stems from a number of issues. Recent experience with wireless, whether it pertains to small or big companies, is mainly related to Wi-Fi. Typically, Wi-Fi is used for office/management type of work and not for applications sensitive to jitter or handover outages. Verticals are aware of cellular technologies for mobile broadband services but not of how the system is fundamentally different from their current Wi-Fi deployments. Moreover, current cellular systems do not provide the same plug-and-play experience as what verticals are used to, be it Ethernet or Wi-Fi. Cellular systems are significantly more complex, and the implementation of the current cellular systems lacks native support of Ethernet. Therefore, additional devices are needed to facilitate the tunneling of diverse traffic types at the cost of extra complexity. Furthermore, the maturity of available terminals is relatively low when not considering cell phones.

Generally, we observed a lack of awareness regarding interference sensitivity and Listen-Before-Talk (LBT), including the significant occurrence of

outage events due to handovers. Some of these events are due to improper configurations of the network Access Point (AP)s and Station (STA)s, while others relate to protocol design. As shown in **Annex A**, we observe a bottleneck in terms of supported traffic with 15 connected robots, limiting the throughput to only 5 Mb/s. This bottleneck occurs because most of the transmitted packets are TCP control overhead packets (66 %) and thus occupy available resources unnecessarily. Via a minor restructuring in the HTTP call sequence, we can reduce the number of packets by a factor of ~ 5 , significantly improving the efficiency of the protocol and the scalability to more than 20 robots. When further comparing this protocol design in both loaded (10 Mb/s additional traffic injected) and unloaded (only robots) conditions for both LTE and Wi-Fi, we observe that LTE exhibits better scalability. The Transport Control Protocol (TCP)-Acknowledgement (ACK) Round Trip Time (RTT) only increases of a few ms for LTE in the loaded condition, whereas with Wi-Fi it increases from ~ 20 ms at the 99 %-tile to 150 ms. This difference would be emphasized even further if the evaluation included mobility.

Characterization of the long-term production impact of e.g. replacing wired links with wireless communication for a production line may be impractical. As mentioned in Section 1.4, empirical measurement can be time-consuming especially for long experiments and with multiple configurations as there may be infrastructure or time limits (the duration of the longest experiment performed at a vertical has been one week). Therefore, we show in **Paper B**, how the communication impact can be included in a simple factory digital twin, which normally assumes instantaneous data transfer. Using the digital twin, we were able to evaluate Ethernet, Wi-Fi, and LTE performance based on empirical measurements and simulated URLLC 5G NR, for a duration of a month in a fraction of the real-world time. Here, we observe less than a 1% difference in productivity between Ethernet-based control operation and any of the wireless technologies for the real-world production line, even for the most demanding services where new production orders are immediately queued. The reason for this outcome is the physical limitations of the production line, e.g. the timing at which a robot arm can perform a task. If the total production time is significantly higher than the average communication delay, the measurable impact is low. However, if this is not the case and production orders start queuing due to communication delays, the impact on production will be much greater.

2.5 Research impact and follow-up studies

At the time of the writing of this thesis, our cross-disciplinary methodological framework presented in **Paper A** has already been recognized by authors from industry and academia in multiple fields, not only related to wireless

communication and its associated performance [1–3]. For example, our work has impacted industry verticals dealing with control of AGVs [4, 5], safety enhancements [6], or even business models [7]. Naturally, our ‘front-runner’ contributions have also been noticed by other colleagues working on the integration of 5G and industrial use cases [8, 9]. Furthermore, this framework has been utilized and iterated upon in each of the collaborations presented in Section 1.5.1. A detailed example of such collaboration can also be found in **Annex A**.

Moreover, the use of factory digital twins to characterize the impact of communication presented in **Paper B** has been referenced, among others, by the authors in [10, 11] dealing with the development of edge intelligence for unmanned autonomous vehicles or optimization of human-machine interaction in smart factories. While this paper presents a relatively simple exercise, it has been necessary as part of the ‘trust-building’ with verticals.

References

- [1] K. Bhimavarapu, Z. Pang, O. Dobrijevic, and P. Wiatr, “Unobtrusive, accurate, and live measurements of network latency and reliability for time-critical internet of things,” *IEEE Internet of Things Magazine*, vol. 5, no. 3, pp. 38–43, 2022.
- [2] L. Chetot, “Activity models and Bayesian estimation algorithms for wireless grant-free random access,” Theses, Université de Lyon, Jul. 2022. [Online]. Available: <https://theses.hal.science/tel-03871656>
- [3] M.-A. Kourtis, A. Oikonomakis, D. Santorinaios, T. Anagnostopoulos, G. Xilouris, A. Kourtis, I. Chochliouros, and C. Zarakovitis, “5g npn performance evaluation for i4.0 environments,” *Applied Sciences*, vol. 12, no. 15, p. 7891, Aug 2022. [Online]. Available: <http://dx.doi.org/10.3390/app12157891>
- [4] L. Wang, Q. Liu, C. Zang, S. Zhu, C. Gan, and Y. Liu, “Formation control of dual auto guided vehicles based on compensation method in 5g networks,” *Machines*, vol. 9, no. 12, p. 318, Nov 2021. [Online]. Available: <http://dx.doi.org/10.3390/machines9120318>
- [5] A.-L. Kampen, R. Cupek, M. Fojcik, M. Drewniak, and K. Øvsthus, “Case study of agv in industry 4.0 environments – an evaluation of wireless communication protocols,” in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2022, pp. 2049–2055.
- [6] X. Yang, H. He, Z. Wei, R. Wang, K. Xu, and D. Zhang, “Enabling safety-enhanced fast charging of electric vehicles via soft actor critic-lagrange drl algorithm in a cyber-physical system,” *Applied Energy*, vol. 329, p. 120272, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S030626192201529X>
- [7] A. Psyrri, A. Kargas, and D. Varoutas, “Mnos business models and roles enabled by 5g technologies and use cases : Transformation, challenges and strategies,” in

References

- 2021 14th CMI International Conference - Critical ICT Infrastructures and Platforms (CMI), 2021, pp. 1–11.
- [8] D. Ficzer, D. Patel, J. Sachs, J. Ansari, G. Soós, and P. Varga, “5g public network integration for a real-life profinet application,” in *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*, 2022, pp. 1–5.
- [9] M. Javaid, A. Haleem, S. Rab, R. P. Singh, R. Suman, and S. Mohan, “Progressive schema of 5g for industry 4.0: features, enablers, and services,” *Industrial Robot: the international journal of robotics research and application*, 2022.
- [10] B. Yang, B. Wu, Y. You, C. Guo, L. Qiao, and Z. Lv, “Edge intelligence based digital twins for internet of autonomous unmanned vehicles,” *Software: Practice and Experience*, 2022.
- [11] F. Longo, A. Padovano, G. Aiello, C. Fusto, and A. Certa, “How 5g-based industrial iot is transforming human-centered smart factories: a quality of experience model for operator 4.0 applications,” *IFAC-PapersOnLine*, vol. 54, no. 1, pp. 255–262, 2021.

Chapter 3

Performance of commercially available wireless technologies

This part of the thesis focuses on the analysis of the off-the-shelf performance obtained using commercially available Wi-Fi 5/6, LTE, and 5G NR Release 15 in real-world industrial applications.

3.1 Motivation

5G NR standard includes URLLC features that allow the support of demanding industrial use cases, such as isochronous operations between robots. Similar capabilities are also included in the roadmap of future Wi-Fi releases (e.g. Wi-Fi 7 [1]). However, currently available commercial equipment does not support URLLC features. Still, current solutions can be used for less demanding use cases, such as supervisory control of industrial machinery, and AMR control. In order to take a first step towards a wireless factory, it is therefore of paramount importance to address the capabilities of current radio products in addressing such use cases.

3.2 Objectives

The specific objectives related to this part of the thesis are the following:

- Empirical investigation of data traffic observed in real factories.

- Design of tools that allow for seamless integration between existing manufacturing equipment and wireless technologies.
- Performance investigation of LTE, Wi-Fi, and 5G NR for static use-cases, i.e., cable replacement.
- Performance investigation and deployment considerations for Wi-Fi, LTE, and 5G NR for use cases requiring mobility.

3.3 Summary of contributions

The contributions in this chapter consist mainly of the content published in three peer-reviewed papers.

Paper C: Empirical IIoT Data Traffic Analysis and Comparison to 3GPP 5G Models

To determine the characteristic of industrial applications, we measure traffic flows between selected entities in different production cells at various Danish factories and in the facilities of an AMR vendor. To characterize applications, we analyze the traffic based on: data rate, i.e., M/k/bps, packets per second (PPS), packet inter-arrival, metadata (protocol type, unicast/broadcast, and survival time). Additionally, we compare how well the traffic models used in 3GPP fit with the observations. This includes modeling metrics used in the 3GPP models for aperiodic and fixed packet arrival and mapping them to the 5G NR QoS parameters.

Paper D: Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0

This paper presents the performance evaluation of Wi-Fi and a public LTE network for use in a delay-tolerant wireless MES deployment. This includes the description of the first iteration of the AAU multi-access gateway tool, which allows seamless integration of wireless connectivity and Ethernet-based industrial production modules. We set up a testbed with time-synchronized ingress and egress points for the performance characterization. We then monitor the one-way delay of the uplink and downlink transmission from the MES to the individual module for different technologies and network configurations. These configurations are dedicated Wi-Fi (20 MHz) where only the modules and MES host are connected, Wi-Fi with an additional 10 Mbps background load, public LTE, and finally, multi-connectivity across Wi-Fi and LTE using packet duplication.

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

This paper investigates the performance of Wi-Fi for applications requiring mobility, such as AMRs. This includes an analysis of the contributing factors to the break-before-make handover mechanism and an investigation of IEEE 802.11 amendments available in COTS-type STAs and APs addressing mobility challenges of Wi-Fi. We place a Wi-Fi STA, managed by the defacto Linux tool WPA_supplicant, on an AMR for the evaluation and let it roam between two factory halls while pinging a server and logging the Wi-Fi connectivity status. This is repeated for multiple STA configurations and network load conditions to characterize different stages' impact on the handover duration.

Paper F: 5G Swarm Production: Advanced Industrial Manufacturing Concepts enabled by Wireless Automation

This paper presents an overview of how current and future wireless technologies can cope with various visions of manufacturing. We derive communication requirements by analyzing use-cases such as cloud PLCs and large-scale AMR deployments based on traffic observations. Using these requirements in conjunction with empirical measurements, we address which radio technologies can be used for supporting the different use-cases. We also introduce the AAU 5G Smart Production Lab that implements some of these visions, including initial measurements, ranging from Wi-Fi to 5G NR Rel. 15, for a few use-cases.

3.4 Main findings

From a functional point of view, some considerations are necessary for current commercial products to support industrial applications. As observed early in **Paper C** and cemented further by measurements in **Paper D**, wireless technologies have to support full layer-2 support to provide one-to-one feature parity with existing networks in factories. This means that current cellular technologies such as 5G NR Release 15 and LTE need extra components to provide tunneling functionalities at the ingress and egress points of the networks. However, for pure IP unicast flows, as seen with, e.g., the MiR AMRs, such features are not necessary. Furthermore, the traffic analysis from operational factories suggests that composite models and packet size variability should be considered as compared to the assumptions for the FTP Model 3 (asynchronous traffic) and single periodicity (periodic traffic).

For static applications, such as delay-tolerant MES traffic with 2s survival time, a one-to-one replacement of Ethernet to Wi-Fi 5 and even public LTE is possible (subject to the MES system requirements and traffic). In general, we observe that current Wi-Fi 5/6 can, under low load in static deployments,

provide latencies under 100 ms at the 99.99 %-tile but not in loaded conditions or miss-configured network deployments with incorrect frequency allocation. Similarly, cellular technologies (5G NR Release 15, LTE) can provide sub 100 ms latencies in the same static environment, where 5G NR Rel. 15 provides the best performance of the two, as detailed in **Paper F**. **Paper E** shows that the main contributor to Wi-Fi handover is latency channel scan time, up to 100 ms per channel. Depending on the STA configuration, in a 3 AP network, we have observed that latencies between 400 ms and 1s+ at the 99.99 %-tile are achievable without relying on specific network features. Here the performance of co-located cellular networks, presented in **Papers A and F**, show that cellular technologies such as LTE and 5G NR can provide a more bounded latency closer to the 100 ms level for the exact use-case.

As are subject to the circumstances under which the measurements were performed, including the traffic characteristics, the number of devices, and network planning, **Papers D, E, and F** have limited scope. However, the results can still provide relevant insights into what kind of performance could be expected by a given technology.

3.5 Research impact and follow-up studies

Paper C presents an analysis of a selected set of traffic traces obtained from both in-house measurements at the AAU Smart Production Lab and a plethora of measurements in live production environments at external collaborators. The learnings from the various measurements have been crucial in formulating the general methodological framework presented in **Paper A**. This includes the continued evolution of the network sniffers first introduced in [2] as well as the design of the multi-access gateway, first presented in **Paper D**. Furthermore, the intended outcome of the paper was to illustrate the potential need for more advanced traffic models for 5G IIoT. The authors in [3] and [4] have confirmed and built upon our findings and proposed more advanced aggregated models for data traffic prediction for multiple industrial machinery states.

Paper D has been a crucial part of the thesis, since the setup and performance analysis of Wi-Fi and LTE in factories for delay-tolerant applications have helped initiate collaborations with verticals. The findings of this paper have been used by others authors as a reference for the selection of wireless technology candidates for industrial control [5, 6], as well as for a reference in terms of their expected performance [7–9].

Similarly, the performance results of commercially available Wi-Fi, presented in **Papers D and E**, have been used when disseminating the performance of Wi-Fi and cellular technologies to verticals. Moreover, the practical learnings regarding STA configuration have been transformed into a set of recommendations for industrial partners whose products rely on well-behaved

Wi-Fi performance.

Paper F presents an amalgamation of various research activities related to the off-the-shelf performance of commercially available implementations (ranging from Wi-Fi 5 to 5G NR). This work has been cited mainly in recent vision papers by authors active in the definition of utilization of 5G (and beyond) for Industry 4.0, such as [10–13]. The results from this paper, together with **Papers A and E**, have been important for knowledge dissemination to industrial verticals to indicate what performance can be expected from various technologies.

For further reference, a combined analysis of the performance of the different off-the-shelf wireless technologies analyzed during the study is presented in **Annex B**. These comparative results between technologies have also been paramount in the ‘trust generation’ and engagement towards industrial partners.

References

- [1] “Wi-Fi 7: The next generation in the evolution of Wi-Fi.” [Online]. Available: <https://6ghz.info/wp-content/uploads/2022/07/SenzaFili-IntelTB-Wi-Fi7-The-next-generation-in-the-evolution-of-Wi-Fi.pdf>
- [2] I. Rodriguez, R. S. Mogensen, E. J. Khatib, G. Berardinelli, P. Mogensen, O. Madsen, and C. Møller, “On the Design of a Wireless MES Solution for the Factories of the Future,” *Global IoT Summit (GloTS)*, p. 6, Jun. 2019.
- [3] A. Lieto, Q. Liao, and C. Bauer, “A generative approach for production-aware industrial network traffic modeling,” 2022. [Online]. Available: <https://arxiv.org/abs/2211.06089>
- [4] M. Lavassani, J. Åkerberg, and M. Björkman, “Modeling and profiling of aggregated industrial network traffic,” *Applied Sciences*, vol. 12, no. 2, 2022. [Online]. Available: <https://www.mdpi.com/2076-3417/12/2/667>
- [5] S. Mantravadi, R. Schnyder, C. Møller, and T. D. Brunoe, “Securing it/ot links for low power iiot devices: Design considerations for industry 4.0,” *IEEE Access*, vol. 8, pp. 200 305–200 321, 2020.
- [6] D. Segura, E. J. Khatib, and R. Barco, “Dynamic packet duplication for industrial urllc,” *Sensors*, vol. 22, no. 2, p. 587, Jan 2022. [Online]. Available: <http://dx.doi.org/10.3390/s22020587>
- [7] K. Bhimavarapu, Z. Pang, O. Dobrijevic, and P. Wiatr, “Unobtrusive, accurate, and live measurements of network latency and reliability for time-critical internet of things,” *IEEE Internet of Things Magazine*, vol. 5, no. 3, pp. 38–43, 2022.
- [8] L. Zhang, Y. Gu, R. Wang, K. Yu, Z. Pang, Y. Li, and B. Vucetic, “Enabling real-time quality-of-service and fine-grained aggregation for wireless tsn,” *Sensors*, vol. 22, no. 10, p. 3901, May 2022. [Online]. Available: <http://dx.doi.org/10.3390/s22103901>

References

- [9] L. Lu, Q. Liang, Q. Zhu, and Y. Zhao, "Synthesis of wireless networked control system based on round-trip delay online estimation," in *2020 Chinese Automation Congress (CAC)*, 2020, pp. 2729–2734.
- [10] A. Mahmood, L. Beltramelli, S. Fakhrul Abedin, S. Zeb, N. I. Mowla, S. A. Hassan, E. Sisinni, and M. Gidlund, "Industrial iot in 5g-and-beyond networks: Vision, architecture, and design trends," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 6, pp. 4122–4137, 2022.
- [11] M. Javaid, A. Haleem, S. Rab, R. P. Singh, R. Suman, and S. Mohan, "Progressive schema of 5g for industry 4.0: features, enablers, and services," *Industrial Robot: the international journal of robotics research and application*, vol. 49, no. 3, pp. 527–543, Jan 2022. [Online]. Available: <https://doi.org/10.1108/IR-10-2021-0226>
- [12] H. Wang and S. Li, ""5g+tsn+dt" solutions for digital factory key issues of networking, precision, automation and digitalization," *Journal of Physics: Conference Series*, vol. 2185, no. 1, p. 012022, jan 2022. [Online]. Available: <https://dx.doi.org/10.1088/1742-6596/2185/1/012022>
- [13] A. Mahmood, S. F. Abedin, T. Sauter, M. Gidlund, and K. Landernäs, "Factory 5G: A Review of Industrial-Centric Features and Deployment Options," 12 2021. [Online]. Available: https://www.techrxiv.org/articles/preprint/Factory_5G_A_Review_of_Industrial-Centric_Features_and_Deployment_Options/17089265

Chapter 4

Transport-layer multi-connectivity for existing wireless networks

This part of the thesis focuses on the performance enhancements that transport-layer multi-connectivity can bring to commercially available implementations or already deployed networks.

4.1 Motivation

As many verticals might already have existing wireless network deployments, one might not expect a migration toward a new radio infrastructure in the short term. For verticals that already have existing wireless networking deployments, a short-term migration to a new radio infrastructure may not be desirable. There can be multiple reasons for this, such as cost in relation to return-on-investments, spectrum availability for e.g. LTE and 5G NR, technology availability, or the fact the technology itself cannot solve the problem. Furthermore, AMR vendors may not have any influence on the network deployments of their customers. Therefore, it can be of interest to reuse existing deployments and enhance their performance. Here, transport-layer multi-connectivity offers a unique opportunity for further leveraging the capabilities of existing wireless networks, without imposing requirements on the underlying technology. Depending on the implementation, potential benefits of transport-layer multi-connectivity include mitigation of handover outages, lowering tail-latencies, providing bandwidth aggregation, and even improving network availability through network and hardware redundancy.

4.2 Objectives

The specific objectives related to this part of the thesis are the following:

- Enhancing MP-QUIC scheduling and congestion mechanisms to provide proper QoS to multiple data streams from the same application.
- Design and evaluation of a mobility management entity for Wi-Fi based on vendor-agnostic radio metrics and features.
- Design and evaluation of a framework that allows for the combination of radio and transport-layer-based scheduling to mitigate handover outage and enable multi-flow QoS.

4.3 Summary of contributions

This chapter's contributions consist mainly of the content published in three peer-reviewed papers.

Paper G: Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications

In this paper, we analyze state-of-the-art transport protocols (MP-TCP and MP-QUIC) and introduce both a novel scheduling scheme and a QoS aware congestion algorithm for MP-QUIC to enable the support both critical and bandwidth-intensive traffic flows simultaneously. When operating in conjunction, selective redundant MP-QUIC (SR-MPQUIC) and the stream aware congestion algorithm (SACA) facilitate the task of duplicating mission-critical application frames while also serving a high bandwidth background stream. To evaluate the effectiveness of the proposed solution, we compare it against the standard shortest round trip time scheduler (SRTT) as well as full redundant (FR) scheduling of all streams in terms of latency of the mission-critical data and protocol goodput. The evaluation of the different techniques is done via emulated network conditions of fixed-delay and packet loss, wide-area internet models, and empirical one-way-delay measurements of public LTE along the Danish highway.

While this paper is not specifically targeted towards industrial applications, the learning's and proposed scheduling scheme for Multipath Quick User Datagram Internet Connection (MP-QUIC) can be valid in other contexts. In industrial settings, it is possible to exploit the diversity of multiple wireless links as well as different data-flows for the same physical entity.

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

As shown in Paper E, even in the best scenario the achievable latency for single-path Wi-Fi is in the order of 400 ms. This will result in data being queued or dropped, significantly lowering the achievable reliability of a 20 Hz loop at the 99.9 %-tile, which limits the scope of the industrial application. Therefore, we design and present a mobility management entity plugin for the AAU multi-access gateway, whose main task is to enable disjoint connectivity across multiple APs and coordination of handovers in a way that the communication outage due to handovers is non-existent. Other than coordinating connectivity, the entity also informs the packet scheduler in the gate about which modem provides the best connectivity based on RSSI. Using this setup, we evaluate optimized single-path Wi-Fi against best path scheduling (BPS) (traffic steering to the modem with the highest RSSI) and full duplication. The evaluation is performed in a three AP setup in the AAU Smart Production Lab using an AMR roaming between factory halls.

Paper I: A Novel QoS-Aware Multi-Connectivity Scheme for Wireless IIoT

If the connectivity of modems is not provided by parallel networks with the correct spacing between APs a pure transport-layer approach such as SR-MPQUIC may be insufficient due to the issue of simultaneous handovers observed in Figure H.7. Therefore, we propose a further enhancement that combines the learnings from both papers G and H. The main concept is enriching the transport-layer scheduler by not only using statistics such as available bandwidth and RTT, but also radio-based metrics such as RSSI via a mobility management entity. For the initial version, presented in this paper, we implement this idea as a two-step scheduling scheme denoted as selective duplication with quality of service (SDQoS). The first step filters eligible paths for transmission based on RSSI to rank them and remove paths that are expected to perform handovers soon. After that, the second step uses the estimated path capacity along with a QoS rule of a given flow to decide whether a transmission should occur. To evaluate this scheme's ability to handle multi-flow QoS, we reuse the same setup of paper H, where instead of a single high priority flow (10 kb/s ping), we also include a background flow of (10 Mb/s UL). To investigate the impact of the different scheduling steps, we compare SDQoS with pure radio-based traffic steering, with- (SD) and without duplication (BPS) of the priority flow and single path Wi-Fi.

4.4 Main findings

As shown in **Paper G**, both MP-TCP and MP-QUIC lack support for scenarios where multiple application flows with different QoS requirements

want to leverage the same underlying connection. Here, we illustrate that in multi-flow applications where the background flow (1 Mb/s) generates significantly more data than a priority flow (100 Kb/s), selective duplication (SR-MPQUIC) offers reliability equivalent to full redundancy (FR-MPQUIC) while maintaining similar goodput as RTT based scheduling (SRTT-MPQUIC). Furthermore, in an emulated LTE network based on data collected in a highway scenario, SR-MPQUIC exhibits a 5x improvement (from 1 s to 220 ms at the 99.9 %-tile) over single path LTE and a 3.5x improvement compared to SRTT-MPQUIC (from 700 ms to 220 ms). This gain is achieved by exploiting the fact that two modems in a public network are unlikely to experience the exact same conditions, both in a connected state and during handovers. This, however, may not be the case for indoor factory networks (usually Wi-Fi). The reason for the potential discrepancy in diversity gain is a mixture of the high correlation of the received beacon signal strength experienced by e.g. two co-located modems and the fact that the modems use the same handover threshold. Whereas in public cellular networks, the network-controlled handover mechanism reduces the risk of simultaneous handovers even for co-located modems connected to the same operator.

Paper H detailed how the lack of coordination in conjunction with the device-controlled handover of Wi-Fi causes modems to connect to the same AP. This renders duplication less effective as the multiple modems contend for the same resources. However, by actively controlling the connection of each of the modems using RSSI, it becomes possible to force non-overlapping handovers. Furthermore, by providing the transport-layer protocol with the best path for traffic steering (BPS), one can achieve a gain of 55 % (from 470 ms to 225 ms) compared to the optimized single path Wi-Fi at the 99.99 %-tile. This gain improves even to 147 ms (64 %) by enabling duplication at the cost of using double network capacity.

Finally, in **Paper I**, we show that the SDQoS scheme is able to maintain a 99.99%-tile performance of 85 ms, for the priority data, in a severely loaded Wi-Fi network (25 Mb/s UL and DL for each AP), offering a more than 2x gain compared to radio-based steering with duplication (198 ms). Furthermore, it yields a 10x performance improvement compared to optimized single-path Wi-Fi (883 ms at 99.99 %-tile). While this framework is evaluated with Wi-Fi, it should still provide a gain in, e.g., LTE. However, due to the network-controlled handover and cell association, precise traffic steering and assured disjoint connectivity may require two different networks. Therefore, depending on the network deployment the gain may not be as significant for LTE.

Based on the above findings, it is clear that multi-connectivity can significantly improve the performance of an existing network. However, depending on the deployment scenario extra steps may be necessary. While simultaneous handover of modems can occur in both LTE and Wi-Fi it is much more

common for Wi-Fi, especially if modems share the same set of APs. In this case, it is possible to control the connectivity of the Wi-Fi modem, whereas in LTE, one is subject to the configuration of the network.

4.5 Research impact and follow-up studies

As both **Paper I and H** are very recent, it is still early to assess their research impact as well as their benefits for verticals. However, **Paper G** is one of the first papers which illustrate how the flexibility of QUIC and, in extension, MP-QUIC can be used to provide service to multiple flows with different QoS requirements via the proposed scheduling scheme SR-MPQUIC. This work has been used by authors further illustrating the flexibility of the proposed schemes for application of diverse nature such as control of unmanned aerial vehicles, web browsing, industrial control applications, video streaming [1–6]. While the paper only contains emulation-based results, the proposed functionalities have been verified in a real-world drive test presented in **Annex C** for further reference.

References

- [1] J. Güldenring, P. Gorczak, F. Eckermann, M. Patchou, J. Tiemann, F. Kurtz, and C. Wietfeld, “Reliable Long-Range Multi-Link Communication for Unmanned Search and Rescue Aircraft Systems in Beyond Visual Line of Sight Operation,” *Drones*, vol. 4, no. 2, p. 16, Jun. 2020, number: 2 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/2504-446X/4/2/16>
- [2] F. Nakayama, P. Lenz, A. LeFloch, A.-L. Beylot, A. Santos, and M. Nogueira, “Performance Management on Multiple Communication Paths for Portable Assisted Living,” in *2021 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, May 2021, pp. 340–348, iSSN: 1573-0077.
- [3] D. Ginhör, M.-T. Suer, M. Schüngel, R. Guillaume, and H. D. Schotten, “Survival Time-aware Dynamic Multi-connectivity for Industrial Control Applications,” in *2021 22nd IEEE International Conference on Industrial Technology (ICIT)*, vol. 1, Mar. 2021, pp. 1193–1199.
- [4] A. Sharma and D. Kamthania, “QUIC Protocol Based Monitoring Probes for Network Devices Monitor and Alerts,” in *Smart Sensor Networks: Analytics, Sharing and Control*, ser. Studies in Big Data, U. Singh, A. Abraham, A. Kaklauskas, and T.-P. Hong, Eds. Cham: Springer International Publishing, 2022, pp. 127–150. [Online]. Available: https://doi.org/10.1007/978-3-030-77214-7_6
- [5] W. Yang, J. Cao, and F. Wu, “Adaptive Video Streaming with Scalable Video Coding using Multipath QUIC,” in *2021 IEEE International Performance, Computing, and Communications Conference (IPCCC)*, Oct. 2021, pp. 1–7, iSSN: 2374-9628.

References

- [6] D. Hasselquist, C. Lindström, N. Korzhitskii, N. Carlsson, and A. Gurtov, "QUIC Throughput and Fairness over Dual Connectivity," in *Modelling, Analysis, and Simulation of Computer and Telecommunication Systems*, ser. Lecture Notes in Computer Science, M. C. Calzarossa, E. Gelenbe, K. Grochla, R. Lent, and T. Czachórski, Eds. Cham: Springer International Publishing, 2021, pp. 175–190.

Chapter 5

Conclusion & Future work

This thesis aimed at investigating the performance and integration of commercially available wireless technologies with regard to industrial applications, as well as designing enhancements to better cope with their use cases and requirements. To achieve this goal, we formulated several hypotheses and research questions in Section 1.6. Such hypotheses are related to industrial readiness and understanding of wireless communications, the performance of technologies such as Wi-Fi 5/6 and LTE/5G NR for use in different applications, and the use of transport-layer multi-connectivity to improve the capabilities of any existing networks. In the following, we recall each hypothesis to prove it based on the answers found for each related research question.

The first hypothesis **H1** relates to the readiness and potential challenges for the adoption and integration of wireless technologies by industrial verticals.

- **Q1: What are the major barriers hindering the adoption of wireless technologies by verticals?**

Based on discussions with verticals, we discovered that prior experience with wireless technologies was mainly limited to the usage of Wi-Fi for office-type applications. Even verticals that depend on customers' on-site Wi-Fi performance, such as AMR vendors, had a limited understanding of the capabilities and limitations of such technology.

When visiting some AMR end-users, we discovered that basic best practices, such as proper frequency planning, had not been followed. As shown in our study, this can significantly impact the performance in loaded networks, where we observe a $\sim x10$ increase in latency for an AMR even with optimized handover settings. Moreover, Wi-Fi APs used for office-type applications were reused for, e.g. AMRs, which is not ideal. Verticals who did, however, have a properly configured network

are aware of the consequences of LBT and handovers in Wi-Fi; still, they are unaware of the root cause, and as such their main mitigation technique is to expand the network with more access points.

Therefore, while we have observed that there exist verticals with elevated knowledge as compared to others with regards to wireless technologies, it is still not at a level where they are able in-house to discern the key characteristics of technologies such as Wi-Fi 6 and 5G NR Rel 15, and correctly deploy and use them.

- **Q2: How to ensure a gradual and efficient adoption of wireless technologies in factories?** Given the knowledge gap of the verticals, we have found that a cross-disciplinary approach would yield the best results. It provides insight into essential aspects such as technical barriers, mindsets, terminology, and understanding of wireless technologies. Here, a step-by-step framework is imperative, i.e., finding the use-case, designing a solution, and showcasing that it works in a real industrial environment. Furthermore, one should keep technology neutrality when investigating an appropriate technology/network deployment for a particular use case. This is because flexibility in choosing the technology is needed to fit all use-cases as, e.g., the spectrum may not be available, or a specific technology may present an over-complicated solution.

Based on the answers to **Q1** and **Q2**, we can conclude that new cross-disciplinary approaches are necessary as industrial verticals do not necessarily have sufficient in-house knowledge to integrate or successfully leverage wireless technologies.

The second hypothesis **H2** relates to the usability of commercially available wireless technologies such as Wi-Fi, LTE, and 5G NR Rel 15 for current industrial applications.

- **Q3: What are the main characteristics of the data traffic for the relevant industrial applications?**

Empirical evidence from real factories shows that a very diverse mix of layer-2/3 traffic can be present for the same use case and even the same application. Additionally, the observed traffic was a mixture of unicast, broadcast, and multicast. This has major implications when integrating e.g. current cellular technologies with legacy use-cases as these only support unicast IP traffic and need a translation layer. Furthermore, as the observed applications exhibit a mixture of different packet sizes and large variability of inter-arrival times, we suggest that more advanced modeling approaches than those currently used by 3GPP (e.g. aperiodic flows having fixed packets size with Poisson arrivals) may be necessary

to increase similarity with real-world observations. Such suggestions include modeling different packet sizes and the inclusion of traffic flow types, i.e., broadcast/multicast or unicast.

- **Q4: To what extent can existing off-the-shelf wireless technologies address the industrial communication needs?**

Based on the use cases and scenarios covered during the study of wireless technologies, we have observed that Wi-Fi can deliver down to ~ 18 ms at the 99.99 % in a low load network scenario. In the same scenario, 5G NR Rel. 15 and LTE can achieve ~ 21 ms and ~ 35 ms respectively. With regards to scenarios with mobility, the device-controlled handover of Wi-Fi is very sensitive to improper configuration of both STA and AP and can vary from 400 ms to multiple seconds at the 99.99 %-tile. Here, studies on 5G NR Rel. 15 and LTE illustrate the performance gain of network-controlled handover by offering a latency of 90 ms and 110 ms at the 99.99 %-tile. As such one could say that given the right conditions, applications even at the SCADA level of the automation pyramid (see Figure D.1) should be supported, though a full scalability investigation (similar to the one conducted in collaboration with the AMR vendor) of 5G NR Rel. 15 and Wi-Fi 6 is left for future work.

Here, the answers to Q3 and Q4 suggest that Wi-Fi can accommodate the industrial communication demands in low load conditions with a proper network configuration, but is very sensitive to the number of users and inter-AP mobility. On the other hand, LTE and 5G NR Rel. 15 present a more deterministic behavior in both loaded conditions and under mobility but can require extra effort for integration depending on the underlying application.

Finally, hypothesis H3 relates to the notion of using transport-layer multi-connectivity to increase the applicability of already deployed networks for the more demanding use-cases.

- **Q5: What are the limitations of current transport-layer protocols in supporting challenging communication requirements?**

Current protocol design is either geared towards applications with a single data flow, e.g. MP-TCP, and therefore unable to provide QoS if multiple flows are present, or do not support it naively, e.g. MP-QUIC. Furthermore, depending on the underlying network technologies there is a significant risk of miss-scheduling data transmission on a bad link. While in LTE the network-controlled handover de-correlates the handovers of modems even using the same operator, it was observed that for pure Wi-Fi deployments the modems, without large physical separation, will perform handovers at the same time leading to significant occurrences of outages even when using packet duplication.

- **Q6: To what extent can radio access technology metrics be used for enhancing multi-connectivity performance with respect to pure transport layer approaches?**

By combining radio quality metrics such as RSSI from multiple Wi-Fi interfaces to ensure disjoint connectivity and to provide handover mitigation via traffic steering, one can achieve seamless handover from an application point of view. Moreover, by combining traffic steering with transport-layer statistics such as the estimated bandwidth one can ensure no packets are miss-scheduled on bad paths while preserving QoS behavior for multi-flow applications. Here, the proposed SDQoS scheduling scheme yields up to a ~ 10 performance improvement at the 99.99 %-tile when compared to optimized single path Wi-Fi in loaded network conditions.

Based on the answers of **Q5** and **Q6** one can say that the proposed enhancements for transport-layer multi-connectivity such as radio-based traffic steering, can be used to mitigate technology-specific limitations, e.g. issues such as outages during Wi-Fi handovers. Furthermore, with the correct scheduling mechanisms in place, one can even support applications with multiple flows requiring different quality-of-service. Such enhancements should also yield improvements in network-controlled handover systems such as LTE and 5G as single links may quickly deteriorate. However, due to the time limitations of this study, such investigations are left for future work.

5.1 Recommendations

The following contains a set of recommendations based on the findings and learnings from the study.

- It is important for providers of cellular networks not to view industrial end-users similar to traditional subscribers of current networks. Given the complex nature of some industrial applications, an in-depth understanding from both sides is needed to facilitate collaboration and deployment successfully.
- When designing a solution for an existing application one should investigate traffic characteristics such as packet size, inter-arrival times, etc., but also consider metadata information. This includes protocol structure, network architecture, etc., as such information can provide insight into whether the application protocol is optimal for wireless transfer. If this is not the case, one could potentially improve the performance of a wireless solution without touching the network.

5.2. Future work

- In case an existing network infrastructure does not deliver satisfactory performance for an intended use case, one should consider the possibility of using transport-layer multi-connectivity. Its correct usage and implementation can indeed significantly improve the wireless usability, especially in the case of Wi-Fi networks.

5.2 Future work

One of the obvious future work directions is the extension of the performance analysis to other radio technologies. Due to the limited scope and time of the study as well as the encountered implementation challenges, an in-depth scalability analysis of 5G NR and Wi-Fi 6 has not been performed. Such analysis would not only enrich the studies presented in this work but also accommodate the wishes of some of the verticals as a direct comparison of the two technologies in different conditions has been requested multiple times. This is especially interesting as the OFDMA scheduling mechanisms of Wi-Fi 6 should provide a more deterministic behavior as compared to Wi-Fi 5, making it possibly closer to current cellular networks.

Furthermore, the use of the presented SDQoS framework should be evaluated in cellular networks as well in order to fully characterize the gain of a combined scheduling solution. Moreover, the framework should also be evaluated in use-cases where a multi-RAT network is present. Furthermore, as one of the applications investigated in the thesis has been AMRs whose location is typically known in a factory floor, it would be of interest to see whether one can fuse the AMRs telemetry together with context information such as intended route and radio metrics to perform prediction of the radio environment as the authors of [1] and [2]. This approach can simplify the handover mechanism in Wi-Fi as in principle the channel scanning is not necessary, and could provide a proactive way of applying traffic steering in multi-connectivity setups.

Also, a further investigation of traffic types that can be observed in real-world factories would be of interest. The investigation would provide input to the observations of discrepancies between 3GPP models and the observed traffic patterns highlighted in this study. This would potentially lead to the definitions of more complex models that may better capture the behavior of industrial traffic, which in turn could lead to an improved consistency between simulation studies and real-world observations.

References

- [1] M. López, T. B. Sørensen, I. Z. Kovács, J. Wigard, and P. Mogensen, "Measurement-based outage probability estimation for mission-critical services," *IEEE Access*, vol. 9, pp. 169 395–169 408, 2021.
- [2] 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 *2012 National Conference on Computing and Communication Systems*, Nov. 2012, pp. 1–5.

Part II

Annex

Annex A

MiR fleet operations: scalability analysis

This annex presents a scalability study conducted in collaboration with Mobile Industrial Robots (MiR), whose goal was to characterize how well the robot communication scaled when considering the challenges of Wi-Fi and the potential upside of providing support of cellular connectivity. While new wireless systems such as 5G NR and Wi-Fi 6 can support more challenging services than previous radio technologies, there may be a gain in further optimizing the utilization of such technologies. For example, one may examine KPIs such as the overhead of higher-layer protocols to suggest potential improvements. This is especially valuable in existing deployments that rely on current infrastructure. Here, a scalability study of the MiR AMRs revealed possibilities for performance enhancements while keeping the same wireless network infrastructure.

In the studied setup, the AMRs are managed through a central orchestrator known as a fleet manager. This entity is responsible for assigning robot tasks, keeping track of their status, and instructing charging operations. Once a robot is registered and configured by the fleet, two communication types occur, i.e., asynchronous and periodic. The asynchronous communication consists of task assignments from the fleet manager and event updates from the robot, such as checkpoint reached or notification of a completed task. The periodic traffic consists of a 1 Hz loop that polls different operational metrics from the robot. Delayed or failed communication will result in the robot standing still (asynchronous) or the fleet manager making decisions on outdated information (periodic).

The scalability investigations have been carried out in two measurement sessions:

1. At MiR premises, where 20 real commercial robots were placed in a warehouse and connected through Wi-Fi to a single AP.
2. At AAU, where we conducted a comparison test of Wi-Fi and LTE using emulators running an improved protocol design (presented later in the annex). While the emulators do not physically move, they have a one-to-one mapping with real robots regarding data exchanged with the fleet manager and are much easier to get.

The measurement setup is depicted in Figure A.1, where a robot can either be real or emulation-based.

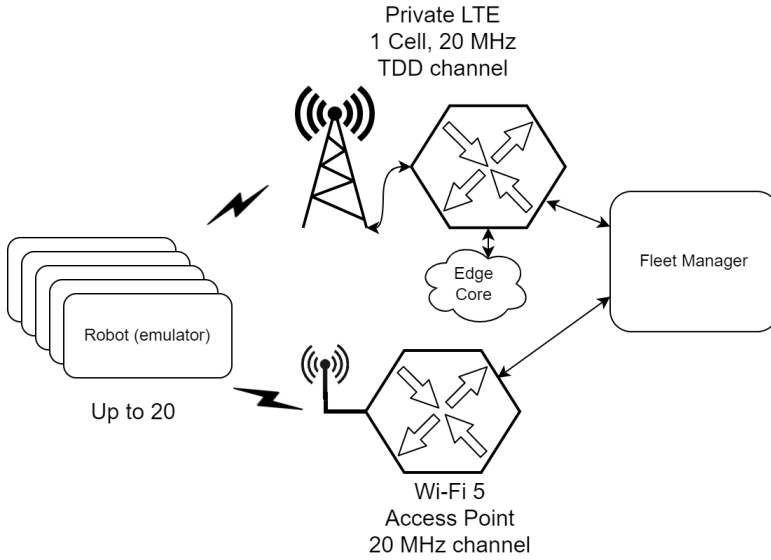


Fig. A.1: Illustration of the test setup at MiR and AAU.

We measured the traffic flows to and from a fleet manager considering an increasing number of robots (1 to 20). While the break-before-make nature of the Wi-Fi would also affect the communication, we limit the investigation to intra-AP performance only. By applying the methodological framework presented in Paper A, we first analyze the traffic patterns as the amount of connected robots increases. Looking at metrics such as packets per second (PPS) and throughput in Figures A.2 and A.3, respectively, it is clear that with the initial 'old' protocol implementation measured, the served traffic stagnates with 15 and more robots.




Fig. A.2: Packets per second for the 'old' and 'new' protocol implementations.


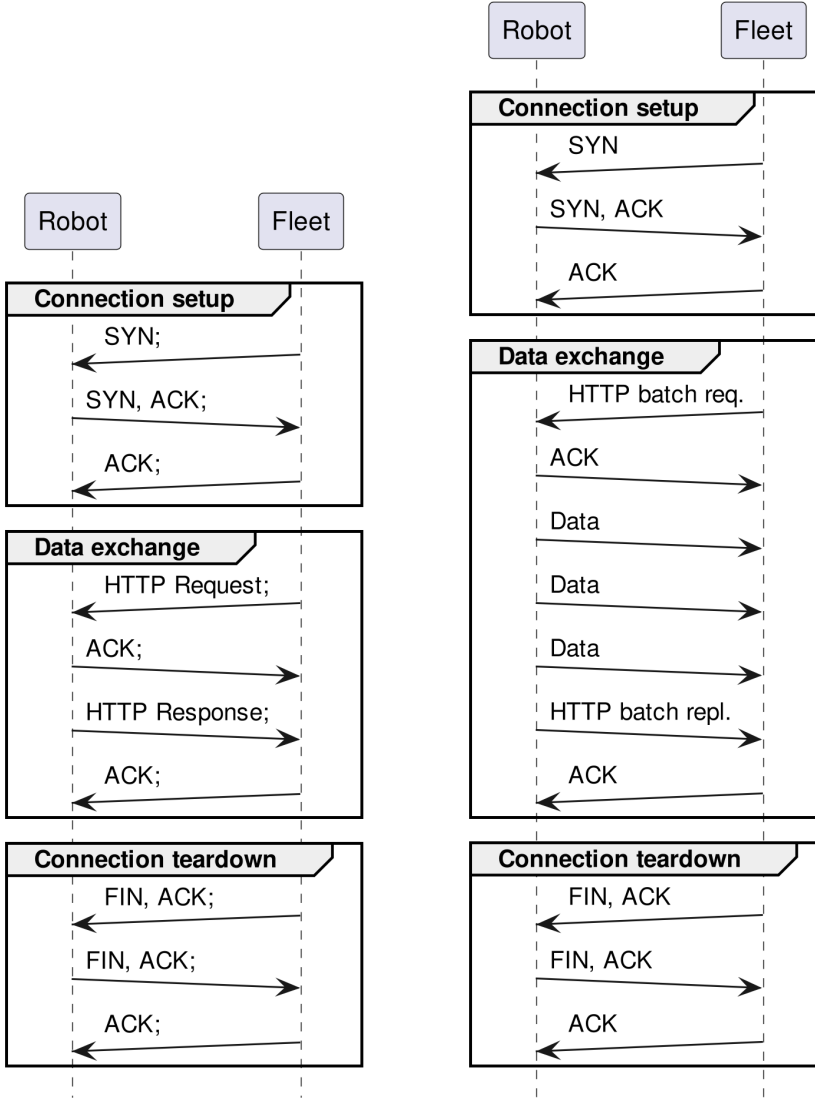


Fig. A.3: Throughput of the 'old' and 'new' protocol implementations.

In other words, these results suggest that, with the 'old' implementation, a single AP cannot serve more than 15 robots without significantly impacting the performance. One of the reasons for this behavior could be attributed to the LBT mechanism of Wi-Fi as there is simply not sufficient capacity to serve the increasing amount of packets that needs to be transmitted through the air, even though the throughput appears to be low. When filtering the traffic for



(a) Sequence diagram for the 'old' protocol design. This is repeated 7 times in total for a full loop. (b) Sequence diagram for the 'new' protocol design. This is done only once for a full loop.

Fig. A.4: Sequence diagrams for the 'old' and 'new' protocol design.

pure TCP control packets, i.e. considering no data payload, we observe that this corresponds to $\sim 66\%$ of all packets transmitted.

Figure A.4a shows a diagram of the 'old' communication protocol. This is a complete REST over HTTP call [1], that includes a full TCP connection setup and tear-down, along with the necessary data acknowledgments. From a communication point of view, this structure is inefficient since such connection setup is repeated 7 times for each periodic status update by the robot, and only 14 data packets are transmitted. Based on this analysis, one could perform the following set of optimizations with varying degrees of implementation implications:

1. Consolidation of the 7 HTTP calls to a singular one.
2. Restructure the information exchange using TCP keep-alive messages such that only a singular connection setup and tear-down can be used, though it may require some tuning to accommodate Wi-Fi handover outage properly.
3. Restructure the protocol from HTTP/JSON to raw binary encoding over TCP for increased application goodput.

For compatibility with their existing protocol design, MiR chose to only implement suggestion 1, referred as the 'new' protocol design, which should also offer the highest gain when compared to the implementation effort. Optimizing the usage of the Wi-Fi by tuning the behavior of the HTTP request as illustrated in Figure A.4b, one can significantly improve the scalability.

As shown in Figure A.5, for this 'new' protocol implementation the loop interval requirement of 1s is kept even for 20 robots, as opposed to the 'old' protocol case, where severe deviations from the target time performance were observed. However, this improvement only holds for intra-AP operations and low load deployments. Cellular communication may provide better service if the network needs to be shared with other entities, or operations happen in a large deployment requiring multiple APs. Therefore, we performed a new test with the 'new' improved protocol design for both Wi-Fi 5 and a private LTE deployment. The performance results in terms of complementary cumulative distribution function (CCDF) of the TCP ACK-RTT are shown in Figure A.6. Here we observe that while Wi-Fi provides better performance than LTE in an environment without additional load, performance degrades fast with increasing load in the network. On the other hand, LTE is not affected by the load, as the 10 Mb/s extra load setting only causes a minor increase in latency. LTE scalability is rather affected by the control channel capacity as limiting factor, since only a certain amount of users can be scheduled in the same time slot.




Fig. A.5: Control loop time for the different protocol implementations and wireless technologies.




Fig. A.6: TCP acknowledgement RTT for the 'new' protocol.

While this investigation is performed in a single cell/AP, the performance difference between LTE and Wi-Fi is expected to be further differentiated in case of mobility due to the way handovers are handled by the individual technologies, i.e., device-controlled vs. network controlled. This is further cemented by the mobility results of LTE and Wi-Fi shown in Figures B.2 (Annex B) and E.7 (Paper E).

A.1 Short conclusion

The work presented in this Annex exemplifies the importance of engaging with industry verticals which go in both ways: the private companies benefit from our expertise in operational wireless technologies for industries, which quickly help them to improve their commercial products; while we benefit from having full access to real commercial use cases and products, which make our research more targeted and valuable. In this case, the collaboration with MiR has continued after this initial study, in further activities such as some of the projects described in Section 1.5.2, where the 'cloudification' of the on-board processing elements via 5G and edge-cloud is being explored.

References

- [1] "REST over HTTP/HTTPS." [Online]. Available: <https://docs.tibco.com/pub/om-ll/5.0.0/doc/html/GUID-0DDBCC2C-2196-42E4-83E0-4A5B2B1A7CFB.html>

References

Annex B

Overview of the out-of-the-box performance of wireless technologies in operational conditions at the AAU 5G Smart Production Lab

This annex presents a brief overview of the out-of-the-box performance of the various wireless technologies examined during the study within the context of the AAU Smart production Lab activities. By out-of-box performance, we refer to the performance achieved in operational conditions right after network installation and commissioning, when no further parameter tuning or network optimization has been performed. These results should set a baseline reference performance for the selected technologies in our industrial reference environment. This overview includes results presented in papers as well as a few side standalone measurements.

In general, performed measurements can be classified into two categories, addressing different types of industrial use cases:

- **Static:** they resemble wirelessly-connected static industrial equipment, such as production line modules or robotic cell controllers.
- **Mobile:** they resemble wirelessly-connected mobile industrial equipment, such as AMRs or AGVs, or even humans.

The out-of-the-box technologies that were evaluated are the following:

- Ethernet: as the standard industrial connectivity reference.
- Wi-Fi: as the de-facto wireless standard in factories. Two versions were considered: 802.11ac (Wi-Fi 5) and 802.11ax (Wi-Fi 6). In this case, reference measurements were performed with different network load conditions: 'ideal' and 'loaded', where a single STA device or multiple un-coordinated STAs were connected to the network, respectively.
- Private 4G: as the first cellular technology allowing for a fully-isolated private enterprise deployment.
- Dedicated 4G: as the first cellular technology allowing for semi-private network infrastructure deployment and use of dedicated network resources over the public 4G network setup.
- Private 5G: as the current cellular standard allowing for a fully-isolated private enterprise deployment, and the reference technology when addressing wireless for industry. The performance tests were performed with Release 15.

The measurements were performed at the AAU 5G Smart Production Lab [1]. Static measurements were done for a static UE/STA placement at the production line modules, whereas the mobile performance was evaluated by equipping a mobile robot with a UE/STA and roaming along a measurement route covering the main aisles of the lab. Both the area where the static measurements were taken and the mobile test route are detailed in Figure B.1.



Fig. B.1: Route for mobility measurement is marked with green and the red dot is the placement during the static cases.

B.1. Short conclusion

The Complementary Cumulative Distribution (CCDF) for the full control close-loop (RTT) latency performance are detailed in Figure B.2. Although Wi-Fi 5/6 offers attractive performance in 'ideal' conditions resulting in latency levels of a few ms, in operational conditions where extra load or mobility is present, performance is degraded to tenths of ms at median level, reaching even up to a few seconds or tenths of seconds at the low percentiles (e.g., 10^{-3}). This is due to the vulnerability to high load and interference of the channel access mechanism, and uncoordinated handover mechanisms in the Wi-Fi technology.

The usage of 4G cellular system brings significant latency reduction due to the more robust scheduled access mechanism, which results in a more deterministic control-loop latency distribution, leading to latencies in the order of a few tenths of ms. This is observed both in the private and dedicated 4G cases. However, the performance in the private case is a few ms better as the latency between the RAN and the core network is negligible. At the 10^{-3} percentile, for the 4G cases, out-of-the-box performance is well contained withing 40 and 60 ms, in the static and mobile cases, respectively.

The benefits of cellular technology are further exemplified by the use of 5G. In particular, with the use of a 5G private network, latency is bounded within 30 ms, even at the 10^{-3} percentile. Future releases of 5G supporting URLLC features are expected to additionally enhance latency performance, leading to the support of closed control-loop latencies in the order of 2 ms, reducing the gap with respect to the reference Ethernet performance (sub-ms latencies). As a reference for this, Private 5G URLLC target performance results were also included in the figure.

B.1 Short conclusion

When focusing on out-of-the-box technologies, without any complex parameter tuning or optimization, cellular technologies have demonstrated to provide superior performance as compared to technologies operating in the unlicensed spectrum (e.g., Wi-Fi). This was particularly evident for the tests considering industrial use cases requiring mobility (e.g., AMRs and AGVs) as they support close to zero-latency roaming between cells. It should be noted that Wi-Fi can provide very good performance for static industrial use cases when the factory spectrum environments are fully-controlled, with light network load or extensive coordination between STAs and APs. However, as this is usually not the case, cellular technologies will provide also better performance in loaded scenarios with static UEs, which will be further emphasized with future releases of 5G (e.g., 16 and above).

Annex B. Overview of the out-of-the-box performance of wireless technologies in operational conditions at the AAU 5G Smart Production Lab




Fig. B.2: Performance overview of all technologies within the AAU Smart Production Lab under different conditions

The choice of the best wireless technology depends on the use case requirements, timeline, and budget capabilities of the specific industry vertical company. As it was demonstrated in Chapter 4, the out-of-the-box performance of the selected wireless technologies can be enhanced, which adds a further dimension to the technological roadmap considerations.

References

- [1] “AAU 5G Smart Production - About the Lab.” [Online]. Available: <https://www.5gsmartproduction.aau.dk/about/>

References

Annex C

SR-MPQUIC: real-world tests and analysis

This annex presents the results of a real-world test that builds on the SR-MPQUIC concept presented in Paper G. The concept was evaluated over a public LTE telecommunications network over a driving measurement route around Aalborg.

As the paper only presents results related to emulated network performance, it is unclear whether the same performance gains would be observed in real networks. As such, we conducted an experimental follow-up study as per the measurement campaign illustrated in Figure C.1, where the performance of the proposed protocol was evaluated under live operational network conditions. The test was conducted partially in the city center of Aalborg and in a limited highway stretch nearby.

We leveraged the framework presented in [1], which was previously used to capture the performance of individual modems. But in this case, we integrated MP-QUIC as a candidate transport protocol and extended the receiver capabilities to also log the individual network path latencies before performing the packet de-duplication. During the campaign, approximately 16k samples were collected, where the majority were taken in the urban area. The traffic model for the test was the same as the one presented in paper G i.e., 100 Kb/s priority data and 1 Mb/s background data.

The results of the priority data one-way-delay for uplink and downlink are depicted in Figures C.2 and C.3, respectively. As observed in the figures, the perceived performance of SR-MPQUIC is significantly better than either of the individual paths for uplink and downlink and the 99.9 %-tile is well contained within 100 ms.




Fig. C.1: Route of drive test in Aalborg city and on a nearby highway stretch.

C.1 Short conclusion

The results presented in this Annex, illustrate that the previously-emulated multi-connectivity gains based on the SR-MPQUIC scheduling concept are also evident in real-world operational network settings. This highlights the effectiveness of the proposed scheme, as well as it validates the applied research methodology in Section 1.4.

References




Fig. C.2: Uplink one-way-delay for the priority data in real-world drive test conditions.




Fig. C.3: Downlink one-way-delay for the priority data in real-world drive test conditions.

References

- [1] G. Pocovi, T. Kolding, M. Lauridsen, R. Mogensen, C. Markmoller, and R. Jess-Williams, "Measurement Framework for Assessing Reliable Real-Time Capabilities of Wireless Networks," *IEEE Communications Magazine*, pp. 1–8, 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8450876/>

References

Part III

Papers

Paper A

An Experimental Framework for 5G Wireless System Integration into Industry 4.0 Applications

I. Rodriguez, R. S. Mogensen, A. Fink, T. Raunholt,
S. Markussen, P. H. Christensen, G. Berardinelli,
P. Mogensen, C. Schou, and O. Madsen.

The paper has been published in the
Energies, vol. 14, no. 15, 4444, July 2021.

© 2021 by the authors; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>).

The layout has been revised.

Abstract

The fourth industrial revolution, or Industry 4.0 (I4.0), makes use of wireless technologies together with other industrial Internet-of-Things (IIoT) technologies, cyber-physical systems (CPS), and edge computing to enable the optimization and the faster re-configuration of industrial production processes. As I4.0 deployments are ramping up, the practical integration of 5G wireless systems with existing industrial applications is being explored in both Industry and Academia, in order to find optimized strategies and to develop guidelines oriented towards ensuring the success of the industrial wireless digitalization process. This paper explores the challenges arisen from such integration between industrial systems and 5G wireless, and presents a framework applicable to achieve a structured and successful integration. The paper aims at describing the different aspects of the framework such as the application operational flow and its associated tools, developed based on analytical and experimental applied research methodologies. The applicability of the framework is illustrated by addressing the integration of 5G technology into a specific industrial use case: the control of autonomous mobile robots. The results indicate that 5G technology can be used for reliable fleet management control of autonomous mobile robots in industrial scenarios, and that 5G can support the migration of the on-board path planning intelligence to the edge-cloud.

A.1 Introduction

Seamless system integration is a key aspect of the fourth industrial revolution (Industry 4.0) [1]. In this respect, 5G, one of the most advanced connectivity options, designed to achieve low latency, high reliability, flexibility, and security [2], is expected to be integrated with different cyber-physical systems (CPS), other industrial Internet-of-Things (IIoT) technologies, and cloud computing. This will bring smart factories and other industrial production environments to the next level in terms of optimization of the production processes and flexibility and re-configuration of the industrial manufacturing systems [3]. Common for most of these concepts is that they require extensive knowledge of information technologies (IT) and operational technologies (OT). Furthermore, as such, they should be viewed as a multi-disciplinary venture, making the transition to I4.0 for a particular company, site or factory not an easy task. To assist in this transition, multiple frameworks/methodologies have been proposed, mainly with the aim of defining roadmaps based on the digitalization and technological readiness levels [4–7].

These frameworks aim mainly at helping industrial manufacturing companies in identifying their levels of digital maturity, as well as guiding their digitalization efforts with the addition of reflective steps throughout the entire process. In general, these frameworks emphasize that reliable wireless con-

nectivity is the key enabler for system interconnection, allowing for tying all the factory assets together in a seamless way, and facilitating the deployment of advanced IT concepts such as big data and artificial intelligence (AI) but also OT concepts such as autonomous mobile robots or matrix production [7]. However, while some of the frameworks include considerations for the underlying communication-specific aspects, mainly in relation to the industrial application requirements, they do not go into details with the integration and performance that specific wireless technologies may have when applied to specific applications. The wireless landscape in industrial scenarios is comprised by multiple different technologies with very different capabilities, with Wi-Fi being today the de facto connectivity option [8]. However, the performance and reliability of Wi-Fi varies largely depending on the specific environment and application scenario due to its operation over unlicensed frequency bands, where the spectrum is shared with other networks. As in such case, it is necessary to “compete” to access the medium for transmission, high reliability and quality-of-service levels might not be fulfilled under certain conditions [9]. As 5G was designed to provide better performance and higher levels of reliability and, since the first commercial 5G networks for industrial use are being deployed [3], it is now time to ensure that the full potential of 5G is properly evaluated and demonstrated within the context of industrial use [10]. However, we are witnessing a lack of proper use case analyses, performance evaluations and development of guidelines addressing the integration of wireless technologies within the different I4.0 applications but also within the overall digitalization processes.

From a 5G perspective, the envisioned flexibility targeted by the I4.0 concept will require a proper assessment of the integration of the networking elements with the individual applications or use cases [11]. Thus, the objective of this paper is to propose an experimental framework for 5G wireless system integration into Industry 4.0 applications. This framework defines an industrial automation operational flow structured in different steps aiming at demonstrating, evaluating and optimizing the feasibility of operating specific industrial use cases over 5G in realistic factory conditions. The framework considers specific needs from factories or industrial entities, industrial-grade hardware, specific communication requirements and protocols, to evaluate the feasibility of operation of selected use cases over 5G and to benchmark it with alternative wireless technologies. In order to achieve that, the experimental framework is complemented by a specific industrial research lab environment and a number of 5G hardware and software prototyping tools. The specific selections of experimental methodology, available in-lab technologies, and even the specific hardware and software prototyping implementations have been partly impacted by generalization of learnings and inputs from our extensive conversations with multiple entities of the Danish manufacturing industry.

The rest of the article is organized as follows: Section A.2 discusses a number of framework design considerations based on the learnings from our conversation with the Danish Industry. Section A.3 describes the proposed experimental framework, including the operational flow, research lab facilities and hardware and software prototyping implementations. Section A.4 illustrates the applicability of the experimental framework for a specific industrial application: autonomous mobile robots (AMRs). Section A.5 provides a discussion on the potential considerations for future versions of the framework. Finally, Section A.6 concludes the paper.

A.2 Learnings from Industry and Related Framework Design Considerations

In order to guarantee that our proposed 5G system integration framework is aligned with the needs of industry, the main learnings from four years of numerous conversations and visits to different Danish factories are analyzed and considered directly into the design:

- The manufacturing industry, and especially the small and medium enterprises, has little or no experience with wireless communications. In general, some experience with Wi-Fi (non-optimized in most cases) was observed, but not with 4G, 5G, and the other wireless technologies. This is mainly due to the fact that, until recently, products and business models for their usage in factory were practically nonexistent. Thus, an initial learning process should be expected when introducing these technologies in manufacturing environments.
 - Design consideration: a common language and understanding needs to be established. Each factory and operational teams are different from others so, ideally, this should be done on an individual factory and use case basis. Benchmarking the performance of the wireless solution, not only over 5G, but also over other technologies such as Wi-Fi should help in improving the understanding of the different wireless operational possibilities. In this case, disseminating performance test results to manufacturing experts is an important aspect to consider.
- Deploying a wireless system is typically not considered in the business models, and is generally perceived as a cost and not as benefit. Moreover, due to the lack of experience with wireless communications, in most cases, a large degree of skepticism about the reliability, capabilities and potential of these technologies was detected.

- Design consideration: In order to build trust and convince the manufacturing industry about the suitability and the potential benefits of wireless applied to production, live demonstrations of wireless-operated industrial use cases are encouraged. This could take form of trials in factories or demos in lab environments, where wireless-integrated production concepts are shown to manufacturing experts.
- In general, current wired industrial factory control networks are not optimized for direct integration with wireless technologies. Integration gaps range from a non-optimal topology, to the use of non-IP traffic or transmission modes, such as broadcast, that may not be directly enabled/supported by wireless technologies. To fully exploit the benefits of wireless applied to industrial production, a re-architecture of such control networks will be an essential step.
 - Design consideration: hardware and software prototypes should be flexible enough to cope with modern (IP-based) and old (non-IP-based) communication protocols. Furthermore, the prototypes should be flexible enough to support different control architectures and modes of operation such as infrastructure mode (where end-devices communicate to a centralized entity or controller) and device-to-device communication (where the end-devices can communicate directly among themselves).
- There is a huge hype among the manufacturing industry about cloud control and cloud monitoring, which lead to a noticeable tendency of installing such systems within their production equipment without a proper performance impact assessment, as high capacity communication links available from each device to the cloud are taken for granted. However, this might not always be true, especially when operating over wireless, which could lead to some operational problems.
 - Design consideration: a proper data traffic analysis should be done before proceeding with the integration of an application with wireless in order to identify all potential communication flows in a given system. Further, a performance evaluation should be done after the wireless solution is deployed, to ensure operational correctness.
- Legacy is extremely important. Not all companies will have the chance to invest in the most advanced solutions, but still introducing a few wireless components for specific communication needs might result in a considerable gain for them.

- Design consideration: hardware and software prototypes should be able to transport over wireless legacy communication protocols that are conceived for industrial wired setups.

A.3 Experimental Framework for 5G Integration

The proposed experimental framework is aimed at providing services to industries, and is expected to be technology-unbiased. In this respect, the main actors of the framework can be, in a first moment, academic institutions interested in cross-disciplinary research on wireless and manufacturing. Initiating a dialogue with the manufacturing sector, and bringing wireless solution to factories, can indeed be mutually beneficial. The manufacturing sector can obtain an unbiased view on the wireless technologies and installation types that would better suit their specific needs. Besides, the manufacturing sector can enrich their vision for a wireless factory by leveraging the knowledge of the scientific trends brought by the academic institutions. Our dialogue with Danish factories has revealed that such unbiased view on wireless technologies is of great value for them. On the other hand, academic institutions can acquire extra knowledge on real factories setups, data traffic and production needs, besides identifying shortcomings of current wireless solutions or manufacturing concepts in addressing the factory demands. This can pave the way of future research activities on both wireless and manufacturing domains, and their integration. In the long term, we foresee opportunities where new business entities can act as middlemen between the manufacturing sector and the wireless sector (vendors and operators) for providing adequate recommendations on solutions and installations which are not tight to specific technologies or equipment/service providers and, thus, this framework can be of relevance for them.

The proposed framework is composed of: an operational flow that describes the overall sequential actions to be taken to achieve a successful operation of a given I4.0 application over 5G, an industrial research lab environment where the experimentation will take place, and the 5G hardware and software prototypes tools that are used to integrate the 5G system and the machinery associated with the selected industrial use case.

A.3.1 Industrial Automation Operational Flow

As illustrated in Figure A.1, the proposed experimental framework is based on a sequential operational flow structured in six steps. Each of the steps has a different objective and makes different use of the lab environment and prototyping tools that are described in Sections A.3.2 and A.3.3, respectively.

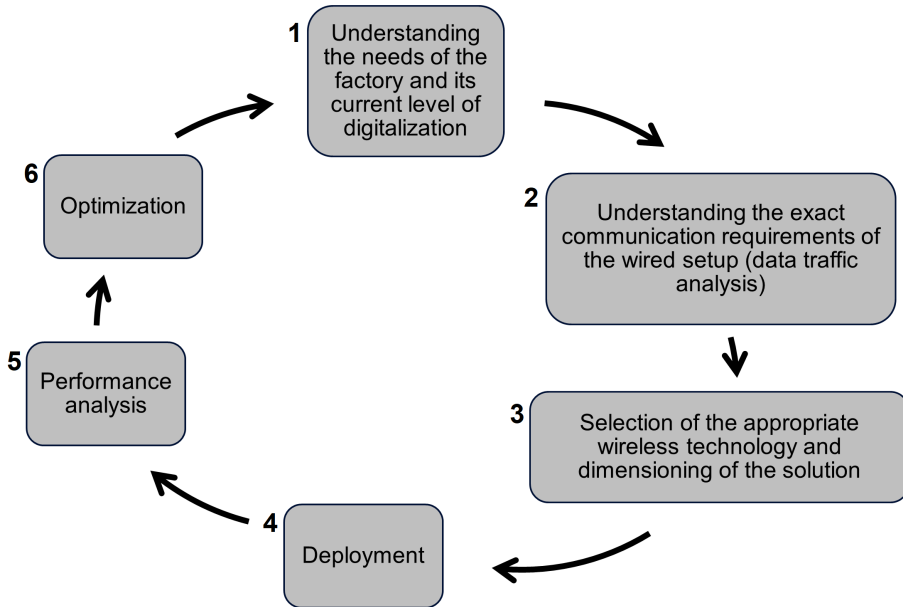


Fig. A.1: Flow diagram illustrating the different steps of the proposed industrial automation operational flow.

STEP 1. Understanding the needs of the factory and its current level of digitalization: the factory is visited in order to get a good overview of overall operations and digitalization status. It is important, that during a first visit, common ground and common languages are established. Conversations are typically initiated with management, digitalization, research and innovation responsables to learn about their views on 5G, the specific applications that they envision to run in their factories in the future, but also about those current applications that they would like to operate over wireless for improving the current production systems. Conversations are also held with IT responsibilities in order to understand the current wired and Wi-Fi network architectures, office and production network splits and network performance. With this information, we can feed the following steps of the operational flow with realistic information about the specific operational conditions required.

STEP 2. Understanding the exact communication requirements of the wired setup (data traffic analysis): in this step, that also takes place at the factory, a specific targeted I4.0 application is selected and the production area is visited—ideally, together with someone from both the IT and manufacturing departments, in order to have an overview of the exact machinery setting and its control network architecture.

A.3. Experimental Framework for 5G Integration

Further, a data traffic analysis is performed over relevant interfaces, candidates to be potentially operated over 5G, to understand and characterize the exact communication flows, protocols used, and the volume of data in different parts of the industrial system. This step is essential, as similar applications in different factories might implement very different control communication systems, and thus the more detailed the information, the better. It is important to check with the IT or production engineers about the control-loop latencies implemented in their systems, as well as survival times (sometimes called control failure times). These two parameters, together with the control architectures and the traffic statistics (throughput, packet sizes and inter-packet times) are essential information for evaluating whether the chosen application would be suitable for operation over 5G.

- STEP 3. Selection of the appropriate wireless technology and dimensioning of the solution: based on the information collected in the previous step, we can analyze the suitability of 5G or other wireless technologies for coping with the targeted application and guide the factory about the most appropriate one. These could be done by simple “pen and paper” exercises based on requirements and wireless capabilities or via simulation/emulation. In general, the decision will be made according to capacity and control-loop latency requirements. High throughput time-critical applications could only be supported by 5G, while those applications that are delay-tolerant could be supported by 5G or other alternative technologies such as 4G or Wi-Fi (under very specific and optimized operational conditions) [9]. It should be noted that some of the most demanding applications might still not be supported by 5G, e.g., applications requiring sub-ms latencies, and in that case, legacy wired setups should still be used.
- STEP 4. Deployment: once the wireless technology has been selected, i.e., 5G, the practical system integration begins. One or more of the relevant interfaces identified in STEP 2 are selected for integration of a 5G interface into them. This process is normally not straight forward as the IT integration itself, and the debugging of connectivity and 5G-integrated interface routing schemes, turns typically into a quite time-consuming task. Ideally, at the end of this step, the result is a fully functional 5G-integrated I4.0 application, both in terms of hardware and software, ready to be tested in operational conditions.
- STEP 5. Performance analysis: the fully functional 5G-integrated industrial application is thoroughly tested. In this step, typically, reliability tests are executed by monitoring the correct operation of the industrial

system during a long period of time (could be from several hours to several days). Performance tests are also done, looking at the communication efficiency in terms of, for example, control-loop latency performance. Other parameters such as throughput or time between failures could be monitored, ensuring that the outcome of the testing is well aligned with the methodology defined in [12], guaranteeing that a fair comparison with future results reported by Industry for similar use cases will be possible. Further, the industrial production efficiency could be evaluated, by mapping the empirical wireless performance to the operation of the underlying industrial process [13]. Scalability tests could also be performed, in case they are relevant for the targeted use case. Ideally, similar tests should be done over other suitable wireless technologies (i.e., 4G, Wi-Fi), in order to benchmark their performance with the 5G one and estimate the potential gains achieved by operating the application over 5G.

STEP 6. Optimization: this final step deals with the potential enhancements in those areas of the system identified based on the analysis of the outcome of the different performance tests. Depending on the results, optimizations could take place in different domains: at the machinery side, in case it is detected that an optimized performance would be feasible by a slight re-architecture of the industrial control network or by a slight tuning of parameters in the control communication protocols; at the 5G network side, by optimizing the radio interface for the specific targeted application by tuning of 5G configuration parameters; or at the 5G-application integration point, in case it is identified that changes in the integration hardware or software are needed.

Once STEP 6 in the applied operational flow is completed for a given application, this should be ready to integrate and deploy with the operational production system within the factory. Then, it is possible to go back to STEP 1 and re-initiate the flow after re-evaluating what the new needs of the factory would be after the successful completion of the previous iteration. The introduction of a wireless link may affect the entire networking ecosystem in a factory (eventually unleashing new opportunities for running wirelessly other applications involved in the same control network infrastructure), and therefore the integration of such other application and their potential conversion to wireless should be considered in the next iteration. This converts our operational flow into a circular loop, and makes the proposed experimental operational framework compatible for direct integration with those digitalization methodologies such as the AAU 360 DMA [6], briefly addressed in Section A.1. Our integration methodology is flexible enough to be applied to any application at any level digitalization level. Furthermore,

despite the fact that in this paper we illustrate and apply it mainly within a 5G context, our framework is technology-neutral. We believe that multiple wireless technologies will co-exist within the industrial ecosystem, and this makes our integration framework unbiased as compared to, for example, those that specific network providers interested in selling a given technology might suggest.

A.3.2 Industrial Research Lab Facilities

Ideally, all steps in the proposed operational flow should take place at the operational factory. However, this is not always a possibility as one may need to interrupt the manufacturing process for experimentation, which might have a negative impact on the productivity. In some cases, factories might even not have a 5G network deployed at the time when they begin to get interested in the potentials of 5G and how it could be applied for their manufacturing use cases. Therefore, having an industrial research lab with 5G capabilities where STEP 3 and beyond of our proposed flow could be executed is of paramount importance. In some specific cases, even STEPS 1 and 2 could take place in the lab, as sometimes it might be possible that representatives from the factory will visit the lab, tell about their digitalization plans and bring some of the pieces of their manufacturing system that they would like to integrate with 5G, e.g., robotic arms, programmable logical controllers (PLCs), AMRs, sensors and actuators, etc.

In those cases where 5G is already available in the factory, the practical 5G integration can take place at either place. For example, for the practical work done in STEP 4, initial efforts can be done at the research lab and, once the system is tested to be stable, the integrated version can be deployed in the factory. This will reduce significantly the impact on the manufacturing process as there would be no need for the factory to interrupt the production for experimentation.

The industrial research lab used in connection with the proposed framework is the “Aalborg University 5G Smart Production Lab”, which is a small-scale industrial factory environment composed of two halls, equipped with a wide range of operational industrial-grade manufacturing and production equipment including production lines, welding machines, robotic arms, etc. The key aspect of this research lab is that it is further equipped with a selection of the most advanced wireless technologies from multiple operators and vendors, which are fully available for our integration and testing efforts. Overview pictures of the two industrial halls of the research lab and the different wireless network deployments are shown in Figure A.2.



Fig. A.2: Overview of the Aalborg University 5G Smart Production Lab, including details on the two industrial halls and the different operational wireless network deployments.

A.3. Experimental Framework for 5G Integration

In particular, the industrial research lab is equipped with the following state of the art wireless research networks:

- 1 × 5G NR private network (5G pNR), 3 cells.
- 1 × 5G NR dedicated operator network slice (5G dNR), 3 cells.
- 2 × 4G LTE private networks (4G pLTE), 3 cells each.
- 1 × 4G LTE dedicated operator network slice (4G dLTE), 3 cells.
- 2 × Wi-Fi 6 networks, 3 cells each.
- 1 × industrial wireless LAN network (IWLAN), 3 cells.
- 1 × LoRaWAN network, 1 gateway.
- 1 × ultra-wideband radio positioning system (UWB), 16 anchors.

Having such a lab with such a selection of industrial equipment and wireless networks is an advantage for us when implementing the integration framework, as it offers a high degree of flexibility and possibilities with respect to choosing where a given step of the operational framework flow should take place. Specifically in STEPS 5, the lab offers a controlled test environment that allows to follow the guidelines from [12] for performance testing of I4.0 use cases over 5G in real-world conditions. Further, the performance achieved over 5G could be benchmarked with that achieved over 4G or Wi-Fi, for example, providing additional input into the suitability of the different wireless technologies for operating a given use case. Similarly for STEP 6, the lab network controlled environments will facilitate performing some optimizations before moving the use case to production in the factory. Further, our research lab can be used as an exhibition room where the operational 5G-integrated prototypes can be showcased to Industry, helping in generating trust on wireless when applied to manufacturing.

A.3.3 5G Hardware and Software Prototyping Tools

In order to execute our work throughout the different phases of the experimental operational flow, we developed a number of customized hardware and software prototyping tool devices:

- Network traffic sniffer (NTS): intended to log all data passing through an ethernet interface and extract relevant network traffic statistics.
- 5G Emulator (5GE): aimed at introducing controlled delays into an ethernet link to imitate the communication performance of a wireless technology, e.g., 5G.

- Wireless multi-access gateway (WMAGW): designed as the main integration element to enable the wireless transport, e.g., over 5G, of data traffic from a given ethernet-based interface.

Although there currently exist commercial devices such as network traffic analyzers [14–16], network emulators [17–19] and 5G industrial devices [20–22], with somehow similar purposes and characteristics to the ones custom-developed by us, this was not the case at the time when our 5G integration experimental activities were initiated 4 years back. As it will be explained later for each of the individual prototypes, our custom designs offer us some benefits and flexibility that would not be possible to obtain by using commercial devices.

Our reference designs consider ethernet interfaces for integration to the industrial machinery, as most of the industrial applications that will make use of 5G are currently operating over ethernet cables [23, 24]. However, it should be noted, that the design could be slightly adjusted to be integrated with other interfaces such as USB or fieldbuses.

As illustrated in Figure A.3, the NTS and the 5GE have a very similar functional structure. Both devices shared an identical hardware (HW) configuration with two bridged ethernet ports that operate by forwarding the data input at one end to the other end. Their software (SW) implementation is, however, different. The NTS (Figure A.3a) works as a “man-in-the-middle attack” [25] by executing a SW that logs relevant information from all the L2 and above data traffic packets passing through the bridge interface. Only headers are recorded. No packet inspection is done in order to preserve the confidentiality of the potential business-critical data contained in the payload of the control data messages flowing across the logged ethernet interface. The main use of the NTS box within the operational flow is done in STEP 2, when the data traffic analysis of the selected industrial application is performed. Compared with existing commercial network traffic analyzers, our NTS has similar form factor and functionality to the one developed by ProfiTAP [14]. Other network traffic analyzer solutions such as the ones from CISCO [15] or Solarwinds [16] are much more enterprise-oriented and offer extensive analysis suites, which makes perfect sense for live enterprise network monitoring but would be an overkill for our traffic characterization purpose. The main advantage of building our own analyzer is that the developed SW can be targeted to extract only relevant parameters of our interest such as number of flows, protocols, packets sizes, or packet inter-arrival times, without logging any payload with factory-critical information. This custom and controlled implementation plays in our favor as some of factories would not allow to connect commercial tools that logs and analyzes all their traffic uncontrollably.

A.3. Experimental Framework for 5G Integration

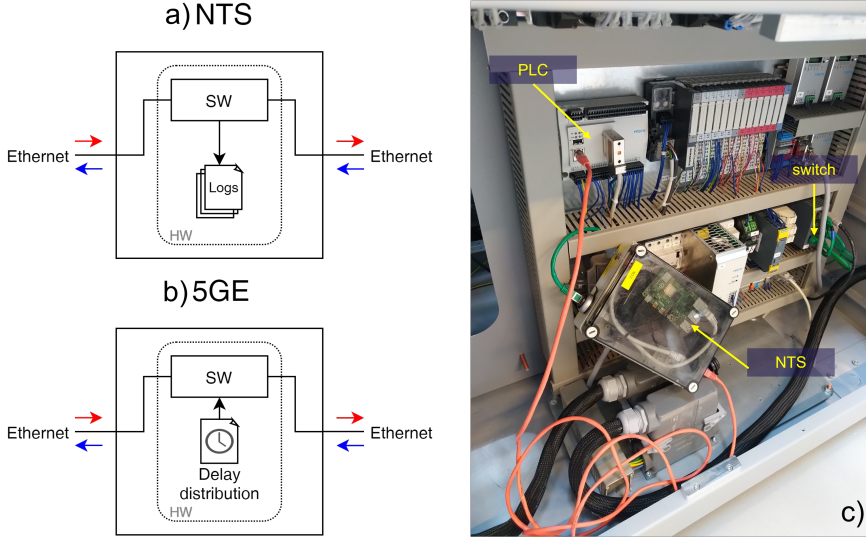


Fig. A.3: (a) Functional block diagram of the NTS. (b) Functional block diagram of the 5GE. (c) Picture of the NTS/5GE (v2) device deployed within a production line module. Under such exact same deployment configuration, the box can be used to either logging the data traffic to and from a PLC, or to introduce communication delays to emulate the connectivity of the PLC over 5G wireless.

The 5GE (Figure A.3b) implements a different SW that introduces specific delays to the bridged communication link. The packets received by the 5GE are restrained for a pre-configured delay which is obtained from the uniform sampling of a specific delay distribution pre-loaded from a configuration file. The 5GE can be configured to emulate the delay performance of different wireless technologies by simply loading specific files containing distributions of empirical delay samples for a given technology, e.g., 5G. Such files are generated by storing empirical samples from real-world lab traces such as the private 5G Rel. 15 ones obtained in [9], or from simulations for those technologies that are not available in the lab, e.g., 5G ultra reliable and low latency communication (URLLC). The interval at which the device adjusts the delay values is fully configurable, although it has not yet been possible to adjust it in a per-packet basis, due to HW limitations. The current implementation adds similar delay in both communication directions, although it should be possible to adjust the SW implementation to introduce different delays for different directions. The 5GE can be used mainly in connection with STEP 3 in the operational flow to estimate the performance of the selected use case over a given wireless technology and validate the suitability of the selection. The 5GE constitutes a simplified intermediate solution between a 5G radio channel network emulator like the one from Keysight [17] and an ethernet network

emulator such as the ones from Spirent [18] or iTrinegy [19]. The main benefit from the 5GE as compared to those commercial devices is that we can emulate 5G performance over ethernet-based links based on empirical traces obtained by ourselves in realistic and controlled operational conditions, without having to rely on pre-configured delay and jitter settings or standard channel models. To date, our 5GE implementation only allows for emulation of a single 5G link, while commercial tools are capable of emulating larger networks consisting of multiple links.

Both devices are “plug & play” and their operational procedures are very similar. In order to plug the NTS or the 5GE into industrial machinery as shown in Figure A.3c, the target ethernet cable should be disconnected from the machinery (e.g., PLC in the picture) and connected to the one of the ports of the device (e.g., the orange cable that spans from the top to the bottom of the picture). Then, the other port is connected back to the machinery with the original PLC Ethernet cable (e.g., the green cable in the middle of the picture that connects to the network switch on the left). Once the device is installed, it should be powered on and it will automatically begin its operation. As a practical recommendation, the NTS and 5GE should be used over the machinery for a time long enough to capture at least one full production cycle in the system as this would ensure that the data traffic analysis or the 5G emulation captures all possible communication states.

The functional architecture of the WMAGW is displayed in Figure A.4a. As detailed in the diagram, the WMAGW enables two different modes of connectivity: IP-based mode or L2-tunnel mode. The use of one or the other mode depends on the characteristics of the selected industrial application and its associated communication requirements, and the capabilities of the machinery that the WMAGW device will be plug into. By using the WMAGW in L2-tunnel model, we ensure that all traffic at all layers running over the ethernet interface of the production machinery is transported over wireless, e.g., 5G. This is done by encapsulating the incoming ethernet frames into UDP packets, which are later routed over the wireless port interface of the source WMAGW towards a destination WMAGW. The data sent over 5G is received at the destination WMAGW via its wireless interface, where the UDP packets are decapsulated into ethernet frames and sent over the ethernet port of the device into the ethernet interface of the production machinery. If all data control traffic in the selected application is IP-based and there are no L2 control mechanisms implemented in the system (i.e., broadcast), the IP-based mode can be used instead. This mode simplifies the integration setup and the use of the WMAGW as the SW is more computationally efficient than for the L2-tunnel mode, since we are simply bridging the ethernet and wireless interface of the gateways and establishing routing paths, but its application is not always possible especially when addressing legacy production control systems.

A.3. Experimental Framework for 5G Integration

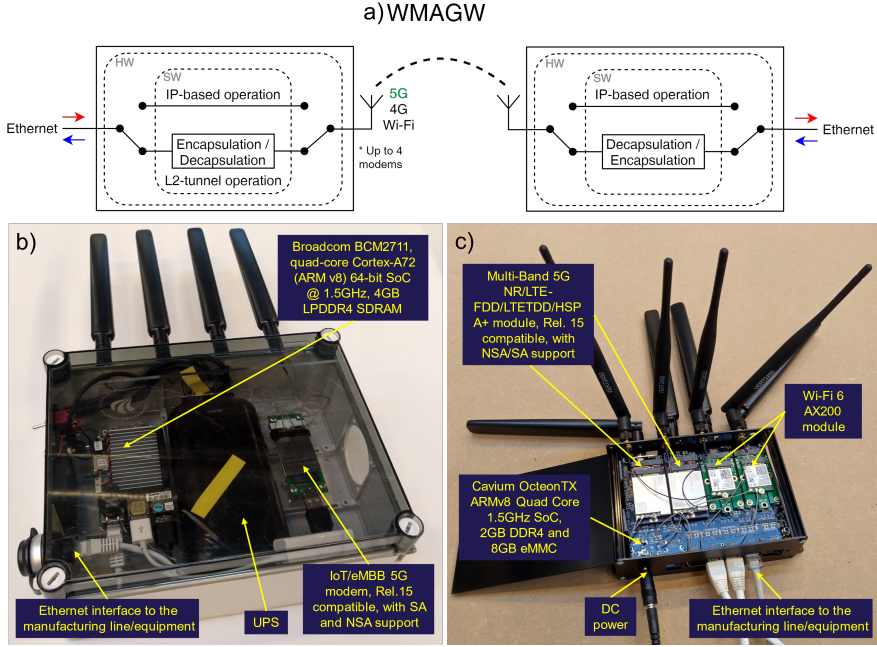


Fig. A.4: (a) Functional block diagram of the WMAGWs (operating in device-to-device L2-tunnel mode). Note: for IP-based operation, the alternative path is chosen via SW configuration. Infrastructure mode is supported by using the exam same functional blocks. (b) Picture of the implemented WMAGW prototype (v2) including details on the different HW components. (c) Picture of the implemented WMAGW prototype (v3) including details on the different HW components.

WMAGWs are used from STEP 4 and above of the operational framework flow. When utilizing the WMAGW in device-to-device communication mode, two boxes are needed; one at each of the machinery devices that are to be connected over 5G. In the case of addressing a centralized infrastructure-mode case, e.g., enabling the 5G communication between a production station and a factory controller located in a server room, only one physical WMAGW device is needed and deployed at the machinery side, while at the other end only the SW is deployed to encapsulate/decapsulate the data.

The implemented functional architecture of the different boxes is the same in all different evolution versions (v). However, the specific choice of HW components has evolved throughout our industrial research and experimentation activities and might be different from version to version, leading to different form factors, as illustrated in Figure A.4b,c for two different versions of the WMAGW (v2 and v3). As a reference for the different designs, Table A.1 summarizes the main hardware and software components implemented in the latest versions of the different devices.

Table A.1: Summary of main hardware and software components, and other relevant characteristics for the different prototyping devices.

Prototype	Hardware	Software	Other Characteristics
NTS (v2)	Raspberry Pi 4 Model B [27] Broadcom BCM2711, QC@1.5 GHz/4 GB OS: Raspberry Pi OS Lite NIC: Realtec USB3.0/Gbps ethernet adapter	tcpdump 4.99 [28]	HW TAP functionality Network transparent
5GE (v2)	Raspberry Pi 4 Model B [27] Broadcom BCM2711, QC@1.5 GHz/4 GB OS: Raspberry Pi OS Lite NIC: Realtec USB3.0/Gbps ethernet adapter	netem 2.6 [29] Custom delay generation	Multi-technology support Network transparent
WMAGW (v3)	Gateworks Newport GW6404 [30] Cavium OcteonTX8, QC@1.5GHz/2 GB OS: Ubuntu 20.04.2 NIC: Gbps ethernet port 5G modem: QUECTEL RM500Q [31] 5G modem: SIMCOM 8200 [32] 5G modem: SIMCOM 8300G [33] Wi-Fi 6 modem: Intel AX200NGW [34]	Custom multi-access tunnel Custom IP bridge	Auto-network discovery Multi-connectivity support Gbps/USB3.0 support L2 BC/MC support 5G SA/NSA support

In the specific case of the WMAGW, the modem and antenna configuration might be different depending on the HW version and the intended use of the device. The WMAGW (v3) allows for the installation of up to four modems, and in its reference design (displayed in Figure A.4c) builds up two 5G modems and two Wi-Fi six modems. With this configuration, the WMAGW allows for operation over either of the technologies individually, using one or two paths, or over both technologies enabling multi-access multi-connectivity (for increased reliability) [26]. The flexibility in the choice of operational modems and technologies is a key advantage compared to the industrial devices in the market such as the ones from HMS [20], SIEMENS [21] or Robustel [22] which typically implement and support only closed configurations. Further, the state-of-the-art 5G equipment for industrial use is still undergoing its initial commercialization phase and it is not mature enough yet and might not operate fully as expected. Currently, commercial devices might experience some of the same challenges that we typically encounter in our integration research. Initial interfacing with recently deployed and configured 5G networks typically results in connectivity issues or performance instability that requires dedicated debugging and firmware updates and tuning. Thus, being in control of the full device implementation, as it is the case with the WMAGW, is of great advantage. Being in control of the SW implementation also plays in our favor when, for example, 5G connectivity is stable but debugging is needed for the advanced L2 tunneling or multi-connectivity schemes.

Capabilities and Calibration of the 5G Prototyping Tools

A thorough calibration was done for each of the boxes in order to ensure a correct and reliable performance, quantify their capabilities, and understand their potential limitations. In the following, a number of calibration test results are presented, illustrating mainly the processing delay introduced by the devices in operational conditions. The NTS, 5GE and WMAGW were tested by feeding data traffic into the systems under different configurations of packet sizes and packet inter-arrival times (PIAT) and benchmarking the output against the input in terms of latency and throughput. A total of 18 different realizations with different traffic input configurations were tested for each device by sweeping over the combination of 6 different packet sizes (64, 128, 256, 512, 1024 and 1470) bytes (B) and 3 packet inter-arrival times (PIAT) (1, 10 and 100) ms. Calibration test results were collected over a long-enough time of operation to guarantee that we had enough statistics to claim a reliable accuracy up to the 99.9% level, i.e., 10^{-3} level in the below complementary cumulative distribution functions (CCDF) [35]. To achieve enough statistical confidence at such reliability levels, more than 100,000 samples were collected per realization, fulfilling the requirements set by normal approximation of

the Binomial distribution [36]. Further, the results illustrate the statistical distribution when considering all samples from all realizations as well as the difference between best/worse realizations to give an overview of the variability of the system.

As the NTS is typically used in connection with operational industrial machinery, it is important to verify that its use does not impact excessively the performance of the underlying industrial control communication processes while the data traffic is being logged. As illustrated in Figure A.5, the processing delay introduced by the NTS device in operational conditions is bounded to 0.45 ms, which would not affect significantly the analyzed control traffic as the introduced latency is deterministic (with small jitter <33 μ s). This jitter sets the practical limitation to what the minimum possible PIAT measured in with the NTS is. Based on the calibrated performance, it can be concluded that industrial operational machines, over which our NTS box is used, will continue to operate reliably as long as the survival-time of the underlying control protocols can tolerate a small misalignment when our device starts to function.




Fig. A.5: CCDF of the processing delay introduced by the NTS device in operational conditions.

In the case of the 5GE, as the device will be used to emulate the performance of a 5G link between two devices, apart from the processing delay of the device itself, the correctness of the generated delay outputs needs to be validated. Figure A.6 displays the results of the latency tests for two different configurations: one where no delay was added, so that the processing delay could be determined, and another one where a 1 ms constant offset delay distribution is configured (i.e., values are read from a file containing 100.000

A.3. Experimental Framework for 5G Integration

delay samples of equal value of 1 ms), in order to confirm the correct performance at the output over a constant reference. The calibrated processing delay of this device is very similar to the one observed for the NTS. This is, essentially, due to the fact that both devices are built over a similar HW configuration with a slightly different SW implementation due to their different functionalities. When the specific 1 ms delay distribution was applied to the emulator, the output was shifted by 1 ms, as expected, validating the main functionality of the device. A similar deterministic behavior was observed with/without input, although a slightly larger best/worse case variability was observed in the case the specific delay distribution was applied. This is mainly due to the extra SW execution time introduced by the implemented uniform sampling mechanisms at configured delay re-configuration intervals (100 ms in this test case). In any case, this variability is limited ($<80 \mu\text{s}$), which ensures that no excessive artificial jitter is introduced, limiting the potential distortion of the device over the input delay distributions, and resulting in a reliable delay emulation.




Fig. A.6: CCDF of the processing delay introduced by the 5GE device in operational conditions.

Finally, the performance experienced when utilizing the WMAGW devices for transmitting data traffic over the configurable L2-tunnel was examined. The total processing delay introduced by the two devices that set up the ends of the tunnel is considered in Figure A.7. In this case, despite of having the combined effect of the processing delay introduced by the encapsulating and decapsulating devices, the delay is bounded by 0.4 ms at the 99.5 percentile ($10^{-2.5}$). This reduced delay as compared to the one in the NTS and 5GE devices is explained mainly by the different HW choice as, as indicated in Table A.1, the WMAGW (v3) is

implemented over a Gateworks Newport GW6404 [30] single board computer, with higher processing power capabilities than the one of the Raspberry Pi Model 4B [27] over which the NTS (v2) and the 5GE (v2) operate. As a reference, the processing delay performance of the L2-tunnel when implemented with the previous Raspberry Pi-based WMAGW HW version (v2) is also displayed in the figure. In that case, as expected, the overall processing delay was bounded by ~ 0.75 ms (approximately twice the maximum processing delay of the NTS and 5GE). The good average processing delay performance in the WMAGW (v3) is slightly worse in some cases, where the processing delay can increase from 0.5 to up to 1.5–2 ms. This happens, at most for 0.5% of the packets transmitted over the tunnel (out of 1000 packets, only five will experience a slightly increased processing delay). As further implementation-specific details, it is of interest to report that, currently, approximately 60% of the WMAGW L2-tunnel processing delay comes from the transmission side of the tunnel where VLAN tagging and the encapsulation takes place, while the 40% remaining is due to decapsulation and packet forwarding at the reception side. There is still some room for SW enhancements, processing coordination, and execution optimization, and future versions of the WMAGW will improve further the processing delay to ensure deterministic performance at all reliability levels.




Fig. A.7: CCDF of the total processing delay experienced in L2-tunneling operational conditions (with two WMAGW devices).

As the WMAGWs are connectivity prototype devices, it was also necessary to make a more detailed analysis of the supported traffic characteristics. A maximum throughput test was performed in this case, trying to push as much data traffic as possible through the L2-tunnel by sending constant-

size packets of (64, 128, 512, 1024 and 1470) B with short PIAT intervals ranging from 10 to 100 μ s. Figure A.8 displays the results of the maximum throughput test for the different combinations of packet size and PIAT in terms of packet loss. Our WMAGW (v3) devices support reliably the encapsulation/transmission-reception/decapsulation of data packets with sizes of up to 1470 B and a minimum PIAT of 40 μ s. For more demanding traffic patterns with shorter PIAT, packet loss would start to happen due to HW buffering limitations. Therefore, the maximum throughput supported by the WMAGW devices in L2-tunnel mode is limited to 294 Mbit/s.




Fig. A.8: PER statistics for different combinations of packet size and PIAT when transmitted between two WMAGW in L2-tunneling operational conditions.

It should be emphasized that the reported WMAGW performance applies to the L2-tunnel mode of operation. When the WMAGW is operated in IP-based mode, the overall link processing delays will be shorter as bridging is a much less processing power-demanding operation than the encapsulation and VLAN tagging operations. Further, in the IP-based operation mode, the maximum supported throughput will not be depending on SW, but mainly on the board and wireless modem HW interfaces. Therefore, in IP-based mode, the WMAGW (v3) can support a throughput of up to 480 Mbit/s.

Based on the calibration test results, we guarantee the compatibility of our prototype solutions with most of the IIoT applications in terms of data traffic [37]. For quick reference, Table A.2 summarizes the main operational performance specifications of the different prototypes.

Table A.2: Calibrated operational performance specifications of the implemented prototypes.

Prototype	Processing Delay	Jitter	Max. Supported Throughput
NTS (v2)	0.45 ms	33 μ s	N/A
5GE (v2)	0.47 ms	33 μ s	N/A
WMAGW (v3) L2-tunnel	<0.4 ms (99.5%)	103 μ s	294 Mbit/s
WMAGW (v3) IP-mode	N/A	N/A	480 Mbit/s

Commercial devices are typically implemented on very optimized HW/SW platforms, which guarantees an optimized performance and thus, processing time calibration values are, normally, not reported in their specification data-sheets. With respect to maximum supported throughput, commercial network traffic analyzers and emulators are typically subject to the capabilities of the ethernet input interfaces. With respect to the commercial industrial 5G devices, in IP-mode, the limitation is typically imposed by the capabilities of the selected modems. As in our implementation, the maximum supported throughput is also reduced for commercial devices if some extra SW is executed, i.e., to operate in L2-tunnel mode. As a reference, the HMS 5G device [20] supports maximum 70 Mbit/s when a 5G virtual private network (VPN) connection is established. Our WMAGW prototype supports higher throughput in L2-tunnel mode as no extra link encryption is implemented (we rely fully on the 5G encryption for security), which results in a smaller communication protocol overhead than in the VPN case.

The described prototype tools have been successfully used in connection with different industrial production systems. The NTS has been used for recording several industrial control data traffic traces for various machines such as production cells, or quality inspection systems or in operational factories [9]. Based on those measurements, we could further validate the correct and reliable operation of the sniffer prototype. Data flow from protocols such as delay-tolerant industrial control protocols like OPC-UA or CIP, but also other more time-critical such as PROFINET, were successfully logged for several hours without detecting any disruption in the operation of the machinery [38]. To date, the WMAGW devices have been used to successfully operate, over 5G (also Wi-Fi and 4G), the control of an operational production line by providing over 5G wireless the connectivity between a factory manufacturing execution system (MES) controller and the different programmable logical controllers (PLC) present inside of each of the stations of the line [9, 39, 40].

A.4 Framework Applicability Example: 5G Autonomous Mobile Robots

In this paper, we exemplify the application of the proposed experimental framework on the integration of 5G technology to one of the most important IIoT applications: autonomous mobile robots [41].

A.4.1 5G-Connected Autonomous Mobile Robots

Following the steps of the proposed industrial automation research flow, the 5G-integration of AMRs was done as follows:

- STEP 1. An AMR manufacturer would like to start exploring the integration of their products with 5G technology. For mobility reasons, autonomous mobile robots are wireless-native elements. However, as current AMR products operate over Wi-Fi, which presents issues in terms of reliability and scalability when applied to mobile applications [42], optimized alternative are to be explored. AMR vendors are, in general, aware of the techniques for optimizing Wi-Fi performance, however, the products that they sell to their customers will operate over the Wi-Fi networks of the customers, which status remains unknown and might be unoptimized in many cases. Apart from the increased reliability, having 5G-compatible versions of their products would create new business opportunities for AMR vendors.
- STEP 2. In this case, our industrial partner brought some of their AMR products and related control elements such as the fleet manager (FM) to our lab. The control elements and system architecture were open for us so that we could perform a thorough data traffic analysis of its mode of operation. The targeted AMRs are managed from the centralized FM deployed at infrastructure-side (typically in a server room) and accessible via the Wi-Fi network. From the FM, it is possible to send specific missions with automated paths or set waypoints to navigate to for the AMRs. Although, as wireless-native components, AMRs are based on a wireless-wired hybrid architecture, different from that of the standard static production industrial use cases based only on a wired architecture, the framework can still be applied by identifying the relevant communication interfaces that are subject to be integrated into 5G. In this case, this interface was the standard AMR communication interface (the same that is used currently in the Wi-Fi version of the product), over which the data traffic flowing was analyzed with the help of the NTS device. We observed that control-loops between the AMR and the FM are executed once per second, and they make use of the TCP protocol with average packet sizes of 4 kB and 100 B, in FM-AMR and AMR-FM communication direction, respectively. The total throughput of the AMR application is, therefore, approximately 32 kbit/s (for a single robot).
- STEP 3. Based on the throughput, packet sizes and inter-packet times values, it was clear that the current AMR implementation is suitable for operation over 5G and its better mobility and reliability support should result in an improved performance of the AMRs in standard mobility operation conditions than when operated over Wi-Fi.
- STEP 4. To integrate the AMR with 5G, a WMAGW (v3) was used. The prototype device was interfaced to the standard communication interface

A.4. Framework Applicability Example: 5G Autonomous Mobile Robots

of the AMR by following the 5G-integrated architecture illustrated in Figure A.9a. The WMAGW was configured in IP-mode to operate in infrastructure mode to route the control data traffic between the AMR and the FM. In this case, as the architecture of the use case is quite simple, and there are not many communication flows, the IT integration and debugging of the 5G-integrated solution was quite quick. A picture of the deployed 5G-integrated AMR solution is shown in Figure A.9b.




Fig. A.9: (a) Architecture of the 5G-integrated AMR solution including the AMR, WMAGW, 5G network infrastructure and AMR fleet manager controller. (b) Picture of the 5G-integrated AMR solution deployed at the lab for the performance testing over the private 5G network and benchmarking with Wi-Fi 6. (c) Lab plan overview including the location of the 5G base stations and Wi-Fi 6 access points, as well as the measurement route followed to test the performance of the AMR in realistic mobility conditions.

STEP 5. To validate the correct operation of the 5G-integrated application, a performance test was carried out. The test was performed by having the AMR automatically navigating within the lab over the measurement route indicated in Figure A.9c. This specific route was chosen to

ensure that the AMR was navigating across the multiple 5G cells and guarantee that a proper assessment of the attachment/re-attachment to multiple cells and the connection reliability in mobility conditions was done. The performance of the control of the AMR over 5G was compared to the one over different Wi-Fi configurations. The performance tests focused on two main different domains: control-loop latency performance and packet error rate. The results from the tests are displayed in Figure A.10. Given the communication requirements of the AMR, with expected control loops every 1 s, all control-loops with latency higher than that will make the AMR stop and enter in a momentary outage. As indicated in the results, no outage was experienced for the 5G-integrated solution, which exhibited a reliable deterministic performance with median control-loop latency of 11 ms and lower than 25 ms at the 99.9 percentile (10^{-3}). This was, however, not the case for Wi-Fi 5/6, over which control-loop latencies higher than 1 s were experienced during 0.1–0.3% of the time in those cases where no frequency planning was considered (which resembles the typical Wi-Fi deployment situation in operational factories where the AMRs are expected to operate). During these outage occurrences, control-loop latencies close to up to 10 s were observed. These long control-loop latencies will keep the AMR stopped and non-operational for a few seconds every time they happen, which will have an impact of the underlying industrial production process in which AMRs are involved. The performance of the AMR over optimized Wi-Fi 6 and ideal frequency planning (i.e., dedicated spectrum channels per access point) was also observed to fulfill the control-loop requirements for the given use case. However, such performance was less deterministic than the one observed for the 5G-integrated solution. While a median control-loop latency, better than the 5G one, of 5 ms was observed; the value at the 99.9 percentile is increased to over 0.5 s, illustrating the higher reliability of 5G as compared to Wi-Fi. Further, the 5G-integrated solution experienced no packet loss, while in the Wi-Fi case, the PER was 0.4–0.26%.

STEP 6. Once it was verified that it is possible to operate AMRs over 5G, the next step would be to look into potential optimizations. Based on the experimental results and the practical observations done along the different operational steps, the potential areas of optimization for the 5G-integrated solution could be: the parametrization of the 5G network configuration to optimize the control-loop latency and reduce the tails of the distribution, or the re-design of the AMR communication schemes to make them more efficient and scalable when operated over wireless in general, or over 5G in particular.




Fig. A.10: CCDF of the AMR-FM control-loop latencies for the different Wi-Fi and 5G network configuration explored in the performance measurement test.

By completion of the operational flow, it was demonstrated the 5G-integration of AMRs. Apart from the commented optimizations, having 5G-connected AMRs opens for new possibilities within the same application ecosystem, such as making use of the 5G capacity and contained latency to migrate some of the intelligence from the AMR to the 5G mobile edge-cloud (MEC).

A.4.2 5G Mobile Edge-Cloud Planner for Autonomous Mobile Robots

Following the operational flow, after completion of the first iteration, a new one can be started prior re-evaluation of the new digitalization needs for AMRs. In this second iteration, the 5G-integration flow was applied as follows:

STEP 1. Once the reliable 5G-control was demonstrated to the AMR manufacturer, they were immediately interested in trying to unleash the full 5G potential by investigating whether other more heavy processes than the FM-based control could be operated over 5G. Moving part of the current on-board intelligence from the AMR to the mobile edge-cloud (MEC) would result in cheaper products and more flexible AMR products as the on-board processing power could be reduced and having a single high-processing power centralized server deployed at infrastructure-side would enable cloud-in-the-loop operation [43] towards all the AMRs, facilitating the exchange of information between AMRs and reducing the burden of having to push SW upgrades to each AMR individually. The specific application selected to be moved to the 5G MEC was the AMR planner. Currently, an AMR navigates by using its on-board planner and does not share any information with other robots. This typically results in sub-optimal navigation routes in the case that the AMR navigates into an obstacle and needs to re-plan its route, even if that same obstacle was previously detected and avoided by another robot. This problem can be mitigated by moving the planner functionalities to the edge-cloud where a centralized virtual shared-world could be built, allowing for the optimization of the navigation for all AMRs at once.

STEP 2. The AMR vendor provided us with an external planner unit and re-architected the AMR communication (mainly the I/O connections) to make use of the external device instead of its on-board planner. The new architecture is illustrated in Figure A.11a, and can be compared for further reference to the standard AMR architecture in Figure A.9a. All internal communication within the AMR happens over ethernet, so the NTS device was used to analyze the data traffic

A.4. Framework Applicability Example: 5G Autonomous Mobile Robots

in the I/O-planner communication link. TCP traffic was found with average packet sizes of 128 B and average PIAT of 0.5 ms in both the I/O-planner and planner-I/O communication directions. The overall throughput was 1.3–1.9 Mbit/s (for a single robot), which confirms that, as expected, the I/O-planner communication is more demanding than AMR-FM one.

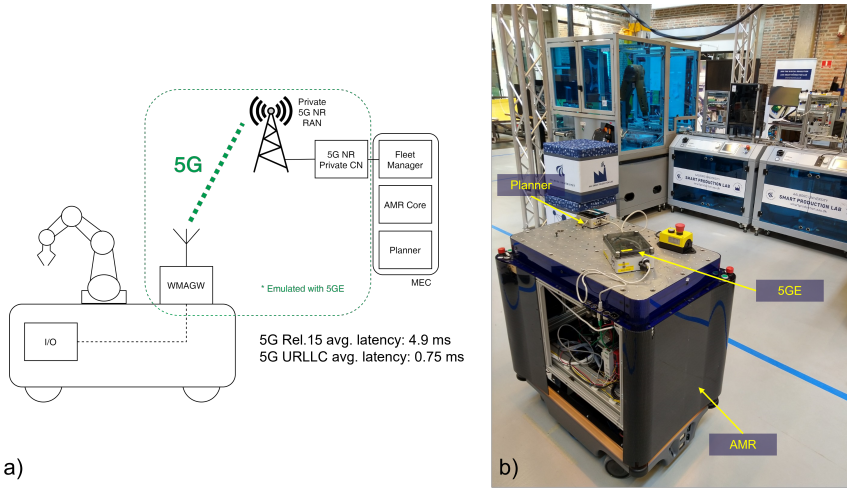


Fig. A.11: (a) Architecture of the 5G-integrated MEC AMR planner solution including the AMR, WMAGW, 5G network infrastructure, and AMR fleet manager controller and planner deployed in edge-cloud configuration. (b) Picture of the 5G-integrated MEC AMR planner solution deployed for 5G emulation at the lab.

STEP 3. Although more demanding than in the AMR-FM case, the I/O-planner link is also suitable for 5G operation, at least based on the throughput, packet sizes and inter-packet times values. However, it was questioned what the impact of the higher 5G delays (compared to the cabled ethernet reference) would be on the performance of the robot. Understanding this is of paramount importance, as the timely and reliable operation of the planner is application-critical, and an excessively long I/O-planner communication delay could result in a sub-optimal navigation performance of the robot. In this case, before proceeding to the deployment step, it was decided to make a 5G-based I/O-planner emulation test to understand the potential impact of 5G delays into the functional operation of the AMR. To do that, the 5GE device was used. As displayed in Figure A.11b, the emulator was placed between the I/O of the robot and the external planner device, mirroring the 5G MEC planner configuration illustrated in Figure A.11a. A navigation performance and a docking accuracy tests were performed for

benchmarking the operation of the reference cabled planner (REF) to the one achieved over 5G. For 5G two different delay distributions were loaded into the 5GE, one with empirical values obtained over 5G Rel. 15 (current 5G commercial version) with an average delay of 4.9 ms, and one with values from a simulation of 5G URLLC [9] with an average delay of 0.75 ms. In the navigation test, a mission was configured to make the AMR navigate over an obstructed route between two selected waypoints within the lab, ensuring that the robot needed to re-plan its path upon detection of the obstacle. The route was covered 40 times in order to ensure that the overall mission times were long enough to observe any potential differences between the reference operation with cabled planner and the 5G-emulated one. For the docking accuracy test, the robot was set to execute a docking maneuver into its charging station from a starting point located at 1 m distance. This test was repeated 15 times in order to obtain an insight into the average experienced accuracy over 5G as compared to the REF. The docking accuracy was evaluated with the assistance of an external optical positioning system with mm-accuracy [44]. The results from the navigation performance tests are presented in Figure A.12a, and they indicate that a very small increase in navigation time of 1.5% is expected in the worst case with the 5G MEC planner as compared to its reference operation with the on-board planner. For the docking accuracy test, the results are shown in A.12b, and they illustrate that the use of the 5G MEC planner would result in a slight loss of accuracy of maximum 1–3 mm, which would be negligible for reliable operation of the AMRs.




Fig. A.12: Results from the MEC-planner performance for the cabled reference case and the different 5G network configurations explored in the emulation tests: (a) navigation performance, and (b) docking performance.

Upon completion of STEP 3, where it was demonstrated by emulation that the AMR would continue to operate in an accurate and reliable manner when having its planner operating in 5G edge-cloud configuration, it would be time to move to the 5G-integration and deployment in STEP 4. However, no further details will be reported for this specific case as the work is still ongoing.

A.5 Considerations for Future Framework Versions

It has been demonstrated that applying the experimental framework leads to successful integration of 5G into industrial applications. In its current shape, the considered operational flow, prototyping tools and research lab environments, provides enough capabilities for compatibility with 5G Rel 15. However, for future 5G releases, such as those introducing URLLC features and close-to-ms control-loop supports [45], the hardware and software of the prototypes (especially the WMAGW) would need to be adjusted in order to improve, mainly, the processing delays. To achieve this, the HW prototypes will be evolved from the current central processing unit (CPU)-based designs to field-programmable gate array (FPGA) architectures to ensure an optimized processing performance. Using FPGA-based HW elements will require the development of new dedicated SW, as not all the functionalities of the CPU-based SW will be compatible with the new architecture. Such upgrade to high-end performance devices will also ensure that our future HW prototypes will be capable of supporting the integration with time-sensitive networks (TSN) aimed at the control of real-time streams with deterministic delay support [46].

The research lab would also need to be upgraded, in order to include all those 5G network release upgrades. In particular for the lab, 5G base stations with mm-wave FR2 support is seen as an essential evolution step. Under the current 5G lab deployments, operating in sub-6 GHz FR1 spectrum, available capacity might be limited for certain I4.0 applications, especially those demanding high uplink throughput for many devices, e.g., the industrial visual quality inspection based on computer vision [47]. In terms of operational flow, no updates are foreseen in the near future. The operational integration steps are described in a detailed, flexible and technology-unbiased style, which ensures the compatibility with different technologies, industrial use cases, and other reference frameworks.

Security has not been explicitly addressed in our experimental framework, as our main focus is on the 5G integration and connectivity aspects, while security is mainly addressed at higher levels [48]. 5G security in I4.0 application scenarios will depend on various aspects such as the specific flavor of 5G network deployment. A fully private 5G network deployment will have the security advantage of being isolated from the exterior (although Internet access might exist via a secure firewall), while other 5G deployments for industrial use, such as dedicated network slices, will rely on several external network

elements, and as such extra security measures need to be taken [48]. 5G radio communication makes use of encryption and is secure by design, and thus the main security challenges arise mainly from the networking perspective of the integrated end-to-end 5G-operational technology networks [49]. In this respect, our WMAGW will play a role and, thus, should be a secure interface at a similar level of current critical industrial networking infrastructure such as router or switches. Therefore, control access through physical and remote-logical interfaces should be protected by strong credentials. In terms of operation, our WMAGW should be already capable of transmitting ethernet traffic protected by high-layer cryptographic protocols such as DTLS over 5G. However, future developments in terms of industrial networking will be followed in order to ensure that our device continues to perform reliably when operating new secure ethernet protocols such as OPC-UA.

Although not related to the current integration framework itself, the experimentation could lead to identifying and highlighting potential bottlenecks of current 5G networks and, therefore, provide input to wireless, manufacturing and also automation&control engineers and researchers. Wireless researchers could then address such limitations in future radio technologies, e.g., 6G [50], while manufacturing and automation&control researchers could adapt their application and control algorithms to the wireless capabilities. Furthermore, joint manufacturing, communication and control co-design is to be explored.

A.6 Conclusions

This paper presented an experimental framework for 5G wireless system integration into industrial applications, aimed at providing service to industries, motivated from the lack of digitalization reference models considering in depth wireless performance integration and performance. The presented experimental framework consists on an operational flow that describes the different steps to be applied to move from the understanding of the needs of a factory, to the deployment and optimization of 5G-integrated applications. The framework is complemented by a number of prototyping tools: a network traffic sniffer, a 5G emulator, and a wireless multi-access gateway, which are of a great utility throughout the integration process. The prototyping devices have been carefully designed and calibrated to ensure an optimal behavior and a negligible impact when interconnected with operation industrial machines.

Further, this paper exemplifies the application of the framework on a real industrial use case: the 5G-integration of autonomous mobile robots. Two looped runs of the operational flow are applied: the first focusing on the overall 5G control of the robots, and the second one focusing on the 5G mobile edge-cloud operation of the robot path planner. In both cases, the potential of 5G for operating reliably the use case of autonomous mobile robots was

demonstrated. It was also shown how the 5G operation of the robots was superior in terms of control-loop reliability to the one of Wi-Fi 6, which resulted in several robot outages of a few seconds.

The current flow, prototyping tools and lab environment are proven to be effective at current, but the experimental framework will need to continue evolving in order to support the integration of future 5G releases, time-sensitive networks, ultra-low latency control-loop features, or even future more advanced wireless technologies such as 6G.

Author Contributions

Conceptualization, I.R.; methodology, I.R., R.S.M. and P.M.; hardware and software, R.S.M., P.H.C., A.F., T.R. and S.M.; validation, R.S.M., A.F., T.R. and S.M.; investigation, I.R., R.S.M., A.F., T.R., S.M., G.B. and C.S.; writing—original draft preparation, I.R., R.S.M. and A.F.; writing—review and editing, I.R., R.S.M., A.F., G.B. and C.S.; supervision, I.R., P.M., C.S. and O.M.; project administration, I.R.; funding acquisition, P.M. and O.M. All authors have read and agreed to the published version of the manuscript.

Funding

This research was partially funded by Industriens Fond Denmark through the project “Real-Time Industrial IoT with 5G”.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

Not applicable.

Conflicts of Interest

The authors declare no conflict of interest.

References

- [1] Aceto, G.; Persico, V.; Pescapé, A. A Survey on Information and Communication Technologies for Industry 4.0: State-of-the-Art, Taxonomies, Perspectives, and Challenges. *IEEE Commun. Surv. Tutor.* **2019**, *21*, 3467–3501.
- [2] Varga, P.; Peto, J.; Franko, A.; Balla, D.; Haja, D.; Janky, F.; Soos, G.; Ficzer, D.; Maliosz, M.; Toka, L. 5G support for Industrial IoT Applications - Challenges, Solutions, and Research gaps. *Sensors* **2020**, *20*, 828.
- [3] O’Connell, E.; Moore, D.; Neue, T. Challenges Associated with Implementing 5G in Manufacturing. *Telecom* **2020**, *1*, 8–67.
- [4] Sanchez, M.; Exposito, E.; Aguilar, J. Industry 4.0: Survey from a system integration perspective. *Int. J. Comput. Integr. Manuf.* **2020**, *33*, 1017–1041.
- [5] Hankel, M.; Rexroth, B.. *The Reference Architectural Model Industrie 4.0 (RAMI 4.0)*; version 1.0; Zentralverband Elektrotechnik- und Elektronikindustrie (ZVEI): Frankfurt am Main, Germany, 2015.
- [6] Colli, M.; Berger, U.; Bockholt, M.; Madsen, O.; Møller, C.; Wæhrens, B.V. A maturity assessment approach for conceiving context-specific roadmaps in the Industry 4.0 era. *Annu. Rev. Control* **2019**, *48*, 165–177.
- [7] Colli, M.; Madsen, O.; Berger, U.; Møller, C.; Wæhrens, B.W.; Bockholt, M. Contextualizing the outcome of a maturity assessment for Industry 4.0. *IFAC-PapersOnLine* **2018**, *51*, 1347–1352.
- [8] Seferagic, A.; Famaey, J.; De Poorter, E.; Hoebeke, J. Survey on Wireless Technology Trade-Offs for the Industrial Internet of Things. *Sensors* **2020**, *20*, 488.
- [9] Rodriguez, I.; Mogensen, R.S.; Schjørring, A.; Razzaghpour, M.; Maldonado, R.; Berardinelli, G.; Adeogun, R.; Christensen, P.H.; Mogensen, P.; Madsen, O.; et al. 5G Swarm Production: Advanced Industrial Manufacturing Concepts enabled by Wireless Automation. *IEEE Commun. Mag.*, **2021**, *59*, 48–54.
- [10] Gidlund, M.; Lennvall, T.; Akerberg, J. Will 5G become yet another wireless technology for industrial automation? In Proceedings of the IEEE International Conference on Industrial Technology (ICIT), Toronto, ON, Canada, 22–25 March 2017; pp. 1319–1324.
- [11] Spinelli, F.; Mancuso, V. Toward Enabled Industrial Verticals in 5G: A Survey on MEC-Based Approaches to Provisioning and Flexibility. *IEEE Commun. Surv. Tutor.* **2021**, *23*, 596–630.
- [12] 5G Alliance for Connected Industries and Automation (5G-ACIA). White Paper: Performance Testing of 5G Systems for Industrial Automation. Available online: <https://5g-acia.org/whitepapers/performance-testing-of-5g-systems-for-industrial-automation-2/> (accessed on 16 July 2021).

References

- [13] Mogensen, R.S.; Rodriguez, I.; Schou, C.; Mortensen, S.; Sørensen, M.S. Evaluation of the Impact of Wireless Communication in Production via Factory Digital Twins. *Manuf. Lett.* **2021**, *28*, 1–5.
- [14] ProfiTAP ProfiShark Network TAPs. Available online: <https://www.profitap.com/profishark-network-taps/> (accessed on 16 July 2021).
- [15] CISCO Security Packet Analyzer. Available online: <https://www.cisco.com/c/en/us/products/security/security-packet-analyzer/index.html> (accessed on 16 July 2021).
- [16] Solarwinds Network Traffic Analyzer. Available online: <https://www.solarwinds.com/netflow-traffic-analyzer> (accessed on 16 July 2021).
- [17] Keysight 5G Network Emulation. Available online: <https://www.keysight.com/es/en/cmp/2020/5g-network-emulation-software.html> (accessed on 16 July 2021).
- [18] Spirent Network Emulator. Available online: <https://www.spirent.com/products/network-impairment-emulation-testing> (accessed on 16 July 2021).
- [19] iTrinegy NE-ONE Professional Network Emulator. Available online: <https://itrinegy.com/ne-one-professional-range/> (accessed on 16 July 2021).
- [20] HMS Wireless Router 5G NV1000. Available online: <https://www.hms-networks.com/docs/librariesprovider6/labs/router-only.pdf> (accessed on 16 July 2021).
- [21] SIEMENS SCALANCE MUM856-1. Available online: <https://new.siemens.com/global/en/products/automation/industrial-communication/industrial-5g/5g-router-scalance-mum.html> (accessed on 16 July 2021).
- [22] Robustel 5G IoT Router R5020. Available online: <https://www.robustel.com/en/product/r5020/> (accessed on 16 July 2021).
- [23] 5G Alliance for Connected Industries and Automation (5G-ACIA). White Paper: Integration of Industrial ethernet Networks with 5G Networks. Available online: https://5g-acia.org/wp-content/uploads/2021/04/5G-ACIA_Integration-of-Industrial-Ethernet-Networks-with-5G-Networks-.pdf (accessed on 16 July 2021).
- [24] Wollschlaeger, M.; Sauter, T.; Jasperneite, J. The Future of Industrial Communication: Automation Networks in the Era of the Internet of Things and Industry 4.0. *IEEE Ind. Electron. Mag.* **2017**, *11*, 17–27.
- [25] Baud, M.; Felser, M. Profinet IO-Device Emulator based on the Man-in-the-middle Attack. In Proceedings of the IEEE Conference on Emerging Technologies and Factory Automation (ETFA), Prague, Czech Republic, 20–22 September 2006.
- [26] Suer, M.T.; Thein, C.; Tchouankem, H.; Wolf, L. Multi-Connectivity as an Enabler for Reliable Low Latency Communications—An Overview. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 156–169.
- [27] Raspberry Pi 4 Model B. Available online: <https://www.raspberrypi.org/products/raspberry-pi-4-model-b/specifications/> (accessed on 4 June 2021).

References

- [28] tcpdump. Public Repository. Available online: <https://www.tcpdump.org/> (accessed on 4 June 2021).
- [29] netem. Network Emulator. Linux Traffic Control. Available online: <https://wiki.linuxfoundation.org/networking/netem> (accessed on 4 June 2021).
- [30] Gateworks Newport GW6404. Available online: <https://www.gateworks.com/products/industrial-single-board-computers/octeon-tx-single-board-computers-gateworks-newport/gw6400-single-board-computer/> (accessed on 16 July 2021).
- [31] QUECTEL 5G RM500Q. Available online: <https://www.quectel.com/product/5g-rm500q-gl/> (accessed on 4 June 2021).
- [32] SIMCOM 5G SIM8200G. Available online: <https://www.simcom.com/product/SIM8200G.html> (accessed on 4 June 2021).
- [33] SIMCOM 5G SIM8300G. Available online: https://www.simcom.com/product/SIM8300G_M2.html (accessed on 4 June 2021).
- [34] Intel Wi-Fi 6 AX200. Available online: <https://www.intel.com/content/www/us/en/products/docs/wireless/wi-fi-6-ax200-module-brief.html> (accessed on 4 June 2021).
- [35] Herlich, M.; Maier, C. Measuring and Monitoring Reliability of Wireless Networks. *IEEE Commun. Mag.* **2021**, *59*, 76–81.
- [36] Brown, L.D.; Cai, T.T.; Dasgupta, A. Confidence intervals for a binomial proportion and asymptotic expansions. *Annu. Stat.* **2002**, *30*, 160–201.
- [37] Navarro-Ortiz, J.; Romero-Diaz, P.; Sendra, S.; Ameigeiras, P.; Ramos-Munoz, J.J.; Lopez-Soler, J.M. A survey on 5G usage scenarios and traffic models. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 905–929.
- [38] Mogensen, R.S.; Rodriguez, I.; Berardinelli, G.; Pocovi, G.; Kolding, T. Empirical IIoT Data Traffic Analysis and Comparison to 3GPP 5G Models. In Proceedings of the IEEE Vehicular Technology Conference (VCT), Online - Virtual Conference, 27–30 September 2021.
- [39] Rodriguez, I.; Mogensen, R.S.; Khatib, E.J.; Berardinelli, G.; Mogensen, P.; Madsen, O.; Møller, C. On the Design of a Wireless MES Solution for the Factories of the Future. In Proceedings of the Global IoT Summit (GIoTS), Aarhus, Denmark, 17–21 June 2019.
- [40] Mogensen, R.S.; Rodriguez, I.; Berardinelli, G.; Fink, A.; Marcker, R.; Markussen, S.A.; Raunholt, T.; Kolding, T.E.; Pocovi, G.; Barbera, S. Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0. In Proceedings of the IEEE Vehicular Technology Conference (VTC), Honolulu, HI, USA, 22–25 September 2019.
- [41] Oyekanlu, E.A.; Smith, A.C.; Thomas, W.P.; Mulroy, G.; Hitesh, D.; Ramsey, M.; Kuhn, D.J.; McGhinnis, J.D.; Buonavita, S.C.; Looper, N.A.; et al. A Review of Recent Advances in Automated Guided Vehicle Technologies: Integration Challenges and Research Areas for 5G-Based Smart Manufacturing Applications. *IEEE Access* **2020**, *8*, 202312–202353.

References

- [42] Fink, A.; Mogensen, R.S.; Rodriguez, I.; Kolding, T.; Karstensen, A.; Pocovi, G. Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility. In Proceedings of the European Wireless (EW), Verona, Italy, 10–12 November 2021.
- [43] Kehl, P.; Lange, D.; Maurer, F.K.; Nemeth, G.; Overbeck, D.; Jung, S.; Konig, N.; Schmitt, R.H. Comparison of 5G Enabled Control Loops for Production. In Proceedings of the IEEE Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Online - Virtual Conference, 31 August–3 September 2020.
- [44] Raunholt, T.; Rodriguez, I.; Mogensen, P.; Larsen, M. Towards a 5G Mobile Edge Cloud Planner for Autonomous Mobile Robots. In Proceedings of the IEEE Vehicular Technology Conference (VTC), Online - Virtual Conference, 27–30 September 2021.
- [45] Pocovi, G.; Shariatmadari, H.; Berardinelli, G.; Pedersen, K.; Steiner, J.; Li, Z. Achieving Ultra-Reliable Low-Latency Communications: Challenges and Envisioned System Enhancements. *IEEE Netw.* **2018**, *32*, 8–15.
- [46] Gundall, M.; Huber, C.; Rost, P.; Halfmann, R.; Schotten, H.D. Integration of 5G with TSN as Prerequisite for a Highly Flexible Future Industrial Automation: Time Synchronization based on IEEE 802.1AS. In Proceedings of the Annual Conference of the IEEE Industrial Electronics Society (IECON), Online - Virtual Conference, 18–21 October 2020.
- [47] Lu, T.; Fan, Z.; Lei, Y.; Shang, Y.; Wang, C. The Edge Computing Cloud Architecture Based on 5G Network for Industrial Vision Detection. In Proceedings of the IEEE 6th International Conference on Big Data Analytics (ICBDA), Online - Virtual Conference, 5–8 March 2021.
- [48] 5G Alliance for Connected Industries and Automation (5G-ACIA). White Paper: Security Aspects of 5G for Industrial Networks. Available online: <https://5g-acia.org/whitepapers/security-aspects-of-5g-for-industrial-networks/> (accessed on 16 July 2021).
- [49] Tuptuk, N.; Hailes, S. Security of smart manufacturing systems. *J. Manuf. Syst.* **2018**, *47*, 93–106.
- [50] Adeogun, R.; Berardinelli, G.; Mogensen, P.; Rodriguez, I.; Razzaghpour, M. Towards 6G in-X subnetworks with sub-millisecond communication cycles and extreme reliability. *IEEE Access* **2020**, *8*, 110172–110188.

References

Paper B

Evaluation of the Impact of Wireless Communication in Production via Factory Digital Twins

R. S. Mogensen, I. Rodriguez, C. Schou,
S. Mortensen, and M. S. Sørensen.

The paper has been published in the
Manufacturing Letters, vol. 28, pp. 1-5, April 2021

© 2021 Society of Manufacturing Engineers (SME). Published by Elsevier Ltd.
The layout has been revised.

Abstract

This paper investigates the impact of wireless communication in manufacturing systems. By using a digital twin representation of a real production line, a production throughput evaluation is performed for various configurations considering different wireless control communication schemes over Wi-Fi, 4G LTE and 5G NR. The results show that operating the manufacturing execution system (MES) over wireless instead of Ethernet will minimally impact production, as the production throughput will not be degraded for lines containing slow stations; or will only degrade by a maximum of 0.41% during one month of continuous production in the case of fast lines involving quick stations.

B.1 Introduction

As factories move towards Industry 4.0, the replacement of control communications based on Ethernet and other wired technologies with wireless becomes a necessary enabler in order to fulfill the envisioned flexibility and easy reconfigurability of the production facilities [1]. Removing cables and providing wireless control of the production does, however, come with an associated cost: currently available technologies such as Wi-Fi (the most widespread wireless technology deployed in factories nowadays) or 4G LTE introduce significantly larger communication delays compared to standard Ethernet [2]. As we reported in [3], the median one-way latency measured in an operational manufacturing execution (MES) system can vary from the 0.09 milliseconds (ms) experienced over Ethernet to 1.40-2.22 ms when operating over Wi-Fi, and to up to 23.8 ms when running over a public 4G connection. While the effect of communication delays in each of the individual modules of the lines appears negligible, their overall effect in the entire production line in the long-term may affect the production outcome.

In order to provide some insight into the long-term effects of the wireless control of the overall production processes, this paper introduces a framework and a methodology for mapping the performance of different wireless technologies to production throughput. For this purpose, we evaluate the efficiency of various production line configurations when operating their MES control over different radio technologies. In particular, we compare the reference production performance achieved over standard Ethernet with those obtained over Wi-Fi, 4G LTE and the novel 5G NR.

The evaluation is performed by means of a digital twin representation of the Aalborg University (AAU) Smart Production Lab FESTO Cyber-Physical (CP) Factory production line [4, 5], which is often used for other manufacturing-related activities such as virtual commissioning [6]. Digital twins are virtual representations of the physical world which, in addition to mirroring of the

physical assets, are characterized by considering the real-time data flows, in contrast to a simulation model [7]. They are one of the main enabling technologies of Industry 4.0 [8], and are moving towards a common resource used in production design and manufacturing planning [9]. However, the communication aspects of the production, and specially those of wireless, are typically not considered in these tools. To the best of our knowledge, we are first authors in leveraging digital twins for evaluating or predicting the behaviour of production lines when controlled by wireless technologies.

B.2 Integrating Wireless in Factory Digital Twins

For our investigation we leverage the digital twin of a FESTO CP Factory production line. However, it should be noted that the framework/methodology presented in the following is applicable to any other factory digital twin setup and communication technologies. For example, the latency distributions provided in this paper could be reused for evaluation over different digital twin layouts. Also, if other latency distributions are available, the production output could be evaluated over those communication schemes instead.

The production system consists of 6 production modules interconnected by conveyor belts which transport the products across the different process stations (2 per module) during their manufacturing cycle. Within the specific considered configuration, the line is set to produce different variants of mock-up phones. An overview of the reference production line composition is displayed in Fig. B.1, and a description of the process stations, including their average processing time is provided in Table B.1.

The overall product order as well as the individual actions of each specific station for a given product are controlled by a MES. When the system is operated over wireless, the Ethernet cables between modules are removed, and dedicated wireless interfaces are applied between the MES and the different workcells. A diagram of such implementation is detailed in [3]. The MES control implies that multiple communications between the MES central controller and the stations happen throughout the overall manufacturing process. On average, 5 control messages are exchanged every time a product arrive to a specific station. For further reference, as described in [10], MES control traffic consists of mainly TCP traffic with packets of 67 B of SDU MAC size (including headers) sent with a frequency of 120 ms, on average. Thus, the system should be considered as a low-throughput and low-density deployment.

In the current digital twin implementation, this flow is converted to a discrete event simulation where both station processing time and communication delays are considered. To the best of our knowledge, communication delays are usually not considered in typical digital twin implementations as Ethernet,

B.2. Integrating Wireless in Factory Digital Twins




Fig. B.1: Overview of reference composition of the digital twin of the AAU Smart Production Lab FESTO CP Factory (3D-view), and 2D-view of the 3 different layouts used in the simulations: line with robotic cell (LRC), closed loop line (CLL), and fast loop line (FLL). The 2D-views depict the overall production cycles, including the specific order of operations and actions performed at each station.

the reference communication technology used in MES, consistently achieve latencies in the order of 0.1 ms; which can be neglected due to the dominance of the processing times (typically in the order of seconds). However, the latencies experienced in some wireless systems are significantly larger (in the order of a few tens of ms) and thus, it is essential that they are considered in the simulation in order to have a real estimate of the overall impact on the system performance when a wireless technology is chosen to operate the MES.

In our proposed framework, the communication delays, every time there is an exchange of information between a station and the MES controller, are determined by uniform sampling of the empirical MES latency traces obtained over the real FESTO CP Factory production line, which were previously presented in [3].

Table B.1: Digital twin scenarios and simulation settings.

Scenario	Station action (avg. process time)	Conveyor speed	Load capacity	Additional notes
Line with robotic cell (LRC,100%)	Lower part dispenser (5 s) Drilling station (6 s) Manual repair (10 s) Robotic cell (80 s) Camera inspection (0.66 s) Flex station (10 s) Top part dispenser (5 s) Manual packing (11 s)	0.1 m/s	100% capacity 16 pallets	Reference configuration. New product orders placed according to real-world intervals. Product type: mixture of mockup phones with different colors. Max. # of simultaneous products being produced = # pallets.
Line with robotic cell (LRC,70%)			70% capacity 16 pallets	New product orders placed at 30% reduced rate compared to LRC,100%
Closed loop line (CLL)	Lower part dispenser (5 s) Drilling station (6 s) Manual repair (10 s) Dummy station (11 s) Camera inspection (0.66 s) Flex station (10 s) Top part dispenser (5 s) Manual packing (11 s)	0.1 m/s	∞ capacity 16 pallets	Robotic cell is replaced with a dummy station. New product orders are placed instantaneously once a pallet is free.
Fast loop line (FLL)	Lower part dispenser (0.66 s) Top part dispenser (0.66 s) Camera inspection (0.66 s) Quick packing (0.66 s)	1 m/s	∞ capacity	Fast-processing production line with fast conveyor belt speed and infinite capacity.

B.3. Digital Twin Simulation Settings

Fig. B.2.a displays the one-way delay MES latency distributions computed over the real-world traces of the line under different communication schemes:

- Ethernet: default MES communication scheme.
- Ideal Wi-Fi: all modules are connected to a dedicated interference-free WiFi channel (this includes the proper isolation from any other rogue and nearby networks), exclusively used for manufacturing control purposes.
- Non-optimized Wi-Fi: all modules are connected to a standard WiFi channel, potentially interfered by other WiFi networks and not dedicated exclusively to production (i.e. regular office or personal use devices are also connected to this channel).
- Private 4G LTE: all modules are connected to a dedicated 4G channel, exclusively used for manufacturing control purposes. A private 4G network differs from the public 4G network in the sense that it is an in-factory dedicated installation that provides better latency performance and security.
- Private 5G NR¹: all modules are connected to a dedicated 5G channel, exclusively used for manufacturing control purposes. A private 5G network with ultra-reliable low-latency (URLL) support has been specifically designed for supporting industrial applications, achieving latency values close to those achieved over Ethernet, significantly lower than the ones experienced in 4G private networks [11].

B.3 Digital Twin Simulation Settings

In order to evaluate the impact of wireless MES control in the overall production process, we consider 4 different scenarios with different digital twin layouts (see Fig.B.1) and simulation settings which are detailed in Table B.1. These scenarios are used to evaluate the impact of running the MES control over different wireless technologies in terms of degradation in overall production throughput (PT), computed over the simulation of a full month of continuous production (730 hours). This impact is evaluated by comparing the PT achieved over a given wireless technology with the reference PT obtained when the MES operates over standard Ethernet. In order to do that, the absolute difference in number of manufactured products (Δ) and the normalized

¹The private 5G NR latency values have been generated via simulations, but we expect to validate the results once the first release of 5G NR URLL is available in our industrial lab.

production throughput (NPT) metric are defined:

$$\Delta [\text{products}] = \text{PT}_{\text{Wireless}} - \text{PT}_{\text{Ethernet}} \quad (\text{B.1})$$

$$\text{NPT} [\%] = \frac{\text{PT}_{\text{Wireless}}}{\text{PT}_{\text{Ethernet}}} \cdot 100 \quad (\text{B.2})$$

B.4 Production performance results & Discussion

The digital twin simulation results are presented in Table F.1 for the different test setup scenarios and MES communication technologies. The table includes the monthly PT as well as the NPT for the cases considering wireless technologies. Fig. B.2.b illustrates the NPT for the different cases, and together with Fig. B.2.a. serves to put in perspective how the choice of particular wireless technology for MES control impacts the overall performance of the manufacturing process.

The results for the reference scenario, LRC, for both capacity load configurations (100% and 70%), indicate that despite the average communication latency for the different wireless technologies are very different from the ones in the reference Ethernet case, this does not appear to affect the overall production throughput ($\Delta=0$ and $\text{NPT}=100\%$). This is due to the presence of the robotic cell on the line, which average processing time is quite large and therefore forces some products to wait in a buffer or circulate around the line until they can be processed by the station [10]. A similar effect would be observed, for example, in lines where other types of slow stations (i.e. manual stations) are present, as the bottleneck in this case is the station processing time and not the control communication delay.

The situation changes when no robotic cell is placed in the line. In this case, as the results from the CLL scenario with infinite load capacity describe, we start to observe some impact of the different communication delays over the overall production throughput. Using ideal Wi-Fi, a small decrease of 0.05% is observed with respect to the reference Ethernet case. This difference is higher for the non-optimized Wi-Fi and private 4G LTE cases (0.31% and 0.38%, respectively). For the private 5G case, it is expected that the overall production is comparable to the Ethernet reference case (0.01%), which is much better than the one predicted for the other wireless technologies.

Even in the case of faster production lines with high conveyor belt speed and quick stations with short processing time (i.e. camera inspection), as the results from the FLL scenario describe, the impact of wireless control over production remains limited (0.02-0.41%).

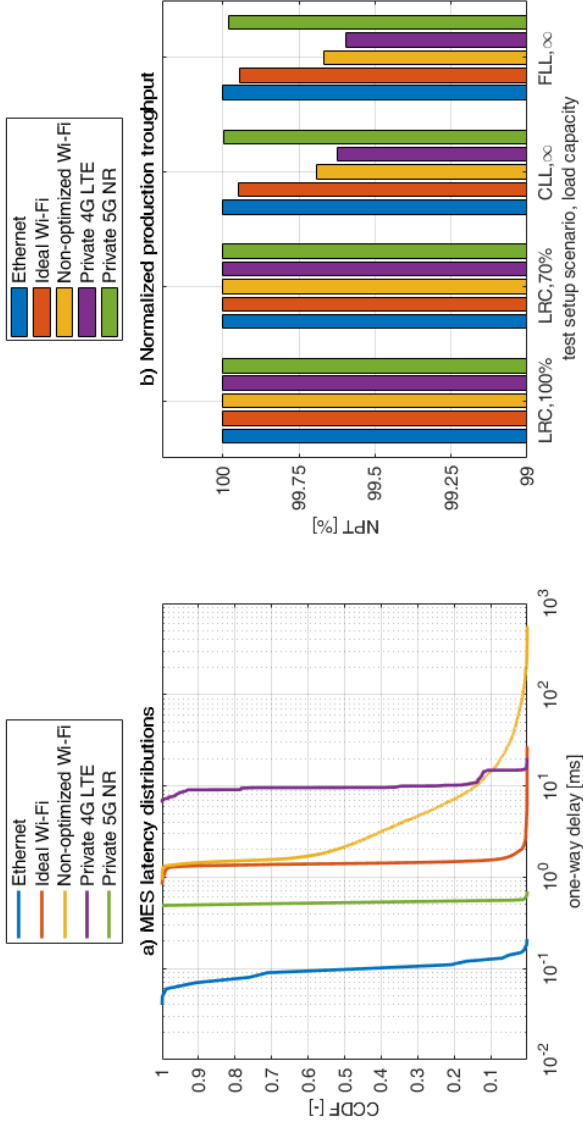


Fig. B.2: a) Complementary cumulative distribution function (CCDF) of the MES latency for the different communication technologies used in the study (Ethernet, ideal Wi-Fi, non-optimized Wi-Fi and private 4G LTE are empirical, while the private 5G NR is simulation-based). b) Normalized throughput simulation results for the different test setups and technologies.

Table B.2: one-month production throughput results for the different test setups and technologies.

Technology	Metric	Line with robotic cell (LRC)		Close loop line (CLL)		Fast loop line (FLL)	
		100% capacity	70% capacity	∞ capacity	∞ capacity	∞ capacity	∞ capacity
Ethernet	# manufactured products NTP [%]	29842 100	22834 100	206999 100	847619 100		
Ideal Wi-Fi	# manufactured products (Δ) NTP [%]	29842 (=) 100	22834 (=) 100	206894 (-105) 99.949	847153 (-466) 99.945		
Non-optimized Wi-Fi	# manufactured products (Δ) NTP [%]	29842 (=) 100	22834 (=) 100	206364 (-635) 99.693	844791 (-2828) 99.666		
Private 4G LTE	# manufactured products (Δ) NTP [%]	29842 (=) 100	22834 (=) 100	206221 (-778) 99.624	844184 (-3435) 99.595		
Private 5G NR	# manufactured products (Δ) NTP [%]	29842 (=) 100	22834 (=) 100	206995 (-4) 99.998	847464 (-155) 99.982		

Despite the fact that the results presented here relate to a specific production line and configuration, the methodology and observations given in this paper can be generalized to other factory setups with different layouts and production configurations. The simulation presented here considers only continuous production; but real production plans typically include re-configuration of the lines and planned maintenance stops, where the additional benefits of using wireless will be more tangible as service times will be highly reduced due to the fact that no control communication cables between modules will be needed. It is therefore expected that using wireless will result even in positive PT gains when considering the full spectrum of production-related activities including maintenance, reconfiguration, changeover and repairs.

B.5 Conclusions

This paper presented an evaluation of the impact of wireless communication technologies on production throughput based on digital twin simulations. The hybrid methodology introduced, which integrated digital twins discrete event simulations and empirical wireless performance, is applicable to any other factory digital twin setup and communication technologies. The study is carried out using different configurations of a digital twin of the CP FESTO Factory production line at the AAU Smart Production Lab, by considering multiple wireless technologies such as Wi-Fi, 4G LTE and 5G NR for the manufacturing execution system (MES) control communication. The production throughput simulation results computed over a full month of continuous production, indicate that the effects of operating the control of production lines over wireless are negligible in the case that slow stations such as robotic or manual cells are present. If no slow stations are present, the wireless MES control can degrade the overall production output by 0.01-0.41% with respect to the standard reference operation over Ethernet, depending on the exact operational configuration of the line and the wireless technology chosen. The results also reveal that 5G NR is expected to provide comparable production performance as the reference one achieved over Ethernet.

References

- [1] B. Chen, J. Wan, L. Shu, P. Li, M. Mukherjee, and B. Yin, "Smart Factory of Industry 4.0: Key Technologies, Application Case, and Challenges," *IEEE Access*, vol. 6, pp. 6505–6519, 2018.
- [2] B. Holfeld, D. Wieruch, T. Wirth, L. Thiele, S. A. Ashraf, J. Huschke, I. Aktas, and J. Ansari, "Wireless Communication for Factory Automation: an opportunity for LTE and 5G systems," *IEEE Communications Magazine*, vol. 54, no. 6, pp. 36–43, 2016.

References

- [3] R. Mogensen, I. Rodriguez, G. Berardinelli, A. Fink, R. Marcker, S. Markussen, T. Raunholt, T. Kolding, G. Pocovi, and S. Barbera, "Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0," in *IEEE 90th Vehicular Technology Conference, VTC-Fall*, 2019.
- [4] O. Madsen and C. Møller, "The AAU Smart Production Laboratory for Teaching and Research in Emerging Digital Manufacturing Technologies," *Procedia Manufacturing*, vol. 9, pp. 106 – 112, 2017.
- [5] Michael Sparre Sørensen, "Refining the Digital Twin on AAU Smart Factory," Aalborg University, Department of Materials and Production, Tech. Rep., Jan. 2019.
- [6] S. T. Mortensen and O. Madsen, "A Virtual Commissioning Learning Platform," *Procedia Manufacturing*, vol. 23, pp. 93 – 98, 2018.
- [7] R. Rosen, G. von Wichert, G. Lo, and K. D. Bettenhausen, "About the Importance of Autonomy and Digital Twins for the Future of Manufacturing," *IFAC-PapersOnLine*, vol. 48, no. 3, pp. 567 – 572, 2015.
- [8] H. Ahuett-Garza and T. Kurfess, "A Brief Discussion on the Trends of Habilitating Technologies for Industry 4.0 and Smart Manufacturing," *Manufacturing Letters*, vol. 15, pp. 60 – 63, 2018.
- [9] F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang, and F. Sui, "Digital Twin-driven Product Design, Manufacturing and Service with Big Data," *The International Journal of Advanced Manufacturing Technology*, vol. 94, p. 3563–3576, 2018.
- [10] I. Rodriguez, R. Mogensen, E. Khatib, G. Berardinelli, P. Mogensen, O. Madsen, and C. Møller, "On the Design of a Wireless MES Solution for the Factories of the Future," in *2019 Global IoT Summit (GIoTS)*, 2019.
- [11] J. Sachs, G. Wikstrom, T. Dudda, R. Baldemair, and K. Kittichokechai, "5G Radio Network Design for Ultra-Reliable Low-Latency Communication," *IEEE Network*, vol. 32, no. 2, pp. 24–31, 2018.

Paper C

Empirical IIoT Data Traffic Analysis and Comparison to 3GPP 5G Models

R. S. Mogensen, I. Rodriguez, G. Berardinelli,
G. Pocovi, and T. Kolding.

The paper has been published in the
2021 IEEE 94th Vehicular Technology Conference (VTC2021-Fall)

© 2021 IEEE

The layout has been revised.

Abstract

5G New Radio (NR) technology is expected to enable the wireless transition of manufacturing processes in modern factories. In order to properly dimension the 5G system, accurate models of the data and control traffic in industrial use cases are needed. In this paper, we analyze the industrial traffic empirically measured in different Danish factories and compare its characteristics to the modeling assumptions used in the 3GPP for 5G NR performance analyses. In particular, three relevant use cases are studied: a unit test cell for actuator calibration, a visual inspection cell, and the control of an autonomous mobile robot (AMR). Our results indicate that some of the relevant 3GPP assumptions for industrial traffic such as exponential packet-interarrival time (PIAT) for aperiodic traffic and fixed packet sizes are only partially corroborated by experimental evidence. Moreover, industrial traffic can exhibit a large burstiness that may lead to conservative admission control and radio resource allocation.

C.1 Introduction

The fourth industrial revolution (I4.0) is the next step in manufacturing. Numerous emerging technologies ranging from big data and artificial intelligence to matrix-based production and autonomous mobile robots, are expected to bring value creation through new services or increased production output [1]. A prerequisite for the agile vision of I4.0 is wireless communication, allowing for reduced deployment costs as well as more flexible and scalable production environments. Dimensioning a wireless system to operate specific industrial applications requires insight into specific communication requirements which in turn must be analyzed, modelled, and validated. 5G New Radio (NR) is considered a strong candidate for providing wireless industrial communication, as it is designed to support a broad set of requirements, including millisecond-level latencies with five nines reliability. The 3rd Generation Partnership Project (3GPP) standardization body and advisory bodies such as 5G-Alliance for Connected Industries and Automation (5G-ACIA) have worked together with industrial verticals to define the proper application scenarios and traffic models that are used in the design and evaluation of 5G Industrial Internet-of-Things (IIoT) systems [2] [3]. However, as highlighted in [4], there is still work to be done in terms of definition of industrial use cases, their associated communication requirements and traffic models.

While the existing models can provide a valid starting point for the analysis of relevant wireless use cases, they may indeed fail in reproducing critical aspects of real-world deployments and scenarios. This may result in incorrect dimensioning of the wireless system accomplishing specific tasks, leading to over-provisioning of the radio resources or in the worst case in under-

performing solutions, which could be showstoppers to the support of critical applications. Experimental verification of the characteristics of industrial traffic in relation to established models is therefore needed. In this respect, there is little scientific literature reporting traffic measurements of industrial use cases under operational conditions. In [5], the authors present models for several different Ethernet-based use cases based on measurements of a testbed for flexible manufacturing processes. While the authors report the final parametrization of the models, there is a lack of insight on the actual measurements and underlying industrial communication. A method for analyzing industrial data traffic is presented in [6]. However, the analysis is limited to supervisory control and data acquisition (SCADA) traffic and Skype calls and lacks discussion on the suitability of existing radio technologies and their respective modelling approaches. Related to data traffic characterization, but in a different non-industrial context, [7] presents a thorough analysis and classification of data traffic patterns in public networks.

In general, the existing literature presents simplified findings from measurements performed over non-industrial networks or semi-artificial system setups. This paper aims at shedding some light into the topic, by providing an extensive empirical analysis of control data traffic flows obtained from real industrial factory machinery in operational conditions. Further, this paper verifies to what extent existing 3GPP modelling assumptions reflect real world scenarios, and identify further relevant aspects in the traffic behavior, not accounted by current models, that are to be considered when dimensioning the wireless system. We base this upon empirical analysis of traffic measurements for three relevant industrial use cases performed in different operational Danish factories. Specifically, we study the cases of an unit test cell for quality insurance, a visual inspection cell, and manual or fleet operated control of an autonomous mobile robot (AMR). Through the analysis we compare traffic distribution, packet sizes, and burstiness characteristics against 3GPP models and assumptions used as part of 5G NR IIoT studies. It is worth mentioning that it is not in the scope of the paper to generalize the results of our analysis, neither to derive new traffic models, as this would require more thorough and extensive studies. We believe however this paper can represent a starting point for traffic characterization in industrial scenarios, which may lead in the future to tailored traffic models to be used in performance evaluation of novel IIoT concepts.

The rest of the paper is structured as follows. Section C.2 introduces the main modelling assumptions and metrics for traffic characterization in industrial scenarios as defined by the 3GPP. Section C.3 describes the measurement approach and the studied use cases. Section C.4 presents the traffic statistics extracted from the traces along with a discussion on how these compare with the 3GPP models and assumptions. Finally, conclusions are summarized in Section C.5.

C.2 Modeling of IIoT traffic in 5G

As baseline for standardization work for 5G in IIoT, many studies were done to define use-cases and typical application traffic behaviors and requirements for the factories of the future [2]. Reference use-cases include for instance motion control, control-to-control communication, mobile robots, human machine interfaces, process automation, and monitoring. The applications vary significantly in terms of their requirements for service availability, packet reliability, packet latency, and throughput levels. In the following, the traffic models used most often in 5G IIoT studies are presented as well as the parameters that are used to describe traffic characteristics and their requirements within the 5G system, i.e. the 5G Quality of Service (QoS) model.

C.2.1 IIoT traffic models

In spite of large differences between IIoT applications, only few simple models are commonly used when evaluating 5G performance for such applications as illustrated in Table C.1, i.e. periodic and aperiodic traffic models. The periodic model uses a fixed packet size and fixed packet inter-arrival time (PIAT) (i.e., fixed periodicity) and the relative phase or time offset between different flows may be modeled [8]. In the case of aperiodic traffic, the standard model assumption used in 3GPP is the FTP model 3 [9]. This assumes fixed packet size with Poisson arrival, resulting in exponentially-distributed packet inter-arrival times as follows:

$$\text{PIAT} \sim \text{Exp}(\lambda), \quad (\text{C.1})$$

where λ is the average packet rate. Using this assumption, the traffic is memory-less; meaning that there is no correlation between PIAT values for consecutive packets.

Table C.1: Typical traffic models used in 5G IIoT standards research

Characteristic	Periodic model	Aperiodic model
Packet size	Fixed	Fixed
Packet spacing	Fixed	Statistical
Transmission	Unicast	Unicast

C.2.2 5G QoS model

For characterizing the aperiodic traffic, the 5G QoS framework defines further parameters and metrics [10]. In 5G, a QoS flow is defined as a stream of packet(s) between two endpoints which pertains to the same application

[10, 11]. A pair of endpoints can have multiple QoS flows in parallel, but each flow uses its own distinct traffic model. In case a packet stream is multicast or broadcast oriented, the 5G system acts as a proxy and translate this single stream into multiple distinct unicast QoS flows, which means that from a modeling perspective there is no difference.

Depending on the traffic flow resource type, only selected QoS parameters are used to describe the packet stream. For traffic flows with strict delay control requirements (such as the ones observed in factories), the Delay-critical Guaranteed Bit Rate (GBR) QoS Flow type is expected to be used. For this resource type, the following metrics are of importance when configuring the 5G system: Priority Level, Packet Delay Budget (PDB), Packet Error Rate, Averaging Window (AW), Maximum Data burst volume (MDBV). The MDBV is defined as the maximum data burst that needs to be delivered within any given PDB. A mathematical model for MDBV can be formulated as:

$$\text{MDBV}(\text{PDB}) = \arg \max_{t_0 \in \mathbb{R} \geq 0} \int_{t=t_0}^{t_0+\text{PDB}} b(t) \cdot dt, \quad (\text{C.2})$$

where $b(t)$ describes the incoming data at time t (Bytes).

For a QoS flow, another metric, the Guaranteed Flow Bit Rate (GFBR) denotes the required maximum data rate of the flow averaged over the AW, where AW has an average value of 2 seconds [10]. For IIoT type traffic, the allowed failure rates are very low, often 0.0001% or lower. A failure for the Delay-critical GBR QoS resource type is defined when all of the following conditions are fulfilled [10]:

- A packet is delayed by a value greater than PDB,
- the data burst over the period of PDB has not exceeded the MDBV, and
- the QoS Flow is not exceeding the GFBR in the current sliding window.

Specifically for periodic traffic, 5G introduces Time Sensitive Communication Assistance Information (TSCAI), enabling each gNB to know which periodicity (Periodicity) and at which time-offset (Burst Arrival Time) traffic arrives [10]. As shown in [8], this information is helpful to configure resources for UE in advance to reduce the latency, especially in uplink, but also in configuring semi-static resources that allows for significant optimization of the 5G spectral efficiency.

For the sake of this study, we introduce the Average Data Burst Volume (ADBv) which is similar to the MDBV above except it is defined as the average volume of all windows within a measurement period where data is present. Comparing the ADBv to the MDVB as well as the GFBR provides some additional insights into the burstiness of the observed data flow. Further,

we have generalized the MDBV as we do not know the effective PDB for the studied applications and thus the specific window for which the MBDV should be evaluated over. Instead, we report MDBV statistics for a selected range of windows [1, 2, 5, 10, 100] ms.

C.3 Industrial Use Cases

Three uses cases are considered for our empirical traffic analysis. Such use cases have been selected according to their IIoT relevance but also due to the granted permissions for accessing the factory premises - it is not often the case in which factories are willing to share their use cases, control systems and grant permission to interface their machinery for obtaining an overview of the control data traffic as, apart from control data, there are typically other production-confidential information in those exchanges of information.

Data traffic measurements were obtained by interfacing a "sniffer" to selected machinery Ethernet-based network links where the control data is flowing. The sniffer device used was an evolution of the one previously presented in [12] and it logs all protocol headers at layers 2 and above while remaining transparent to the underline application. No packet inspection is performed on the packets, conserving the confidentiality of the production business-critical data. The sniffer introduces a calibrated extra latency not larger than 0.5 ms due to the input-output bridging and the data logging process. The sniffer performs reliable measurements (no packet losses) for packets of up to 1500 B size and packet inter-arrival times larger than 10 μ s. The introduced latency is generally constant (with a processing delay jitter <30 μ s) and therefore, inter-arrival packet measurements are not affected by the measurement procedure.

C.3.1 Unit Test Cell (UTC)

We measure the traffic in a quality assurance unit test cell, whose main task is to provide calibration and tolerance for an actuator that the company uses in their final product. The network architecture related to this use case is illustrated in Fig. C.1.a, where the position of the sniffers (S) is highlighted with green circles. For further reference, a picture of the machinery control networking setup is shown in Fig. C.2.a. A soft Programmable Logic Controller (PLC) acts as supervisor of the cell operations. An actuator is placed in a test rig, where it is activated by the soft PLC according to a specified control sequence. A set of sensors continuously measures the device performance and report such measurements to the soft PLC. Once the test sequence is finished, the collected results are reported by the soft PLC to a factory level common database.

The first sniffer is placed in between the soft PLC and the sensor aggregator (SA), thus capturing the sensing part of the control cycle. The second sniffer is placed in between the cell and the factory network capturing all the traffic going to and from the cell from a factory network perspective.

C.3.2 Visual Inspection Cell (VIC)

This use case consists of a cell which is also performing quality assurance, but this time, by means of video feeds from a camera located in proximity of a conveyor belt transporting parts for final assembly. The network architecture of this use case is illustrated in Fig. C.1.b. The live feed from the camera is streamed to a image processing unit (IPU), that performs quality control tasks. Here, the images are compared to the ones stored in a collection of successfully assembled product, in order to detect possible relevant discrepancies. Based on the outcome, a command is forwarded to the line controller/soft PLC, which is the centralized reference or main brain of the cell, that then dictates the operation of the robots in the assembly line by sending proper commands to the actuator controllers. For example, in case the product part is deemed to be faulty, the robots remove it from the conveyor belt. Conversely, in case the quality of the product part fulfills the specification, the robots forward it to the next manufacturing cell. Throughout the process, status updates are sent from the soft PLC to a human machine interface (HMI) device, for general monitoring purposes.

In this case, we measure two traffic flows by probing with the sniffer the link between the line controller and the cell main switch: one between the soft PLC and the IPU, and another between the soft PLC and the HMI.

C.3.3 Autonomous Mobile Robot (AMR)

The third use-case considers AMRs, which are robot vehicles, typically used for logistics within production, that can move materials across the factory space without using pre-marked routes, but rather relying on Light Detection And Ranging (LiDAR) or optical technologies to self-navigate between target positions. The network architecture and control traffic measurement setup for this use case are illustrated in Figs. C.1.c and C.2.b, respectively.

The considered AMR use case considers has two different operational configurations:

1. Fleet manager (FM) mode, where a central entity (manager) allocates robots to perform certain automated tasks in the factory.
2. Manual controller (MC) mode, where a human is manually-controlling and guiding the robot via tablet or phone.

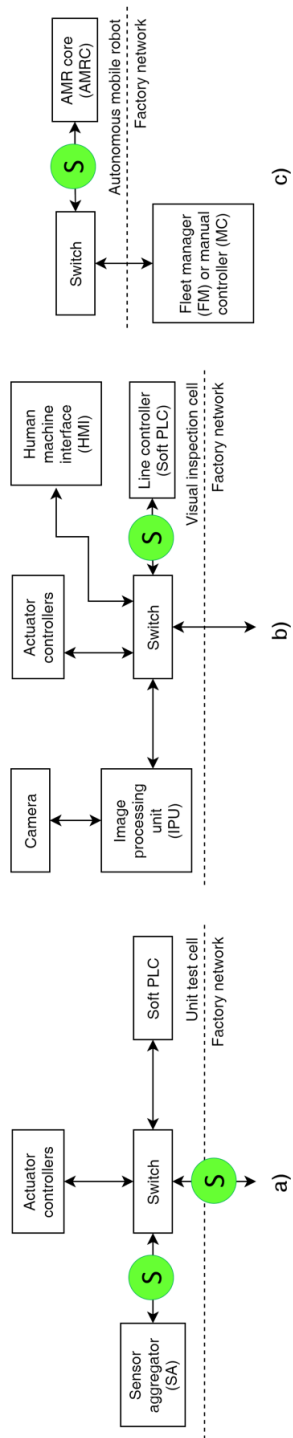


Fig. C.1: Overview of the network architectures and traffic sniffing measurement points (S) for the selected use cases: a) unit test cell, b) visual inspection cell, and c) autonomous mobile robot.




Fig. C.2: Overview of industrial measurement setups for two of the use cases: a) unit test cell, and b) autonomous mobile robot. Unfortunately, we are not allowed to disclose any picture related to measurement of the visual inspection cell use case.

In fleet mode, the manager is requesting telemetry and status information at regular intervals from each robot and receives updates from the robots as they reach objectives or complete assigned tasks. In manual mode, the robots transmit to the operator a map of the local environment together with its position, including live updates of any detected obstacles. From the operator side, communication consists of manual control input from a supervision tablet, besides task commands from a predefined set of actions.

The AMR represents a clear example of a wireless use case. In live operation, the connections between the AMR operation core (AMRC) and the FM and MC would happen via factory floor Wi-Fi. Therefore, to capture the application control data traffic without the intrinsic effects of a wireless connection, a sniffer was placed in between the AMR core and the PC hosting both the manual control and the fleet manager software, interconnected via gigabit Ethernet.

C.4 Results & Discussion

In this section, the traffic analysis of the different use cases is presented. Measurements are obtained via our traffic sniffers as described in Section C.3. Performance is discussed based on the parameters presented in Section C.2 on a per use case basis and based on the traffic flow to (\leftarrow) and from (\rightarrow) a specific target. Relevant bi-directional (\leftrightarrow) data flows between two devices are also addressed in some cases. Emphasis is given to those cases exhibiting very tight inter-arrival times - in the order of 1-10 ms and below, which can represent cases of critical IIoT traffic if associated to a tight PDB.

Results for the unit test cell, the visual inspection cell and the autonomous mobile robot are displayed in Figs. C.3, C.4 and C.5, respectively. In each of these figures, sub-figure a) refers to the PIAT statistics, while figures b) and c) refer to the packet size statistics, and the MBDV and ADBV behavior, respectively. As we do not know the specific effective PDB for most of the cases and, therefore, the specific windows over which the MBDV should be evaluated over remains also unknown; we report statistics for a selected range of windows [1, 2, 5, 10, 100] ms. This is meant to capture the network requirements in terms of traffic to be supported for a diverse set of PDBs.

C.4.1 Individual Use Case Results

Unit test cell

When examining the results for the unit test cell as presented in Section C.3, there is a clear distinction between the cell \leftrightarrow factory and the PLC \leftrightarrow sensor aggregator (SA) flows - see Fig. C.1.a. The staircase-like behavior of the cell \leftrightarrow factory traffic PIAT in Fig. C.3.a, stems from the presence of periodic inter-cell broadcast messages which make up the majority of the packets observed within the traffic traces ($\sim 90\%$). Similar tendencies are also observed for the packet sizes in Fig. C.3.b as there are three distinct values (60, 234, and 637 bytes) that occur each with a different probability. When examining the MDBV as compared to the ADBV in Fig. C.3.c, a different of an order of magnitude is observed, suggesting that dimensioning the system based on the MDBV would result in a potential over-dimension of resource allocation. Regarding the PLC \leftrightarrow SA flows, the periodic nature of the traffic is clear. A 2 ms cycle is observed in Fig. C.3.a, although with minor deviation at the lower percentiles which is due to the initial setup and control messages. The flows have fixed packet sizes (96 and 113 bytes) that are directional-dependent. The ADBV and the MDBV curves in Fig. C.3.c are, in this case, relatively close to each other compared to the cell \leftrightarrow factory case, due to the traffic being nearly periodic and with limited variability in the packet sizes.

Visual inspection cell

As the traffic in this cell appears clearly non-periodic, together with the PIAT statistics in Fig. C.4.a, their associated 3GPP model-based exponential (exp) inter-arrival time distributions are also displayed - the model was fit based on the mean PIAT calculated upon the empirical data. Such exponential distributions exhibit a good match with the IPU \leftrightarrow PLC traffic for inter-arrival times in the order of 1 ms and above, while representing a pessimistic estimate of the effective traffic for more critical timings. For the HMI \leftrightarrow PLC case, the exponential distribution appears even more pessimistic, with the exclusion of the PIAT values at the top end. The HMI \leftrightarrow PLC packet sizes displayed in Fig. C.4.b, range from 128 bytes up to 512 bytes in downlink and 100 bytes to 250 bytes in uplink. For the IPU \leftrightarrow PLC traffic, the packet sizes ranges from 100 bytes to 200 bytes for both uplink and downlink. The ADBV and MDBV analysis in Fig. C.4.c, shows higher data volume for HMI \leftrightarrow PLC traffic. The difference between ADBV and MDBV reflects a certain traffic burstiness for both HMI \leftrightarrow PLC and IPU \leftrightarrow PLC flows, but not as dramatic as in the previous UTC case.

Autonomous Mobile Robot

The fleet manager is generating periodic polling commands (FM \leftrightarrow AMRC flow); however, the PIAT values shown in Fig. C.5.a are not periodic due to variability of response time of each AMR, which affects both input and output communication flows. This is due to non deterministic processing time at each robot, as well as to the implementation of the control application (based on REST interfaces). MC \leftrightarrow AMRC traffic shows critical components with tight inter-arrival times; such inter-arrival times are rather heterogeneous due to the specific underline application process (e.g., image transmission) and behavior of the REST APIs. Packets to and from the fleet are either very small (around 60 bytes) or in the order of 1 kilobyte and above (Fig. C.5.b). For the MC \leftrightarrow AMRC mode, such big packets can be due to the images that are streamed to the operator, as well as to the potential inefficiency of underline protocol design. For this AMR use case, the data volume is fairly insensitive to the window size (Fig. C.5.c). This is due to the nature of the traffic, composed mainly of periodic components. Some dependency is only visible for the manual case, due to the inherent variability of the data traffic (image transfer). Also, no major differences between MDBV and ADBV are visible.




Fig. C.3: Measurement results from the unit test cell flows: a) statistical distribution of the PIAT, b) statistical distribution of the packet size, and c) ADBV and MDBV estimation for different window sizes.




Fig. C.4: Measurement results from the visual inspection cell flows: a) statistical distribution of the PIAT, b) statistical distribution of the packet size, and c) ADBv and MDBv estimation for different window sizes.




Fig. C.5: Measurement results from the AMR control flows: a) statistical distribution of the PIAT, b) statistical distribution of the packet size, and c) ADBV and MDBV estimation for different window sizes.

C.4.2 Summary of traffic statistics

Table C.2 summarizes the main statistics extracted from the figures above, together with some extra reference information related to flow throughputs and other metadata such as protocol types and broadcast/unicast ratios. In particular, the classification of traffic nature and layer 3 and above characteristics are highlighted.

Regarding the unit test cell, as expected the sensor aggregator traffic is mainly dominated by unicast packets, and traffic type is industrial, mainly consisting of the PROFINET protocol. Conversely, the ingress/egress cell traffic is dominated by broadcast packets, and mainly composed by UDP packets due to keep-alive heartbeat signals. Regarding the visual inspection cell case, the IPU traffic is mainly unicast, and composed in a major part by industrial specific protocols, i.e. Ethernet Industrial Protocol (ENIP), besides standalone TCP packets. The HMI traffic is almost entirely unicast and also composed by ENIP packets. AMR traffic is only unicast and entirely TCP.

C.4.3 Comparison with current 3GPP 5G IIoT models

In some use cases (unit test cell, autonomous mobile robot), the traffic appears to be dominated by periodic components. A composite model of the traffic flow is therefore needed to capture such dynamics. This could be obtained by combining multiple 3GPP periodic models with different periodicity. Generally, when measuring between two end-points, the overall traffic is a combination of multiplexed data flows, as for e.g., the "Cell→" link in the UTC use case. To link those measurements to the 3GPP models, which work on a per-flow basis, it is important that those flows are first demultiplexed and only then mapped to 5G QoS flows, each with more balanced parameters and proper setting of the PDB.

For aperiodic traffic as in the visual inspection cell - see Fig. C.4.a, the results indicate that the used FTP Model 3 for 5G IIoT studies in the 3GPP has only a partial match with the traffic flows of HMI and IPU. For very critical PIATs, the exponential assumption has a different curve slope compared to the real inter-arrival data extracted from the measurements. In some cases, the fitted FTP Model 3 leads indeed to a significant over-estimation of the inter-arrival times of the traffic. The model therefore does not reflect accurately the dimensioning requirements for the network and is hard to tune to the observed traffic flows. Further, variable packet size instead of fixed packet size should be considered as part of the traffic modeling as this was the case for most of studied data flows.

Table C.2: Summary of the main data traffic statistics, including flow meta data, for the different use cases and control flows.

Use case, control data flow	Inter-arrival time [ms]			Packet size [B]			Protocol statistics [%]		Throughput [kB/s]
	min	max	median	min	max	median	BC/UC	UDP/TCP/Other/Industrial	
UTC, SA←	0.007	2000	2	60	1506	96	4.5/95.5	3.3/1/<1/95.6 (Profinet)	49/48
UTC, SA→	0.002	300	2	42	1506	113			50/56.6
UTC, Cell←	0.002	1055	26	60	950	234	90.3/9.7	66.8/<1/31/1.7 (Profinet)	5.7/26.2
UTC, Cell→	0.02	382	2	42	1496	174			1.5/7.9
VIC, IPU←PLC	2.86	508	47	82	596	140			4.1/29.5
VIC, IPU→PLC	1.58	508	46	82	262	204	8.3/90.7	4.1/0/<1/95.5(ENIP)	4.2/8.2
VIC, HMI←PLC	0.009	382	9	82	592	554			42/57.7
VIC, HMI→PLC	0.128	443	10	82	582	108	0.8/99.2	0.9/0/<1/98.2(ENIP)	9.4/13.5
AMR, AMRC←MC	0.021	3584	193	42	901	678	1.7/98.3	<1/98.3 (REST)/0/0	10.3/19.9
AMR, AMRC→MC	0.007	251	44	155	1422	1047			3.8/8.8
AMR, AMRC←FM	0.021	15817	1003	42	883	104			4.3/10.7
AMR, AMRC→FM	0.030	1094	996	1396	4292	4292	0/100	0/100 (REST)/0/0	0.8/1.3

Examining the GFBR for the various use-cases and flow directions, presented, we observe that the GFBR is generally similar to the average throughput measured over the entire measurement period, with the exception of a few use cases (e.g. "IPU←" link in the VIC use case) where GFBR can be up to 8 times higher. The deviations suggest that there is some burstiness in the data traffic as also indicated by the MDBV and ABDV analysis. To fully explain the burstiness, deeper insight into the application and protocol behavior is required which is left for future work.

The window size for MDBV calculations in 3GPP typically corresponds to the PDB. When considering for instance a 10 ms window size, the MDBV in Fig. C.1.c for the "Cell→" UTC link is just above 1000 bytes meaning that data arriving within a single window have a corresponding instantaneous data rate of about 100 kB/s an ABDV rate of 29 kB/s, whereas the data rate averaged over the entire measurement period is only 1.5 kB/s. Such a large discrepancy between the MDBV, the ABDV and the average bit rate is challenging in the modeling as well as for the dimensioning of the wireless system and may lead to conservative admission control and resource allocation mechanisms. Better characterization could be done by including an additional parameter for the delay-critical GFBR resource type which describes the discrepancy between MDBV and ABDV defined over the PDB.

C.5 Conclusion

In this paper, we have studied the control data traffic characteristics of three industrial use cases (unit test cell for quality assurance, visual quality inspection of raw products, and autonomous mobile robot) based on empirical measurements obtained in Danish factories. Traffic has been statistically analyzed in terms of inter-arrival packet time, packet size, and burstiness along with the used protocols, and compared with 3GPP assumptions for 5G IIoT system design.

Our analysis revealed that the traffic in industrial scenarios is significantly more heterogeneous than what has been considered as baseline in the 3GPP. Exponential distribution for packet inter-arrival times has in the best case a limited match with the actual experienced traffic, and it may lead to an underestimation of the required resources for successful wireless communication. Moreover, packet sizes are rather diverse. Furthermore, analysis of the burstiness show large variability between average and maximum data burst sizes which may lead to e.g., conservative admission control and resource allocation. Protocol analysis showed that broadcast packets represent a major fraction of the overall traffic in the unit test cell use case. As in 5G NR broadcast traffic is divided into unicast flows, proper care should be taken when handling wirelessly this type of traffic.

This paper is intended as a starting point for industrial traffic characterization activities. In this respect, it was not in the scope of the paper to propose an updated model for 5G IIoT traffic, but rather to identify the need for a more careful analysis of data traffic in real use cases, moving beyond the generic 3GPP models. A further set of measurement campaigns in a comprehensive set of scenarios is needed for the sake of designing accurate traffic models for 5G IIoT.

References

- [1] I. Rodriguez, R. S. Mogensen, A. Schjørring, M. Razzaghpour, R. Maldonado, G. Berardinelli, R. Adeogun, P. Christensen, P. Mogensen, O. Madsen, C. Møller, G. Pocovi, T. Kolding, C. Rosa, B. Jørgensen, and S. Barbera, "5G Swarm Production: Advanced industrial manufacturing concepts enabled by wireless automation," *IEEE Communications Magazine*, vol. 59, no. 1, pp. 48–54, 2021.
- [2] 3GPP, "3rd Generation Partnership Project; Technical Specification technical specification group services and system aspects; service requirements for cyber-physical control applications in vertical domains; Release 17," 3rd Generation Partnership Program, Tech. Rep. 3GPP TS 22.104 V17.4.0, Sep. 2020.
- [3] 5G-ACIA, "A 5G Traffic Model for Industrial Use Cases," 5G Alliance for Connected Industries and Automation, Tech. Rep., Nov. 2019.
- [4] J. Navarro-Ortiz, P. Romero-Diaz, S. Sendra, P. Ameigeiras, J. J. Ramos-Munoz, and J. M. Lopez-Soler, "A survey on 5g usage scenarios and traffic models," *IEEE Communications Surveys Tutorials*, vol. 22, no. 2, pp. 905–929, 2020.
- [5] M. Glabowski, S. Hanczewski, M. Stasiak, M. Weissenberg, P. Zwierzykowski, and V. Bai, "Traffic modeling in industrial ethernet networks," *International Journal of Electronics and Telecommunications*, vol. 66, no. 1, pp. 145–153, 2020.
- [6] G. Soós, D. Ficzer, and P. Varga, "Investigating the network traffic of Industry 4.0 applications—methodology and initial results," in *2020 16th International Conference on Network and Service Management (CNSM)*. IEEE, 2020, pp. 1–6.
- [7] G. Soos, D. Ficzer, and P. Varga, "Towards traffic identification and modeling for 5g application use-cases," *Electronics*, vol. 9, no. 4, 2020.
- [8] R. B. Abreu, G. Pocovi, T. H. Jacobsen, M. Centenaro, K. I. Pedersen, and T. E. Kolding, "Scheduling enhancements and performance evaluation of downlink 5g time-sensitive communications," *IEEE Access*, vol. 8, pp. 128 106–128 115, 2020.
- [9] 3GPP, "3rd Generation Partnership Project; Technical Specification group radio access network; study on licensed-assisted access to unlicensed spectrum; Release 13," 3rd Generation Partnership Program, Tech. Rep. 3GPP TS 36.889 V13.0.0, Jun. 2015.
- [10] 3GPP, "3rd Generation Partnership Project; Technical Specification technical specification group services and system aspects; system architecture for the 5g system (5gs); stage 2 Release 16," 3rd Generation Partnership Program, Tech. Rep. 3GPP TS 23.501 V16.7.0, Dec. 2020.

References

- [11] D. Chandramouli, R. Liebhart, and J. Pirskanen, *5G for the Connected World*. Wiley, 2019.
- [12] I. Rodriguez, R. S. Mogensen, E. J. Khatib, G. Berardinelli, P. Mogensen, O. Madsen, and C. Møller, "On the Design of a Wireless MES Solution for the Factories of the Future," in *2019 Global IoT Summit (GloTS)*, 2019.

Paper D

Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0

R. S. Mogensen, I. Rodriguez, G. Berardinelli, R. Marcker,
S. Markussen, T. Raunholt, T. Kolding, G. Pocovi, and S. Barbera.

The paper has been published in the
2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)

© 2019 IEEE

The layout has been revised.

Abstract

This paper presents the implementation framework and performance evaluation of a wireless Manufacturing Execution System (MES) for Industry 4.0. The proposed solution is based on self-configuring multi-access gateways that enable seamless transport of delay-tolerant industrial Ethernet control data traffic over LTE or Wi-Fi (or both, using hybrid-access techniques). The wireless MES solution has been deployed at the Smart Production Lab facilities at Aalborg University, allowing the removal of Ethernet cables between modules in a production line setup, and thus enabling a faster re-configuration of the production facilities. The performance of the wireless MES solution is benchmarked against the reference Ethernet case in terms of latency and packet loss. The trials revealed that, despite the increase in latency and packet loss as compared to the reference operation over Ethernet, both LTE and Wi-Fi under different conditions were able to reliably support the production control operations at MES level without discernable difference on the overall functionality of the system.

D.1 Introduction

The usage of wireless communication in industrial environments and applications has received significant interest in recent years. Replacing cables by wireless offers the possibility of reduction of cost and operational expenditures as well as the promise of supporting new use cases which previously were infeasible and is seen as a key technological enabler for the 4th industrial revolution (I4.0) [1]. In general, as illustrated in Fig. D.1, the different levels of automation within industrial scenarios present different communication requirements. Typically, the bottom layers include the communication between physical sensors and actuators and their associated programmable logic controllers (PLC). Above these sits the supervisory control and data acquisition (SCADA) layer which acts as interface to the higher-level management systems of the factory. This includes the manufacturing execution system (MES) which handles and monitors the overall manufacturing/production process (i.e. the progress of a specific product from raw independent components to a finished product). Finally, the top of the pyramid consists of the enterprise resource planning (ERP), which integrates the overall business management [2].

As depicted in Fig. D.1, from a communication point of view, the required latency often increases as one goes up the pyramid; going from real-time (RT) communication at sensor and actuator level to delay-tolerant (DT) communication at higher levels. Currently, the majority of industrial control systems use wired schemes based on fieldbus or Ethernet technologies, where the latter is the current leading solution in the market [3]. Ethernet offers comparable reliability, latency but greater flexibility as compared to fieldbus technologies.




Fig. D.1: Automation pyramid and associated typical communication latency requirements for each level – adapted from [2].

As part of this flexibility, Ethernet allows for interoperation between multiple systems and platforms, which enables easy integration across the different layers in the automation pyramid. Therefore, any industrial wireless solution should be interoperable with Ethernet to ensure compatibility when deployed in an existing setup.

While industrial wireless communication solutions do exist (ISA100.11, WirelessHART, WISAN, etc.) [4], they are mainly used in the context of wireless sensor networks or human machine interfaces, e.g. monitoring and analytics of the production performance but not for controlling the industrial equipment itself. Those which are more suited for the control are often based on the IEEE 802.11 radio interface such as for example Wi-Fi or IWLAN [5]. Common for all these existing wireless solutions is that they operate in unlicensed spectrum and are, therefore, prone to interference and subject to Listen-Before-Talk (LBT) mechanisms and constraints. This means that they do not scale well in environments with interference from other sources or a large number of devices [6]. While solutions have been proposed to mitigate the scalability issue of unlicensed spectrum-based technologies, they are not commercially available yet. These solutions may include for example RT-Wi-Fi [7] and [6], which proposes a modified MAC layer to better accommodate the more demanding real-time exchange of information, or MulteFire [8], which aims at providing LTE-like functionality in the unlicensed spectrum. As they are not subject to the same constraints, e.g. with regards to interference, commercial technologies operating in licensed spectrum such as 4G LTE and 5G NR can be considered as potential solutions. In particular, 5G NR will introduce features and enhancements applicable to I4.0 use cases as specified in the 3GPP Release 16 [9]. Concepts such as private network deployments and network slicing can also help to ensure the levels of reliability needed in industrial scenarios.

This paper focuses on how Wi-Fi and public LTE networks can facilitate the delay-tolerant communication of industrial systems at the MES level of the automation pyramid. To this extent, we build on the framework presented

in [10], where we introduced our incipient vision for a wireless MES solution for industrial manufacturing, and present the implementation details of a novel multi-access gateway (MAGW) which aims at providing seamless and technology-agnostic wireless control communication, including 802.1CB like functionality that enables frame replication and elimination across disjoint communication paths. Further, this study evaluates how the solution compares to the existing wired Ethernet system in a realistic I4.0 scenario.

The rest of the paper is structured as follows. Section D.2 presents an overview of the proposed wireless MES design, with focus on the integration with existing industrial Ethernet systems. Section D.3 describes the implementation aspects of the multiaccess gateway. Section D.4 details different aspects related to the test scenario, setup and methodology used to evaluate the performance of the system over LTE and Wi-Fi as compared to the baseline Ethernet. Section D.5 presents the results of the real-world trials in terms of packet error rate and latency. Finally, conclusions are resumed in Section D.6 along with future directions.

D.2 Integration of Wireless Industrial Ethernet

As detailed in [10], we consider the Aalborg University (AAU) Smart Production Lab assembly line as our reference scenario in this work. This is a small-scale representation, based on commercial industrial-grade equipment and interfaces, of the operational setups in real factories. As displayed in Fig. D.2, the line physically consists of 5 FESTO production modules, each containing two independent stations. Each station consists of a PLC with different input/output (I/O) capabilities (analog, digital, fieldbus etc.), a display screen and an extension port (IN), all of them interconnected via a switch. They are equipped with different industrial-grade components which are in charge of performing different actions or tasks within the manufacturing process. Such components are connected to the integrated PLC either using Ethernet or the I/O. All these modules are daisy chained by Ethernet via the internal switches, to which the MES PC is also connected to orchestrate the overall operation. When operational, the system system can be configured to manufacture different variants of “mock smartphones”, each consisting of three main components: a bottom cover, a circuit board with fuse holders and a top cover. During the assembly process, these phones are transported on trays via conveyor belts from station to station where different operations are performed. The MES keeps track of the specific products in progress based on RFID tags placed on the product trays. When a tray enters a station, the RFID is read, and then the PLC reports the associated product ID to the MES. Then, the MES replies by issuing a command according to the action to be performed over the product at the particular station, according to its current

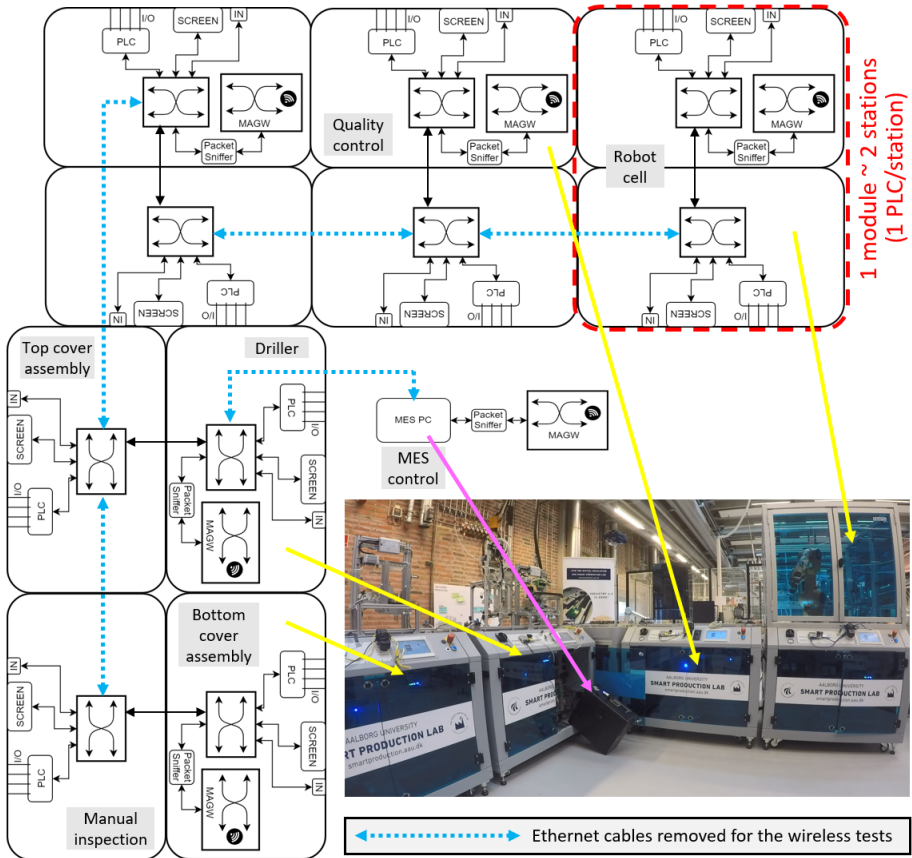


Fig. D.2: Overview of the fully-automated test assembly line at the AAU Smart Production Lab composed of 5 modules (10 process-specific stations), including reference network architecture and picture of the real setup.

state. The production of a singular product consists of 7 steps: 1) bottom cover placement at stack magazine, 2) hole drilling, 3) placement of a circuit board and fuse at robot cell, 4) video-based quality assurance, 5) top cover placement at stack magazine, 6) final product manual inspection and packaging. As a safety feature, even if they are not busy, the stations continuously check their connection status towards the MES PC, verifying whether the communication is active and reliable. In the case that the maximum survival time period of 2 s is exceeded, the station is stopped, and an alert is displayed at the MES. It should be noted that the described system tolerates a slightly higher latency than the one described for typical systems in Fig. D.1.

From a communication perspective, the only time-sensitive communication is the one happening at PLC I/O and fieldbus level. In contrast, the overall

inter-module and MES communication is the delay-tolerant [10]. Therefore, in this paper, we focus only on providing communication at MES level between the different modules and MES controller PC by replacing the Ethernet connection with wireless, while keeping the internal components of the modules wired. By doing this, both greater flexibility and re-configurability of the production line is enabled.

As previously mentioned, an important aspect to be considered, from an industrial network architecture point of view, is backwards compatibility and seamless integration such that no existing equipment needs to be modified or replaced to enable wireless connectivity. This is an essential aspect, especially when dealing with legacy industrial setups (which are the vast majority in the current real-world operational deployments). This requires that the wireless network is self-contained and configurable independently from the wired network. To achieve this, we introduce our proposed multi-access gateway (MAGW) solution, which is presented in detail in Section D.3. To integrate our wireless setup in the system, the cables between modules are removed and a total of 6 gateways are installed in the system (as depicted in Fig. D.2): one of them is connected to the MES PC, while the other 5 are installed at each of the 5 modules (serving two stations each).

The gateways provide seamless connectivity at MES level over both LTE and Wi-Fi. LTE connectivity is provided by one of the major mobile operators in Denmark via its commercial wide area network, optimized for mobile broadband use. A dedicated access point name (APN) is deployed at the operator's core network allowing for inter-communication between gateways based on static LTE IP configurations. Under such configuration, the traffic between gateways is routed at the core without being forwarded to the public internet. In other words, a "private network slice" is created using the public infrastructure, similar to what 5G NR will consider for the support of vertical use cases, but without prioritized resource allocation in the radio access network (RAN). With respect to the Wi-Fi connectivity, it is provided through dedicated infrastructure, in a similar way to what it is done nowadays in any enterprise deployment. Our gateways have the further capability of providing transport layer hybrid access (or multi-connectivity) by combining the connections over both technologies. Combining LTE with Wi-Fi provides the additional benefit of having fallback options in case of network errors, thereby increasing the availability and reliability of the overall system [11].

D.3 Multi-access Wireless Gateway

As previously stated, the purpose of the multi-access gateway is to provide seamless wireless connectivity via LTE, Wi-Fi or both. To create these multi-access gateways, the functionality and components depicted by the block

diagram in in Fig D.3.a, were implemented on UP-boards equipped with LTE and Wi-Fi dongles via USB ports (as illustrated in Fig. D.3.b). Details on the specific hardware and software used are summarized in Table D.1.

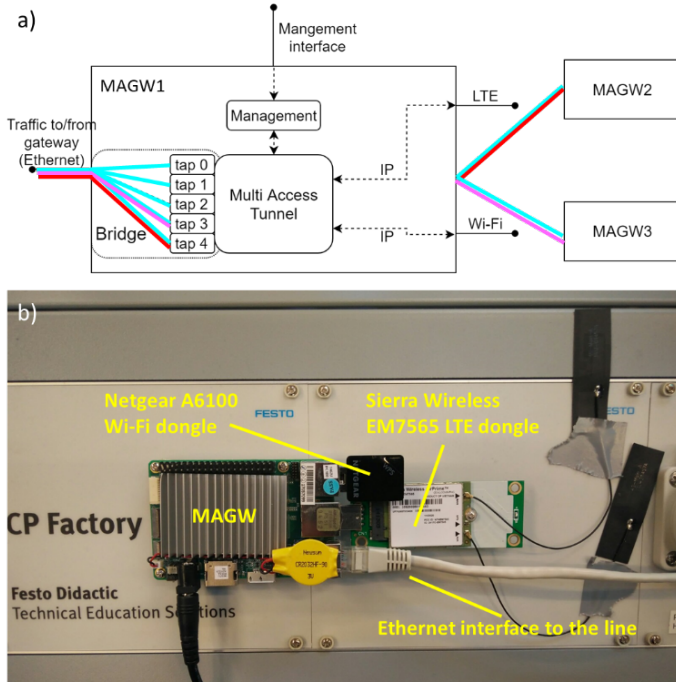


Fig. D.3: Overview of the multi-access gateway: a) block diagram including the internal components, b) picture of the prototype implementation illustrating the different Ethernet, LTE and Wi-Fi interfaces.

Table D.1: Hardware and software details for the MAGW and traffic sniffers.

HW/SW	Details
Gateway	UP-Board, intel x5-Z8350, 4 GB RAM
OS	Ubuntu 18.04 LTS
SW	Custom Multi-access Tunnel
NIC model	RealTek 8111G, 1 Gb/s
Wi-Fi dongle	Netgear A6100, 802.11ac
LTE dongle	Sierra Wireless EM7565, Cat 12
Sniffer	Raspberry Pi Model 3B, Broadcom BCM2837, 1 GB RAM
OS	Raspbian Stretch
SW	Tcp-dump v4.9.2
NIC model	Apple 10/100 USB ethernet adapter

To setup the inter-gateway network, a list containing the address and port information from all gateways in the network is distributed via the management interface (logical interface used for configuration that can be accessed either physical via USB port or over any of the wireless networks when connected). To establish the actual communication, each gateway creates a virtual port that is bridged to each of the remote gateways. Thereby, a gateway is seen as a switch port with a set of devices (MAC addresses) behind it from the perspective of the Ethernet network. This means the physical Ethernet traffic is automatically routed through the correct gateway based on the destination MAC address, as illustrated in Fig. D.3.a. Each of the taps illustrated in the figure corresponds to a unique remote gateway. As an example, the traffic flow associated with the taps with destinations MAGW2 and MAGW3 are color-coded in red and magenta, respectively. The cyan flow indicates broadcast traffic and is therefore routed to all remote gateways.

When an Ethernet packet is routed to the correct virtual port, the multi-access tunnel module will extract the Ethernet frame, and encapsulate it in a user datagram protocol (UDP) packet along with a sequence number and a unique ID used to identify which virtual port it should be routed to at the remote side. Once received at the remote side, the received UDP packet is decapsulated into an Ethernet packet and forwarded to the Ethernet interface of the gateway. A complete sequence diagram is depicted at the top part in Fig. D.4.

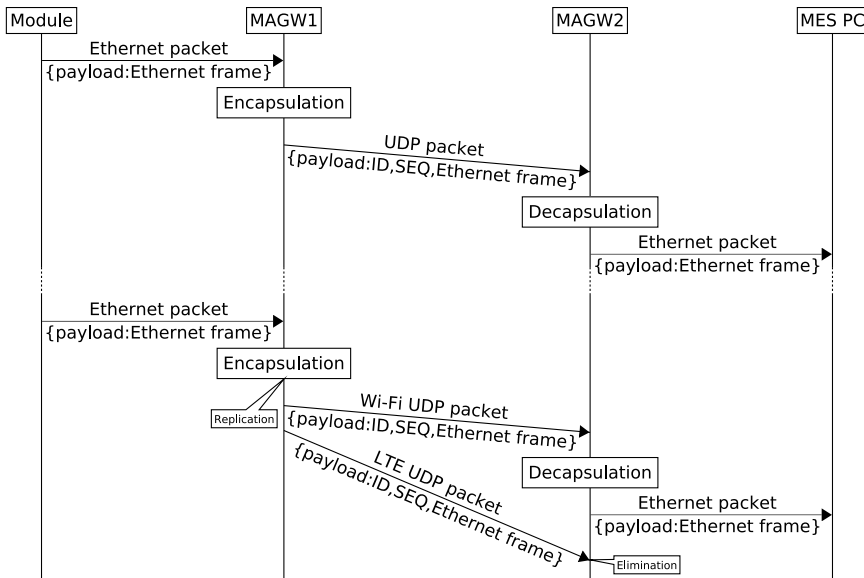


Fig. D.4: Module to MES PC sequence diagrams. The top one illustrates the single connectivity case, while the one in the bottom shows the multi-connectivity case.

If the distributed network configuration specifies that both the local and remote gateways have multiple connection options available between them (hybrid access is enabled), the encapsulated packet is replicated across all specified connections, i.e. the same packet is sent over LTE and over Wi-Fi. At the destination gateway, packets with the same origin network are identified by the unique ID and potential duplicates are eliminated based on the sequence number of the packet. This mechanism is similar to the 802.1CB FRER [12] but at IP-level. An example of this multi-connectivity sequence is illustrated in the bottom part of Fig. D.4.

D.4 Test setup

To evaluate the performance of the wireless MES communication over LTE and Wi-Fi and benchmark it against Ethernet, we have examined the following cases:

1. Ethernet (reference MES performance).
2. Public LTE: the system operates over the wireless gateways, being these connected to the commercial Telenor LTE network. The network operates at 2.6 GHz with 20 MHz bandwidth. The closest macro base stations are placed outside the test lab at an approximate distance of 1 km. The performance in this case is subject to the LTE network load.
3. Wi-Fi: the system operates over the wireless gateways, being these connected to a dedicated Wi-Fi infrastructure based on an 802.11n TP-Link N750 AP (with customized OpenWrt 18.06.2 OS) located inside the lab in the vicinity of the production line at approximately 5-10 m distance. The dedicated Wi-Fi network is operated in the 5 GHz band over the unoccupied channel 161 with 20 MHz bandwidth. This setup emulates an enterprise industrial deployment with strict frequency channel planning.
4. Wi-Fi with interference (w.i.): the system operates over the wireless gateways connected to the dedicated Wi-Fi infrastructure described in 3. However, in this interference case, the dedicated Wi-Fi network is operated in the 5 GHz band over channel 132 with 20 MHz bandwidth. This network overlaps with the Aalborg University Wi-Fi network, which was loaded by starting different traffic flows from 5 devices (streaming service, file download and periodic traffic) while the manufacturing line is running. This setup emulates a non-optimized Wi-Fi case, where the network in charge of controlling the equipment is affected by external interferers.
5. Hybrid access (HA) over public LTE and Wi-Fi: the gateways are simultaneously connected to the networks detailed in 2. and 3.

D.4. Test setup

6. Hybrid access over public LTE and Wi-Fi with interference: the gateways are simultaneously connected to the networks detailed in 2. and 4.

Each case is evaluated by configuring the production line to manufacture several products during a period of 30 minutes, where actions are continuously occurring according to the flow presented in Section D.2. In terms of generated data traffic in the reference configuration operating over industrial Ethernet, this corresponds to 17k-20k packets sent between the MES controller and the FESTO modules, with an average load of 6.7 kb/s, when aggregating all the samples to and from the modules. In terms of transmission behavior, no distinction between TCP and UDP traffic was done, and 75 % of the traffic is unicast, while the remaining is multicast/broadcast traffic. The average packet size is 63-70 bytes, which will translate into 88-108 bytes when considering wireless MES access over the gateways (excluding the technology-specific MAC/PHY overhead for Wi-Fi and LTE). On average, the modules (grouping two stations/PLCs) and MES PC produce one packet every 148 ms and 112 ms, respectively. A similar behavior was observed for all modules, with minor variations mainly in the upper bounds (above 99%-iles). It should be noted that similar traffic patterns have been observed in larger systems in real factory scenarios. These are not very demanding numbers in terms of data traffic nature, which further back up our view with respect to the potential of using existing wireless technologies as LTE or Wi-Fi for provisioning this delay-tolerant communication at MES level.

To evaluate the performance of the system under the different configurations, the key performance indicators (KPIs) are the packet error rate and one-way latency [13]. The latter can be related to system survival time i.e. the maximum time over which the system can function without new control information being received (in this case commands to and from the MES controller), which is configured to be 2 s in the analyzed setup. To compute the KPIs, we use the measurement method presented in [10] and illustrated in Fig. D.5. This method considers a synchronized measurement setup consisting on packet sniffers that are placed at the Ethernet interfaces of the multi-access gateways on each module and the MES controller PC, logging all the incoming and outgoing traffic passing through. During the measurements, these sniffers synchronize (with an accuracy of 8 μ s on average) to a common network time protocol (NTP) server via a dedicated external network. This allows us to extract the precise departure and arrival time of each packet in the network with respect to a common reference. Further details about the hardware and software of the sniffers are given in Table D.1.

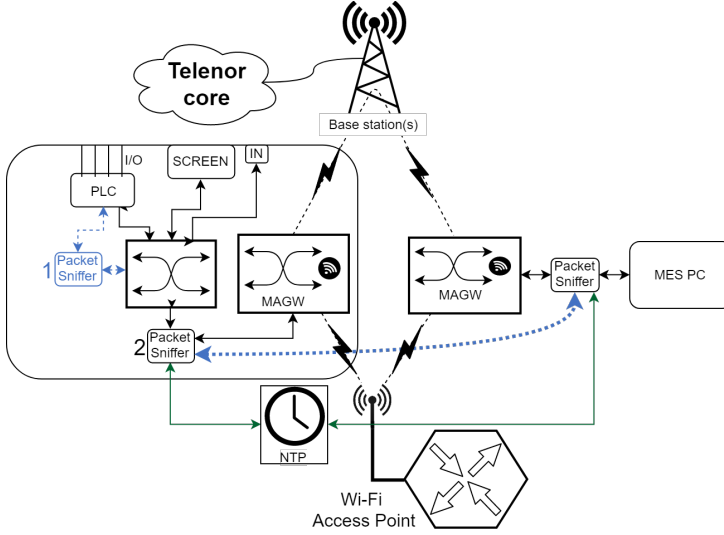


Fig. D.5: Simplified measurement setup between MES and a single station, considering the Ethernet reference case (blue dotted line) and both wireless cases: LTE and Wi-Fi.

D.5 Performance Results

The empirical complementary cumulative distribution function (CCDF) of one-way MES communication latencies is displayed in Fig D.6 for the 6 different test cases along with a summary of the latency values for different percentiles in Table D.2. As expected, the reference system over Ethernet presents the best performance (low latency, low jitter), presenting a maximum one-way latency of 0.43 ms. This indicates that such system and particular configuration could even cope with the more demanding latency requirements of the lower layer systems dictated by the automation pyramid in Fig. D.1.

Communications via public LTE experiences higher latency than the dedicated Wi-Fi case: 47.6 ms median one-way latency with LTE vs 2.81 ms for Wi-Fi. This can be explained from the fact that the considered LTE network was a public network, optimized for mobile broadband, and without control of the load; whereas the considered dedicated Wi-Fi configuration operated over unoccupied non-interfered spectrum and with a small number of devices (modules + MES controller).

Table D.2: Summary of one-way latency and packet error rate measurement results for the 6 test cases.

Technology	Selected latency statistics					Packet loss burst size					PER [%]			
	Min [ms]	Median [ms]	99%-ile [ms]	99.9%-ile [ms]	Max [ms]	σ [ms]	[%] <100 ms	[%] <1s	1	2		3	4	5
1. Ethernet	0.09	0.18	0.31	0.36	0.43	0.04	100	100	-	-	-	-	-	0
2. Public LTE	27	47.6	105	234	468	17.9	98.41	100	500	29	4	5	1	0.5
3. Wi-Fi	1.67	2.81	4.35	10.5	88	1.42	100	100	15	1	-	-	-	0.02
4. Wi-Fi w.i.	1.74	4.43	217	488	1187	43	97.07	99.98	71	1	-	-	-	0.06
5. HA over Public LTE & Wi-Fi	2	3.2	5	9	44	0.8	100	100	1	-	-	-	-	0.001
6. HA over Public LTE & Wi-Fi w.i.	1.8	3.1	31	59	217	5.25	99.99	100	32	-	-	-	-	0.03




Fig. D.6: One-way latency distributions from the trials for the various technologies.

The behavior of the wireless MES over Wi-Fi with interference is, in general, also better than LTE – with a median one-way latency of 4.43 ms. However, for the higher percentiles, it presents a degraded performance due to the interference from the external co-channel network. This degradation resulted in a maximum latency of 1.1 s (still lower than the allowed survival time within the system). This has a large impact on the jitter (σ), reaching up to 43 ms, which is much larger than the one from LTE (17.9 ms) and Wi-Fi without interference (1.42 ms). Without external interferers, Wi-Fi latency lies well below the most demanding MES communication requirement (100 ms) and could provide such service with a 99.99+% of reliability. The tested public LTE connection, however, would only be suitable for such demanding MES deployments if the required reliability level is below 99%.

When examining the packet error rate (PER) behavior during the various tests, LTE presents a significantly higher rate than Wi-Fi with and without interference (0.5% and 0.02-0.06%, respectively). LTE and Wi-Fi with and without interference present mostly sporadic single packet losses. However, the wireless MES communication over LTE also present larger bursts of losses in some cases, reaching up to 2-5 consecutive lost packets. As the LTE is a public network, it is difficult to determine the exact cause of these burst of errors since the connection is not only subject to the radio channel but also to any optimizations and specific configurations made by the operator.

Combining the two technologies through hybrid-access or multi-connectivity, yields to a significant improvement in terms of both latency and PER, even in the case where Wi-Fi is subject to interference. The best performance of the wireless MES system was achieved with configuration 5 (combination of public LTE and dedicated Wi-Fi) with a maximum latency of 44 ms and a single packet loss (resulting in a PER of 0.001%). Multi-connectivity techniques help in both tested scenarios, even in the highly interfered scenario of case 6 (where multi-connectivity over public LTE and Wi-Fi with interference was tested), achieving a much better performance than the LTE-only or interfered Wi-Fi-only cases. This is due to the fact that the access over each technology is highly uncorrelated, which leads to high diversity gain.

The evaluated wireless configurations experienced higher latencies than the reference wired Ethernet case. However, the latency can still be contained below 100 ms on average, and as the considered industrial traffic can tolerate delays of 100-1000 ms [2], using wireless is nevertheless attractive from a flexibility point of view.

D.6 Conclusions

In this paper, we evaluated the suitability of wireless technologies such as public LTE and Wi-Fi in supporting communication at Manufacturing

Execution System (MES) level in industrial scenarios. The study is carried out in a production line at the Smart Production Lab at Aalborg University. Wireless connectivity is obtained via our designed multi-access gateway that seamlessly leverage LTE, Wi-Fi and their combination via hybrid access with frame replication and elimination. For Wi-Fi, both cases of free and busy operational channels are considered, reflecting the cases of strict frequency planning and non-optimized deployment, respectively. Measurement results reveal for all the studied solutions one-way latencies significantly below the 2 s survival time set by the MES of our setup and fitting to the 100-1000 ms range of typical MES. Wi-Fi provides lower median latencies than public LTE, although the presence of interference increases 99.9 %-ile latencies up to 490 ms, whereas public LTE offers 230 ms. The usage of hybrid access leveraging both LTE and Wi-Fi via packet duplication leads to 99.9 %-ile latencies lower than 100 ms, thus coping even with the stricter delay communication requirements seen in MES deployments.

Next research steps include testing Wi-Fi scalability for larger MES deployments, deploying and optimizing a private LTE network for industrial automation, and assessing the impact on production capability by looking at trade-off among higher latency and better flexibility of wireless.

Acknowledgment

This research is partially sponsored by the EU H2020-ICT-2016-2 project ONE5G. The views expressed in this paper are those of the authors and do not necessarily represent the project views. Also, the authors would like to thank the Manufacturing Academy of Denmark (MADE) and the AAU Smart Production Lab for their interest in the presented proof-of-concept and the support throughout the different activities. The wireless equipment used in the setup is partly funded by the Innovation Fund Denmark (IFD) under File no. 7039-00009B.

References

- [1] B. Chen, J. Wan, L. Shu, P. Li, M. Mukherjee, and B. Yin, "Smart Factory of Industry 4.0: Key Technologies, Application Case, and Challenges," *IEEE Access*, vol. 6, pp. 6505–6519, 2018.
- [2] A. Nasrallah, A. S. Thyagaturu, Z. Alharbi, C. Wang, X. Shao, M. Reisslein, and H. ElBakoury, "Ultra-Low Latency (ULL) Networks: The IEEE TSN and IETF DetNet Standards and Related 5G ULL Research," *IEEE Communications Surveys Tutorials*, vol. 21, no. 1, pp. 88–145, 2019.

References

- [3] T. Carlsson, "Industrial Ethernet is now bigger than fieldbuses." [Online]. Available: <https://www.anybus.com/about-us/news/2018/02/16/industrial-ethernet-is-now-bigger-than-fieldbuses>
- [4] V. K. L. Huang, Z. Pang, C. A. Chen, and K. F. Tsang, "New Trends in the Practical Deployment of Industrial Wireless: From Noncritical to Critical Use Cases," *IEEE Industrial Electronics Magazine*, vol. 12, no. 2, pp. 50–58, 2018.
- [5] "Industrial Wireless LAN." [Online]. Available: <https://siemens.com/iwlan>
- [6] D. K. Lam, K. Yamaguchi, Y. Shinozaki, S. Morita, Y. Nagao, M. Kurosaki, and H. Ochi, "A fast industrial WLAN protocol and its MAC implementation for factory communication systems," in *IEEE Conference on Emerging Technologies Factory Automation (ETFA)*, 2015, pp. 1–8.
- [7] Y. Wei, Q. Leng, S. Han, A. K. Mok, W. Zhang, and M. Tomizuka, "RT-WiFi: Real-Time High-Speed Communication Protocol for Wireless Cyber-Physical Control Applications," in *IEEE Real-Time Systems Symposium*, 2013, pp. 140–149.
- [8] "MulteFire Release 1.0 Technical Paper: A New Way to Wireless."
- [9] 3GPP, "Study on Communication for Automation in Vertical Domains," Tech. Rep. TS 22.804 V16.0.0, Apr. 2018.
- [10] I. Rodriguez, R. S. Mogensen, E. J. Khatib, G. Berardinelli, P. Mogensen, O. Madsen, and C. Møller, "On the Design of a Wireless MES Solution for the Factories of the Future," *Global IoT Summit (GIoTS)*, 2019.
- [11] R. S. Mogensen, C. Markmøller, T. K. Madsen, T. Kolding, G. Pocovi, and M. Lauridsen, "Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications," *VTC Spring* 2019.
- [12] "IEEE Standard for Local and Metropolitan Area Networks – Frame Replication and Elimination for Reliability," *IEEE Std 802.1CB-2017*.
- [13] M. Felser, "Real Time Ethernet: Standardization and implementations," in *IEEE International Symposium on Industrial Electronics*, Jul. 2010, pp. 3766–3771.

Paper E

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

A. Fink, R. S. Mogensen, I. Rodriguez, T. Kolding,
A. Karstensen, and G. Pocovi.

The paper has been published in the
26th European Wireless Conference; European Wireless 2021

© 2021 VDE

The layout has been revised.

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.

E.1 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 a realistic 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 E.2 presents an overview of Wi-Fi with focus on mobility, including approaches to improve its performance. Section E.3 details aspects related to tests of the performance with different network setups. Section E.4 presents the results of real-world measurements and identifies areas with potential for improvements. Finally, Section E.5 concludes the paper.

E.2 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. E.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.

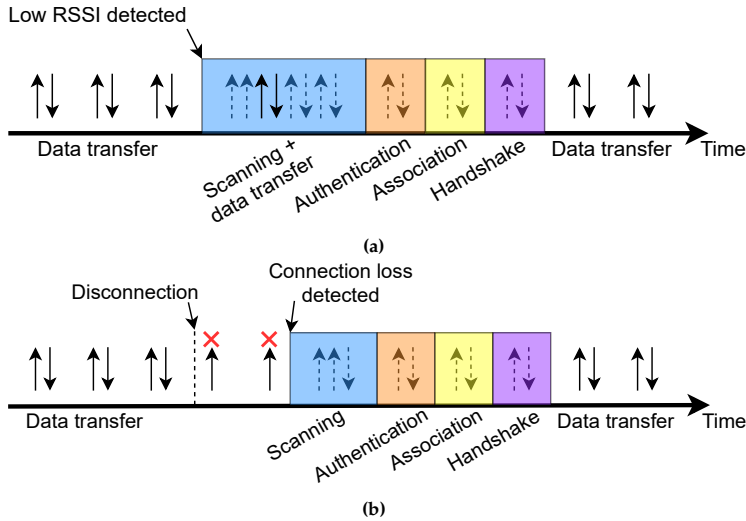


Fig. E.1: Simplified illustration of the 4 stages which occur during a handover for the two different situations: a) handover triggered by low RSSI, and b) handover triggered by detecting that the AP cannot be reached (disconnection).

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. E.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. E.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.

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.

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.

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.

E.3 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. E.2.



Fig. E.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.

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. While the performance achieved over a single testing environment may not resemble the expected performance in all industrial settings, the structured choice of scenario, deployment, network and mobility conditions, and traffic and average interference models are realistic. This allows for capturing, characterizing and explaining main Wi-Fi latency performance trends expected in IIoT mobile use case scenarios.

Measurements were collected using the STA described in Table E.1. 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

Table E.1: 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

E.3. Test Setup

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.

Under the current configuration and measurement route (see Fig. E.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 E.2, 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 above 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. E.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.

E.4 Test Results

A heatmap of the RSSI for a single measurement lap under idle network conditions is illustrated in Fig. E.3. Here it is shown that, in this particular example, the STA did not roam to AP 1 since the RSSI was approximately -70 dBm. This was, however, not the case for all of the measurements, as the STA would occasionally roam to it when different propagation conditions applied or higher interference was present.

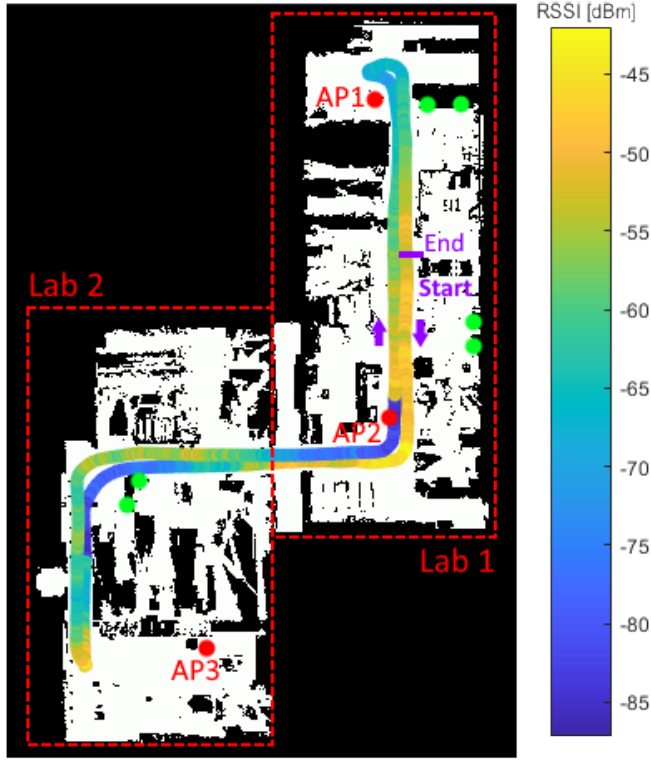


Fig. E.3: RSSI heatmap for a single measurement lap for baseline scheme and idle network configuration. The location of the APs and the controlled traffic load source/sink STAs are indicated by the red dots and green dots, respectively.

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 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. E.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-

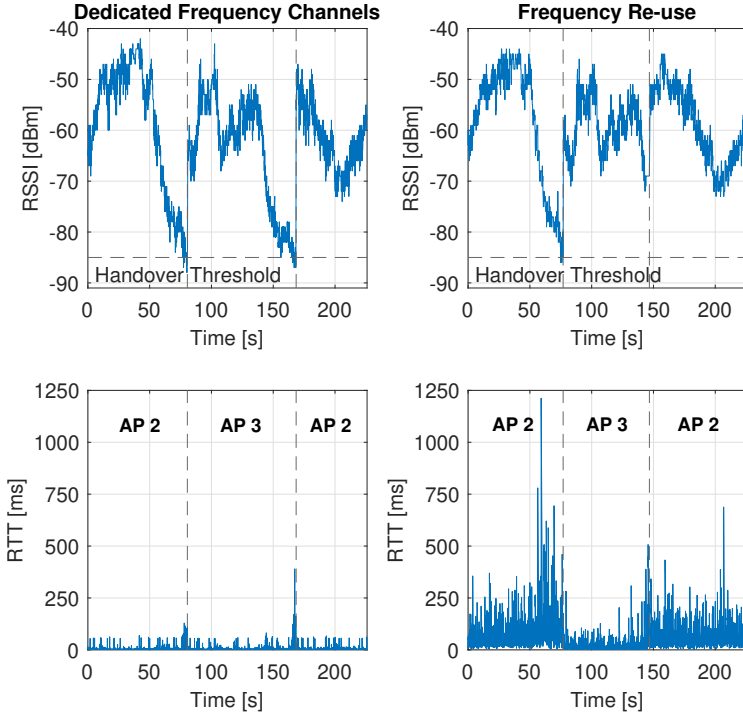


Fig. E.4: RSSI and RTT measurements for a single measurement lap for the baseline scheme and controlled load network configurations with dedicated frequency channels (on the left), and with frequency re-use (on the right).

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. E.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. E.5 for the different Wi-Fi 6 configuration schemes in idle network conditions.




Fig. E.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.

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

Fig. E.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 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.




Fig. E.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.

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. E.7 with their key statistics and Packet Error Rate (PER) summarised in Table E.2. 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. E.4.




Fig. E.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 E.2: 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%
	Idle network (single STA)	1.6 ms	3.2 ms	58.3 ms	1.5 ms	0.045%
Optimised Scanning	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%
	Idle network (single STA)	1.6 ms	4.0 ms	118.0 ms	2.4 ms	0.046%
IEEE 802.11r	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%

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 $\sim 0.06\text{--}0.07\%$ for the dedicated frequency channels case and to the $\sim 0.04\%$ for the idle network single STA case.

E.5 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 a realistic 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 ~ 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 ~ 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, R. S. Mogensen, A. Schjørring, M. Razzaghpour, R. Maldonado, G. Berardinelli, R. Adeogun, P. H. Christensen, P. Mogensen, O. Madsen, C. Møller, G. Pocovi, T. Kolding, C. Rosa, B. Jørgensen, and S. Barbera, "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

References

- (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/>

Paper F

5G Swarm Production: Advanced Industrial Manufacturing Concepts enabled by Wireless Automation

I. Rodriguez, R. S. Mogensen, A. Schjørring, M. Razzaghpour,
R. Maldonado, G. Berardinelli, R. Adeogun, P. H. Christensen,
P. Mogensen, O. Madsen, C. Møller, G. Pocovi, T. Kolding,
C. Rosa, B. Jørgensen, and S. Barbera.

The paper has been published in the
IEEE Communications Magazine, vol. 59, no. 1, pp. 48-54, January 2021.

© 2021 IEEE

The layout has been revised.

Abstract

This paper presents an overview of current Industry 4.0 applied research topics, addressed from both an industrial production and a wireless communication points of view. A roadmap towards achieving the more advanced industrial manufacturing visions and concepts, such as the "swarm production" (non-linear and fully de-centralized production), is defined, highlighting relevant industrial use cases, their associated communication requirements, as well as the integrated technological wireless solutions applicable to each of them. Further, the paper introduces the Aalborg University 5G Smart Production Lab, an industrial lab test environment specifically designed to prototype and demonstrate different industrial IoT (IIoT) use cases enabled by the integration of robotics, edge-cloud platforms and autonomous systems operated over wireless technologies such as 4G, 5G, and Wi-Fi. Wireless performance results from various operational trials are also presented for two use cases: wireless control of industrial production and wireless control of autonomous mobile robots.

F.1 Introduction

The ongoing fourth industrial revolution (Industry 4.0) relies on the integration of cyber physical systems, industrial internet-of-things (IIoT), and cloud computing technologies, as a major driver for achieving a highly flexible and reliable manufacturing in the factories of the future [1]. On top of this integration, optimized wireless technologies will play a pivotal role. Wireless technologies will allow to replace cables (favoring a faster re-configuration of the production facilities and an overall reduction of the cost of deployment), and also to enable new industrial use cases requiring full mobility support [2].

However, under these premises, research addressing Industry 4.0 domains has been typically done in an isolated manner by vertical and horizontal sectors, without jointly accounting for all the components required to succeed in the long-term visions. For example, the industrial production and manufacturing sector has tended to focus on developing concepts and visions while slightly overlooking the communication aspects [3], by giving reliable control data flows for granted, in their advanced manufacturing systems visions, even when considering cloud-based soft programmable logic controllers (PLCs), or autonomous mobile robots (AMR) operating over wireless [4]. On the other hand, due to lack of strong direct interaction with the vertical sectors, it took some time for the wireless communication sector to gather relevant data, such as typical traffic patterns, data rates and tolerable latency thresholds applicable to different industrial use cases [5], which are key elements in the design of systems targeting ultra-reliable and low-latency communications (URLLC) for Industry 4.0, such as 5G. Fortunately, the situation has changed in the last years, and the new releases of 5G targeting time-sensitive networks, have had

more direct impact from verticals than ever.

This proves that it is of paramount importance nowadays, to have a double helix approach between both the manufacturing and communication sectors, working together on the practical integration of wireless solutions with the different manufacturing use cases. Integrated wireless solutions can be optimized by having a better understanding of current and envisioned scenario-specific use cases and associated communication requirements. This will ensure that, an accurate mapping between specific actions in the manufacturing process and wireless technologies, capable of supporting such application requirements, is done. As legacy industrial systems will continue being of importance, and not all industrial use cases will require URLLC, deploying 5G might be an overkill in certain cases, which leaves some room in the industrial wireless ecosystem for other technologies such as 4G and Wi-Fi. In the other direction, advanced manufacturing concepts and their associated architectures or control protocols, could be optimized or evolved by carefully considering the available wireless communication and cloud-computing capabilities of the envisioned integrated systems. Although 5G will be able to support down to 500 us latency with high reliability, there will still be certain industrial use-cases requiring much lower latencies, making some space for beyond 5G (B5G) or 6G technologies in future wireless manufacturing. Similarly, other wireless systems such as indoor positioning systems based on ultra-wideband (UWB) technologies, might also be relevant in those use cases requiring precise location information, in case it cannot be obtained from any other wireless source.

In this paper, we present advanced Industry 4.0 visions, where the ultimate goal is to achieve and demonstrate the so-called "swarm production" where, differently from current traditional manufacturing systems (based on a linear and centralized production concept), in which products are manufactured sequentially over production modules where their respective PLCs and input/output (I/O) systems are connected by wires or buses to a centralized controller; wireless is integrated with the manufacturing system, allowing to have production modules distributed across the factory hall with their PLCs and I/O systems operated in remote edge-cloud configuration, and AMRs are used to move items between them (non-linear and de-centralized production). Swarm production will allow for the maximum level of flexibility and re-configuration of the production process, and will require robust automation and ultra-reliable cloud-based control. Thus, 5G is considered as the baseline technology for this use case. As part of the presented visions, outlined through a double helix approach between the Department of Materials and Production and the Wireless Communication Networks Section at Aalborg University, a roadmap with multiple steps is defined, addressing different sub-components of the swarm production concept, characterized by relevant related use cases which, as it will be explained later, can be realized over

different wireless technologies. To demonstrate the different steps, different prototypes are designed, built and tested in an unique Industry 4.0 wireless testing ecosystem.

F.2 Manufacturing Industry Goals and Wireless Automation Evolution

The manufacturing industry envisions factories as highly flexible facilities, where it will be possible to cope with the increasing demand of highly customized products, while reducing or at least maintaining resource and cost efficiency [4]. Flexibility can be achieved by leveraging the swarm production concept, i.e. de-centralized and non-linear production processes where products are transported by AMRs between manufacturing stations distributed across the factory hall. In order to successfully achieve such level of adaptable intelligent production, the integration of different technological components such as cloud-computing, 5G or B5G communications (able to provide URLLC to the different components in the system), robotics, autonomous systems and highly-accurate localization systems are essential.

As jumping directly from existing production schemes to swarm production might be difficult, we propose a simple reference evolution roadmap for the production process, where we define the transition path from traditional manufacturing systems to swarm production by the 4 steps summarized in Table F.1. Such steps have been carefully selected by analyzing the specific manufacturing and high-level communication needs and the availability of technological components, and can be taken sequentially, but also independently, as they have been defined around different areas of focus. In a sequential manner, the first step (1), applicable to the traditional production systems, would be to replace part of the cables in the production lines with wireless communication links, and set up a cloud-based manufacturing control server, replacing the current local line controllers [6]. By doing this, a first level of flexibility is achieved by enabling an easier re-configuration of the production facilities as compared to wired setups. In order to achieve this step, a robust wireless communication, capable of coping with the control traffic, to and from the different modules, is necessary. The second step (2) targets PLCs which are, nowadays, programmed to perform specific actions and require manual software upgrades if a change is needed. Migrating the intelligence of the PLCs to the cloud, by relying on cloud-computing and the URLLC capabilities of the applied wireless technologies, will add an extra degree of flexibility to the production system by enabling a faster deployment of new or product-specific functionalities via software to the different production modules [7]. Architectures based on cloud-PLCs will be more scalable and

will allow for having lighter production modules in terms of hardware, as now the processing power is moved to the cloud [8]. The communication needs are more demanding in this second step compared to the first one, as much more information will need to be transported from each of the industrial modules to the cloud controller.

The third step of the roadmap (3) focuses on the evolution of AMR fleets, where the main objectives are related to moving some of the robot functionalities to the cloud (similarly to the previous PLC case). In particular, the interest is put on having cloud-based localization and navigation. By doing this, a number of expensive and processing power-hungry onboard sensors (i.e. LiDaRs) could be removed, making the robots much more cost efficient and easier to manage [9]. In order to achieve this, different strategies can be used, such as for example, integrating the robots and the fleet manager server with an external high accuracy positioning system (such as UWB), and/or relying on the connectivity of the robots over a high bandwidth reliable technologies (such as 5G), allowing for the transfer of real time HD videos or pictures from each of the robots to the cloud fleet manager server. All the available information in the cloud can later be combined to create a shared world model (cloud robotics) [10]. In this case, the communication needs are even a bit more demanding than in the second step, as ultra-reliability is essential due to the safety-related critical communication aspects of the control of mobile robots. Once steps (2) and (3) are completed, the integration of both, would result into (4). Swarm production is founded on the intelligent cloud control of PLCs and AMRs [11], allowing to transport products between production modules via AMRs instead of relying on product carriers running over conveyor belts. In terms of communication needs, this step groups all the needs from the previous steps, i.e. accurate positioning, high-bandwidth, as well as ultra-reliable communication to both mobile and static units, to guarantee a synchronized performance between all the entities that integrate the advanced industrial manufacturing scenario. Thus, 5G is seen as a key technology for achieving this futuristic production concept.

The proposed roadmap is not universal. Not all industrial production entities consider advanced manufacturing concepts such as the swarm production in their digitalization strategies. From our three years of conversations with different entities of the Danish manufacturing industry, we realized that, in general, there is a huge hype about cloud control and cloud monitoring, but also that they give paramount importance to legacy machinery. Not all companies will have the chance to invest in the most advanced solutions, but still introducing a few wireless components for specific communication needs, might result in a considerable gain for them. This creates a vast diversity and heterogeneity of use cases with very specific communication needs and requirements, but also opens for the potential use of other wireless technologies, apart from 5G, in the future industrial automation ecosystem [12].

Table F.1: Simplified Manufacturing and Production Roadmap towards Swarm Production

Step	Focus Area	Actions
(1)	Wireless Production	<p>Remove cables between manufacturing line modules. Cloud-based production control.</p> <ul style="list-style-type: none"> • Manufacturing target: Flexibility, reconfiguration. • Communication needs: robust low-throughput delay-tolerant wireless communication links to static units.
(2)	PLC	<p>Remove hardware, use cloud-based soft PLCs instead.</p> <ul style="list-style-type: none"> • Manufacturing target: Faster and cheaper adaptation of new functionalities. • Communication needs: reliable high-throughput low-latency wireless communication links to static units.
(3)	AMR	<p>Move functionality (localization and navigation) to the cloud. Investigate new localization techniques.</p> <ul style="list-style-type: none"> • Manufacturing target: more efficient fleet management/cheaper robots and shared world model (cloud robotics). • Communication needs: ultra-reliable high-throughput low-latency wireless communication links to mobile units.
(4)	Swarm Production	<p>Remove conveyor belts, make product carriers into small mobile robots.</p> <ul style="list-style-type: none"> • Manufacturing target: More flexible and robust automation. • Communication needs: ultra-reliable high-throughput low-latency wireless communication links to both static and mobile units.

In order to illustrate the heterogeneity of applicable wireless technologies, Table F.2 gathers a number of industrial use cases and associated communication requirements, mapped over the applicable technological candidates and roadmap steps. The table considers the following use cases: the manufacturing execution system (MES) - links between the centralized manufacturing controller and PLCs - of a FESTO CP Factory research production line [6], the MES and PLC-I/O links of an operational setup in a real factory, the MiR200-based AMRs - control links between the fleet operation manager and the PLC in the robot, and the envisioned swarm production, orchestrated over optimized robot communication and cloud PLC architectures. All communication-related parameters in the table are based on measurements over operational industrial-grade manufacturing equipment for the research production line MES, MES and I/O in the operational factory, and current baseline AMRs, while the rest, those related to evolution of the current systems and the targeted swarm production, are based on our own visions and educated research analysis. The presented use cases serve as reference to illustrate the applicability of the roadmap steps. The current implementations of both the research production line and the production line at the operational factory, can be evolved by applying the roadmap steps (1) and (2): replacing cables with wireless and moving intelligence from the lines to the cloud control. In the case of the AMRs, steps (2) and (3) would be applicable as the desire in this case is improving the current wireless navigation control, and moving most of the robot intelligence to the cloud. Finally, in the swarm production case, all previous use cases could be combined and optimized as part of steps (3) and (4), by coordinating and synchronizing the operation of production line modules and AMRs, by making use of advanced automation algorithms. Clearly, based on the different requirements, some of the use cases can be operated over wireless technologies different to 5G, such as for example the PLCs and AMRs current control schemes, which could be operated reliably over 4G or Wi-Fi, as it will be illustrated in Section IV. On the other hand, it should be noted that some of the evolution use cases will require B5G technologies as control closed loops in some I/O cases demand stringent deterministic sub-ms latencies, which are not achievable over 5G [13]. In general, we believe that Wi-Fi 6, will also play a role in some of these use cases - mainly in the static ones, while in those requiring mobility, its suitability will be subject to a tight coordination between access points in order to ensure a reliable handover management.

Table F2: Summary of industrial use cases and associated communication requirements

Use Case	Research Production		Operational Factory		AMRs		Swarm Production	
	Line - MES	UL/DL	MES UL/DL	I/O UL/DL	baseline UL	evolution UL	robots UL/DL	cloud PLC UL/DL
Link								
Number of devices	10		50	25	20	20	20	10
Aggregated throughput	10 kbps		100 kbps	1 Mbps	20 Mbps	0.2-2 Gbps		
Individual user throughput					1 Mbps	10-100 Mbps		
Average packet size	64 B		128-256 B	128 B			256 B	128 B
Average inter-packet time	200 ms		20 ms	1 ms			100 ms	5 ms
Maximum control loop latency (survival time)	2 s RTT		N/A - <i>but in the order of s</i>	2 ms RTT jitter-critical	1 s RTT	10-20 ms RTT	10 ms RTT	10 ms RTT
Applicable steps towards Swarm Production	(1) (2)		(1) (2)	(1) (2)	(3)	(2) (3)	(3) (4)	(2) (4)
Candidate wireless technologies	Wi-Fi 4G, 5G		Wi-Fi 4G, 5G	5G B5G	Wi-Fi 4G, 5G	Wi-Fi 6 4G, 5G	Wi-Fi 6 5G	Wi-Fi 6 5G, B5G

UL: uplink, DL: downlink, RTT: round-trip time

E.3 The AAU 5G Smart Production Lab

With the aim of building and demonstrating the swarm production concept, and the associated use cases, in realistic industrial environment conditions, an advanced Industry 4.0 wireless playground was established at AAU. The AAU 5G Smart Production Lab is a 1200 m² factory industrial lab, with access to a wide range of operational industrial-grade manufacturing and production equipment from different vendors including production line modules, robotic arms, AMRs, etc. The lab is currently equipped with multiple networks from different wireless technologies, ranging from local private deployments of 4G LTE, 5G NR, and different flavors of Wi-Fi (including the last version, Wi-Fi 6), to dedicated operator-managed network slices of 4G LTE and 5G NR, and a dedicated positioning system based on UWB radio technology. A summary of all available wireless technologies is given in Table F.3, along with a few technical details.

Table F.3: AAU 5G Smart Production Lab Wireless Capabilities

Network type	Wireless technology	Details
Local Private Network	5G NR Private (pNR)	Private 5G NR mini-core + pico BSs. 3.7 GHz SA, 100 MHz, TDD, 3 cells.
	4G LTE Private (pLTE)	2x Private mini-core + micro BSs. 3.5 GHz, 20 MHz, TDD, 3 cells each.
	Wi-Fi 6	2x Coordinated IEEE 802.11ax deployment. 5 GHz, 3 cells each, cloud management.
	Wi-Fi 5	Uncoordinated IEEE 802.11ac deployment. 2.4/5 GHz, 3 cells.
Dedicated Operator Network Slice	5G NR Dedicated (dNR)	Public core + dedicated APN/BS setup. 1.8/2.1/2.6 GHz, FDD, 3 cells.
	4G LTE Dedicated (dLTE)	Public core + dedicated APN/BS setup. 2.6 GHz, 20 MHz, FDD, 3 cells.
Positioning	UWB	Enterprise TDOA positioning solution. 8 anchors, <10 cm accuracy.

Fig. F.1 depicts the high-level architecture of the research testbed, exemplifying how the different industrial components, such as production modules or AMRs, can be connected and controlled over the multiple available wireless networks. The integration between the machinery and the different networks, is achieved via wireless multi-access gateways (GW), which also allow for simultaneous multi-connectivity over multiple networks [14]. Quite some effort has been done in designing the network management back-end that integrates all the deployments with the local edge-cloud, where the management of the production systems, AMR fleets and GW devices is centralized.

F3. The AAU 5G Smart Production Lab

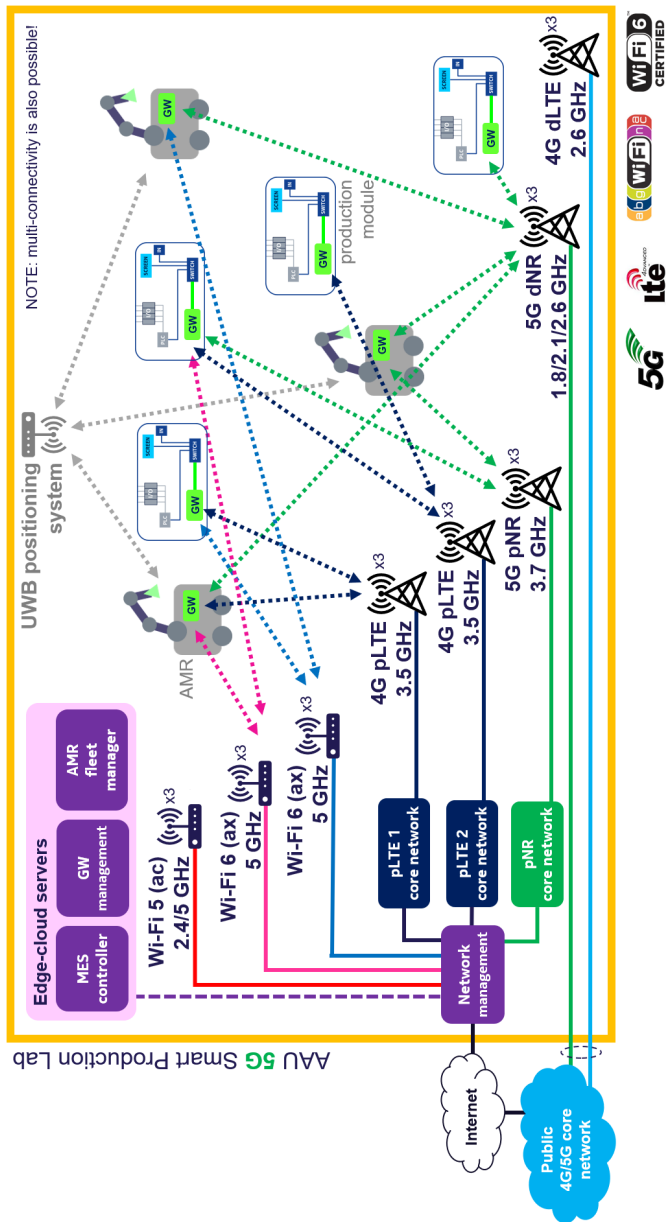


Fig. F.1: High-level overview of the Industry 4.0 wireless production testbed deployed at AAU, illustrating the swarm production concept with production modules and AMRs distributed over the production facilities, and the dedicated network infrastructure available for the testing of industrial use cases.

Such architecture allows to monitor the different networks and configure them as controlled test environments, with the possibility of recording network traces during the testing of the different use cases, enabling the opportunity of optimizing the network for the specific traffic and mobility patterns associated to that particular use case. This setup also allows to benchmark the performance of the different use cases under different 4G/5G licensed spectrum radio access network and core combinations, by comparing the performance of private network solutions with dedicated edge-cloud servers, with that from a dedicated network slice operating over a public core, for example. Moreover, the same use case could be tested over different Wi-Fi settings, providing a benchmark of the performance over unlicensed spectrum technologies.

F.4 Performance of Industrial Wireless Use Cases

Two of the industrial use cases described in Table F.2 have already been demonstrated over the testbed: the wireless control of industrial production and the control of AMRs. These tests were performed based on the testbed elements depicted in Fig. F.2, and their associated wireless performance results are presented in Fig. F.3, in terms of control-loop latency empirical complementary cumulative distribution functions (CCDF). This metric is of paramount importance, for understanding whether the communication requirements of a given use case can be fulfilled (i.e. its control-loop latency operated over a certain wireless technology is contained within the communication protocol bounds specified by its survival time), at high levels of reliability. For further details, values of average latency, jitter and packet error rate (PER) are also given in the legend of the figure for each of the tests.

F.4.1 Wireless Control of Industrial Production

For the first use case, step (1) of the roadmap was demonstrated. The cables between modules in the FESTO CP Factory research production line were removed and wireless GWs were installed instead, to provide control communication from the centralized MES controller deployed in edge-cloud configuration. More details about this specific use case are given in [14]. This industrial static use case has been evaluated over different Wi-Fi 5 configurations, 4G, and also more recently, 5G and Wi-Fi 6. The performance results presented in Fig. F.3 for this case, consider the full line individually operated over the different technologies with all its 7 modules connected over wireless, and compared to the results obtained when the line was operated over its standard Ethernet-based control configuration.

The best wireless performance for this use case was achieved over optimized Wi-Fi 5/6, with an average control-loop latency only 2-2.5 ms higher

F.4. Performance of Industrial Wireless Use Cases




Fig. F.2: Picture of some of the industrial machinery elements of the testbed, including the FESTO CP Factory production line in its standard wired configuration and one of the MiR200-based AMRs. The picture also depicts one of the wireless multi-access gateways used in the various experiments to interface the industrial equipment to the different networks.

than the reference one achieved over Ethernet. These results were obtained with single access point deployments with non-interfered Wi-Fi channels dedicated to the particular use case; which is usually not the case in operational factory scenarios. Operational factory scenarios are better represented by the non-optimized Wi-Fi 5 case, which exhibits a good average performance, but also presents much longer unbounded tails, reaching even up to 1 s at the lower percentiles. The performance over 4G is much more contained, over both the dLTE and pLTE configurations, ensuring a more deterministic low-jitter communication pattern, with control-loop latencies below 35 and 20 ms at median level, and 55 and 40 ms at the 99.99%-ile, respectively. The very similar shape of the dLTE and pLTE distributions can be explained due to the use of equipment from the same vendor in the radio access, and the offset difference between them, is due to the core network configuration: the dLTE relies on the public core of the operator, while the pLTE is based on a local core, which reduces the overall latency. The 5G pNR configuration, also based on a local core network, offers better control-loop performance than 4G, with less than 10 ms median and 22 ms at the lower percentiles, achieving the same reliability than the optimized Wi-Fi configurations at the tails. It should be noted that, the 5G used in the test, was a first out-of-the-box release, and thus, there is still plenty of room for optimization. In any case, as the control-loop latency requirements dictated by the maximum survival time (2 s) were al-

ways fulfilled, the wireless manufacturing system operated reliably (without interruptions) over all Wi-Fi, 4G and 5G technologies. This was the case even when packet loss was observed, meaning that the higher layer mechanisms were able to correctly handle the communication errors. From a manufacturing performance perspective, it is difficult to evaluate the impact of the increased latency introduced by the wireless technologies, without the use of simulation tools. Based on the wireless performance numbers observed in the reported tests, it is expected that the degradation in production throughput will be maximum 0.01-0.41%, depending on the exact line configuration and technology chosen [15].

F.4.2 Control of Industrial Mobile Robots

For the second use case, the baseline for step (3) of the roadmap was demonstrated. In this initial exercise, the AMRs run a light-weight control communication algorithm with feedback to the cloud-edge fleet manager server; but most of its localization functions remain operating locally at the AMR. This mobile use case has been evaluated, for the moment, only over Wi-Fi 5 and 4G, considering a single controlled AMR, while roaming at default speed (maximum 0.8 m/s) around the multiple cells deployed in the industrial hall following pre-configured mobility patterns (based on way-point definitions).

For this mobile use case, the best control-loop performance was achieved over the pLTE configuration. As compared to the non-optimized Wi-Fi 5 case, 4G outperforms Wi-Fi at both the median (20 vs. 35 ms, respectively), and low percentiles (157 ms vs. 4.2 s, respectively). As the maximum latency tolerable by the control-loop in this use case was 1 s, the operation over Wi-Fi resulted in sudden interruptions in the normal operation of the robot during 0.14% of the time. By comparing the performance of the mobile use case and the static use case over pLTE and non-optimized Wi-Fi 5, it is possible to quantify the effect of mobility and handover management for both technologies, being the impact in the order of 1-117 ms for 4G and 4 ms-3.2 s for Wi-Fi. It is clear that, due to their operation in dedicated licensed spectrum and in-built scheduling and handover mechanisms, the cellular technologies (4G, 5G) offer a much more contained and deterministic control-loop latency with lower packet error rates, than Wi-Fi.

F.4. Performance of Industrial Wireless Use Cases




Fig. E3: Wireless performance results for the wireless control of industrial production [MES/static] and for the control of AMRs [AMR/mobile] industrial use cases over different communication technologies.

E.5 Conclusion

The factories of the future will be equipped with flexible manufacturing equipment enabling the mass production of highly customized products. In order to achieve the maximum level of flexibility, a complete transformation of the traditional sequential centralized production paradigm is needed. In this respect, we envision the swarm production, (non-linear de-centralized production), enabled by the integration of advanced wireless technologies, cloud-computing, and autonomous mobile robots, which can be made a reality by following a simple roadmap and implementing the different steps and associated use cases. We have put in practice some of the steps of the roadmap, and have successfully demonstrated the wireless control of industrial production, as well as the control of mobile autonomous robots in a dedicated industrial wireless research setup, the AAU 5G Smart Production Lab, considering multiple technologies such as Wi-Fi 5, Wi-Fi 6, 4G and 5G.

Acknowledgement

The authors would like to express their gratitude to Andreas E. Fink and Søren A. Markussen, M.Sc. students at AAU, for their contribution to the collection of the initial 5G and Wi-Fi 6 performance results.

References

- [1] B. Chen *et al.*, "Smart Factory of Industry 4.0: Key Technologies, Application Case, and Challenges," *IEEE Access*, vol. 6, pp. 6505-6519, 2017.
- [2] B. Holfeld *et al.*, "Wireless Communication for Factory Automation: An Opportunity for LTE and 5G Systems," *IEEE Commun. Mag.*, vol. 54, no. 6, pp. 36-43, 2016.
- [3] M. Wollschlaeger *et al.*, "The Future of Industrial Communication: Automation Networks in the Era of the Internet of Things and Industry 4.0," *IEEE Ind. Electron. Mag.*, vol. 11, no. 1, pp. 17-27, 2017.
- [4] J. Kruger *et al.*, "Innovative control of assembly systems and lines", *CIRP Annals - Manufacturing Technology*, no. 66, pp. 707-730, 2017.
- [5] 3GPP TR 22.804 V16.2.0, "Study on Communication for Automation in Vertical Domains (Release 16)", December 2018.
- [6] I. Rodriguez *et al.* "On the Design of a Wireless MES Solution for the Factories of the Future", in *Proc. GIoT*S, 2019.
- [7] T. Goldschmidt *et al.*, "Cloud-Based Control: A Multi-tenant, Horizontally Scalable Soft-PLC", in *Proc. IEEE CLOUD*, 2015.

References

- [8] H. Tang *et al.*, "A Reconfigurable Method for Intelligent Manufacturing Based on Industrial Cloud and Edge Intelligence", *IEEE Internet of Things J.*, vol. 7, no. 5, pp. 4248 - 4259, 2019.
- [9] E. Garcia *et al.*, "The evolution of robotics research", *IEEE Robot. & Autom. Mag.*, vol. 14, no. 1, pp. 90-103, 2007.
- [10] G. Hu *et al.*, "Cloud robotics: architecture, challenges and applications", *IEEE Network*, vol. 26, no. 3, pp. 21-28, 2012.
- [11] A. Vick *et al.*, "Control of robots and machine tools with an extended factory cloud", in *Proc. IEEE WFCS*, 2015.
- [12] A. Varghese *et al.*, "Wireless requirements and challenges in Industry 4.0", in *Proc. IC3I*, 2014.
- [13] G. Berardinelli *et al.*, "Beyond 5G Wireless IRT for Industry 4.0: Design Principles and Spectrum Aspects," in *Proc. IEEE GLOBECOM Wksp*s, 2018.
- [14] R. S. Mogensen *et al.*, "Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0", in *Proc. IEEE VTC-Fall*, 2019.
- [15] R.S. Mogensen *et al.*, "Evaluation of the Impact of Wireless Communication in Production via Factory Digital Twins", [to-be-published] *Manufacturing Letters*, 2020.

References

Paper G

Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications

R. S. Mogensen, C. Markmøller, T. K. Madsen,
T. Kolding, G. Pocovi, and M. Lauridsen.

The paper has been published in the
2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring)

© 2019 IEEE

The layout has been revised.

Abstract

In this paper, a new multi-connectivity protocol is introduced for mission critical applications that use wireless 5G multi-UE devices. It extends the MP-QUIC transport protocol by adding selective data duplication with strict prioritization to support a mix of critical and non-critical applications on a single end-to-end connection; e.g. suitable for self-driving cars getting mission critical drive control commands from the network (priority) while receiving map-updates (background). The Selective Redundant MP-QUIC solution simultaneously provides bandwidth aggregation and higher speeds for background data while ensuring maximum reliability and availability for the priority data by duplicating it on the available paths. The proposed algorithm is first evaluated using an emulated fixed-delay channel as well as a well-known Internet connection model. Next, a highway drive test is conducted using two LTE modems, resembling a typical autonomous car scenario. The tests show that the algorithm can protect the critical data flow also in presence of a background load of 2 Mb/s and across a wide range of channel conditions. Using our proposed MP-QUIC solution over two LTE connections, the 99.9%-tile of critical traffic latency shows a 5 times improvement versus using a single-path connection and 3 times versus using state of the art MP-QUIC protocols.

G.1 Introduction

5G wireless communication is poised to add tremendous value in everything from autonomous vehicles such as cars and drones to Industry 4.0. To support mission critical systems with very high reliability and availability requirements, multi-modem devices are studied as part of 3GPP 5G Release 16 [1]. Having multiple modems improve the system's reliability and availability. In Fig. G.1, one possible instantiation is shown where the device contains two separate user equipment (UE) that are connected via two different networks towards the remote side of the service (here marked by data network DN), e.g. connected through a separate 5G base station (gNB) and 5G core network user-plane function (UPF). Transport algorithms that combine multiple connections may be located within the 5G network (offered as a service by the operator) or outside at the transport or application layers and be transparent to each of the 5G networks. Redundancy may be limited to the radio layer (gNB) by various multi-connectivity means or also include full UPF redundancy as shown in Fig. G.1.

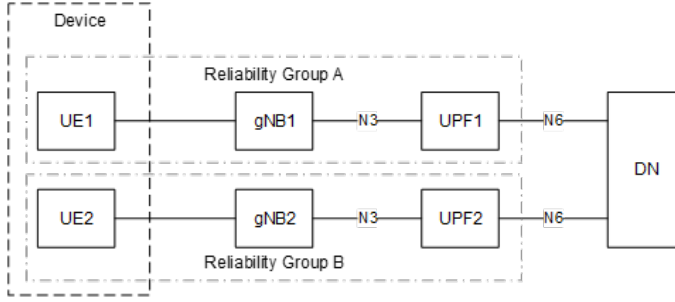


Fig. G.1: 5G high-reliability architecture with multi-UE device [1]

Solutions for access aggregation is a widely researched topic and can be realized in different layers of network stack where flexibility and availability generally decreases when conducted at the lower layers of the network stack [2], [3]. Hence, multi-connectivity on the network or transport layer is beneficial if a mix of different access technologies are used or if sub-networks come from different domains. Of these two, a transport layer solution is the most flexible as this is agnostic to the network provider as well [4]. The benefit of multi-modem devices with transport layer multi-connectivity solution in wireless environments has been well documented, e.g. for LTE networks along Danish freeways in [5] and for German high speed trains in [6]. By mitigating the outage related to e.g. handovers and coverage gaps, the reliability of achieving a certain low delay packet delivery is increased significantly. Common for many of the applications in need of multi-UE devices is that they simultaneously produce or consume multiple types of information, where some require low latency and very high reliability whereas others require higher bandwidth but relaxed latency/reliability requirements [7]. For instance, a common connection to a self-driving car may have critical drive control commands and background map updates that have very different latency and reliability requirements. The most relevant multi-connectivity transport layer solutions include Multipath Transport Control Protocol (MP-TCP), Multipath Quick User Datagram Protocol Internet Connection (MP-QUIC) and the Stream Control Transport Protocol (SCTP), where the latter is often excluded as it exhibits low compatibility with middleboxes [8]. MP-TCP [9] is a multipath protocol that utilizes multiple TCP connections which carry traffic according to different schemes. For increased reliability, the Redundant Scheduler can be used to send duplicated packets on all available TCP connections [6]. The drawback of MP-TCP is however the lack a proper prioritization mechanism when handling of multi-streamed applications which limits its scope of applicability [10]. MP-QUIC [11] on the other hand utilizes multiple QUIC connections [12]. QUIC is a reliable transport protocol based on UDP which provides similar functionality to TCP with the extension of data multiplexing

and prioritization. While still in early development stages, MP-QUIC provides a framework that enables multiple application data sources as well as prioritization of these sources, inherited from QUIC, but does not support multi-path redundancy [11], [13]. State-of-the-art multi-connectivity solutions therefore lack good support for a mix of (a) low latency reliable traffic that requires data duplication across multiple interfaces and (b) simultaneous lower priority traffic. We therefore propose a new Selective Redundant MP-QUIC concept which can provide low latency and full redundancy for priority application streams and high aggregated bandwidth for background application streams in the same connection. To test the performance of the proposed algorithm, the developed software is deployed on a client and a server connected with multiple parallel Ethernet connections. A network emulation tool is used to reproduce different delay and packet loss models on each connection; including an IEEE/IETF Internet delay model, and LTE delay traces from a highway drive test. The rest of the paper is structured as follows. Section G.2 introduces the Selective Redundant MP-QUIC concept and reference concepts used for benchmarking. Section G.3 contains an overview of the methodology used to evaluate the new scheduling algorithm. Section G.4 contains the results analysis followed by conclusions in Section G.5.

G.2 Selective Redundant MP-QUIC

In Figure G.2, the basic concept behind Selective Redundant MP-QUIC (SR-MPQUIC) is illustrated. Note that for simplicity only two UE connections are described and illustrated, but the algorithm supports more parallel connections. The algorithm uses bandwidth aggregation utilizing both UE connections to improve the data rate of the background traffic. To protect the priority traffic, data is duplicated on both UE connections and a new congestion algorithm ensures that the priority data flow is exempt from being stalled even if total bandwidth of priority and background data momentarily exceeds the network capacity.

The MP-QUIC implementation in [11] is used as baseline. Special focus is put on the path scheduler responsible for scheduling and the stream framer responsible for packet creation. The path scheduler selects the path for packet transmission using either Shortest round-trip time first (SRTT) or Round Robin (RR) [14], whereas the stream framer selects frames from the different application streams using a prioritized or RR scheme. The reference scheme used for benchmarking is the SRTT MP-QUIC with prioritization, i.e. packets are always filled with frames from the priority traffic stream before frames from the background traffic stream. To implement redundant transmission, a new stream attribute is introduced that indicates whether data from a stream should be duplicated. This attribute is used whenever scheduling or packet




Fig. G.2: Selective Redundant scheduling concept with two available UEs/paths and two distinct application data sources.

creation is performed. For each scheduling cycle, the path scheduler will check which kind of data the scheduled packet shall contain depending on the duplication attribute of the streams. If one stream in the current scheduling instance has this set, it will schedule a packet transmission on the N available paths. When creating the packet on a path the stream framer checks for this attribute and ensures that the frame is duplicated across the N packets that have been scheduled for transmission. If no stream has the duplication attribute in a scheduling instance, the regular SRTT scheduler is used. In Section G.4, the benefit of SR-MPQUIC over an implementation that does duplication of all packets from all streams (i.e. both priority and background data) is also analyzed. This latter method is referred to as Fully Redundant MP-QUIC (FR-MPQUIC), and can be considered a counter-part to the MP-TCP Redundant Scheduler [6].

G.2.1 Stream aware congestion algorithm (SACA)

Prioritization of priority frames during packet creation process does not guarantee priority in actual packet transmission, e.g. if (a) a path is congested or (b) there is a pending retransmission that per default is selected first. The latter is an intentional impact, while the former results in a potential latency spikes for a priority data stream in case the congestion window (CW) is used completely by packets containing background data frames. In other words, buffered background data may prevent new packets containing priority data frames to be immediately sent to the lower layers. To mitigate this, a stream aware congestion algorithm (SACA) is introduced to ensure that priority data is not impacted by packet scheduling limitations imposed by the congestion algorithm while maintaining the long-term desired effect. An illustrative example of the behavior of the algorithm is shown in Fig. G.3 From time instances t_1 to t_3 , the background data (BD) is getting scheduled according to the limits estimated by the CW algorithm. At time instance t_3 , data from the




Fig. G.3: Example of stream aware congestion algorithm (SACA).

priority traffic (PD) also arrives and the limit is exceeded but still logged in the bytes-in-flight counter. To protect the basic CW scheme, the BD is lowered by the exceeded amount (in t3) during the time instance t4. This ensures that priority traffic is exempt from the CW boundary momentarily, which is immediately compensated by throttling the BD on the next scheduling interval.

For convenience, the different evaluated multipath algorithms are summarized in Table G.1.

Table G.1: Multipath algorithms considered.

Abbreviation	Method
SR-TT-MPQUIC	MP-QUIC using Shortest RTT to aggregate connections (no redundancy).
FR-MPQUIC	MP-QUIC algorithm that uses Redundant Scheduler for all data flows and uses new stream aware congestion algorithm
SR-MPQUIC	MP-QUIC algorithm that allows selective redundant scheduling and uses new stream aware congestion algorithm

G.3 Methodology and Test Setup

To test the performance of the proposed algorithms and reference schemes, the developed software is deployed on a client and a server computer. The used equipment and software is listed in Table G.2. Both client and server are equipped with multiple network interface cards (NICs), three on the client and two on the server, which are used to create two separate IP networks

via Ethernet switches. One network is used for time synchronization and management (single-path i.e. one NIC for the client and server), and the other for transmitting the priority and background streams (multi-path i.e. two NICs for the client and one for the server). Each packet generated at the application layer contains a time-stamp which is used to measure the one-way delay (OWD) at the receiver side. Synchronization is needed to measure OWD accurately and the approach, among other aspects of the measurement framework, is documented in [15].

Table G.2: Hardware/software for network/algorithm emulation

HW/SW	Details
PCs	Dell Optiplex 9020/7010 (server/client) i7, 16 GB RAM
OS	Ubuntu 18.04 LTS
NIC model	Intel 82574L, 1Gb/s
Emulator	NetEm [16]
Switches, cables	Zytek 1Gb/s, Cat 5e

Throughout the tests, the priority data (PD) is modelled as a periodic traffic flow with a 1200-byte payload and a packet inter-arrival time of 100 ms (corresponding to an effective data rate of 96kb/s). This model is inspired by the 3GPP 5G data model assumed for autonomous vehicles in [17]. The default background data (BD) rate is set to either 0 Mb/s or 2 Mb/s. Our target reliability range is 99.9% [18] and the number of emulated packets has been set correspondingly to at least 10k samples per run. A network emulation tool, NetEm [16], is used to reproduce different delay and packet loss conditions to packets egressing from each network interface. There is no bandwidth limitation imposed on the underlying network which means that throughput is governed by the congestion algorithm of the adopted MPQUIC implementation. Table G.3 shows the three considered network scenarios with an increasing level of realism. The first scenario has a static emulated network latency where focus is on the algorithm's ability to tackle different packet loss probabilities. In the second scenario, packet losses are disabled and sensitivity of variable delay effects is studied using the IEEE wired domestic networks model [19]. Finally, performance under packet delays typical for a wireless deployment is considered, using drive test measurements obtained from a Danish highway trial with multiple LTE connections [15]. Since the implementation only supports multihomed clients and latencies can only be emulated on the egress using NetEm, only the client to server links are modeled according to the different models. As MP-QUIC relies on feedback mechanisms, the reverse link (Server-Client) is modeled with a fixed delay of 50 ms such that instantaneous feedback is not possible. Although this asymmetric condition does not reflect reality, it still serves as an appropriate

emulation methodology as only the latency for the traffic flowing from the client to the server is considered. Also, statistics such as losses and RTTs that are relevant for the congestion algorithm and path scheduler will still be affected even if only the network conditions on the client to server links are changed as these statistics are independent of the flow direction from a client's perspective. Hence, only one-way delay results for the client to server link are presented.

Table G.3: The three network test scenarios.

Setting	Server-Client	Client-Server
A: Static delay, Variable packet loss		
Latency	50 ms	50 ms
Packet loss	0 %	0-1 %
B: Wired Domestic IEEE Internet model		
Latency	50 ms	IEEE model
Packet loss	0 %	0 %
C: Wireless model - LTE highway drive test		
Latency	50 ms	Realistic delay traces
Packet loss	0 %	0 %

G.4 Results

First, results are presented to verify the functionality of our stream aware congestion algorithm (SACA). To isolate the effects from the multipath components, a single-path QUIC connection is setup. To stress the system both 96 kb/s PD and 10 Mb/s BD is sent. The fixed 50 ms delay model is used and both 0% and 1% packet loss probabilities are emulated. The single-path QUIC components are adapted from [11] where SACA has been added and is either enabled or disabled during the test. The complementary cumulative distribution function (CCDF) of the achieved latency of the priority traffic is shown in Figure G.4 As expected, there is no impact visible in the 0% loss case and the latency therefore correspond to 50 ms including variations from processing and jitter. For 1% loss probability, it is observed how SACA improves the achieved latency distribution, i.e. priority packets experience a generally better latency performance. Due to the emulated 1% packet loss, even the priority data nevertheless experience a packet delay tail caused by head-of-line-blocking due to retransmission delay [20]. Although not shown, when lowering the BD to 2 Mb/s the benefits of SACA are less apparent, as the link is not constantly congested, but still present. In the following, we attempt to reduce the tails of the distribution by introducing the multipath component.




Fig. G.4: CCDF of the priority-data one-way latency achieved with a single-path QUIC connection with and without the proposed SACA algorithm and different packet loss probabilities (BD 10 Mb/s).

G.4.1 Static Delay, Variable Packet Loss

Next, multipath performance using two connections is considered, see Table G.1 and Table G.3. Tests are done with BD ranging from 0 Mb/s to 2 Mb/s as higher BD rates exhibit constant congestion under the presence of packet losses. The latency CCDF for the priority traffic using two network interfaces is shown in Figure G.5. As with the single path QUIC in Figure G.4 there is no significant difference between the different schedulers for 0% loss. For 1% loss probability, the single interface reference is showing similar performance to the SRTT-MPQUIC as this only leverages the interfaces for aggregation and not redundancy, and therefore does not prevent the tail imposed by the packet losses. In contrast, the two redundant schedulers, FR-MPQUIC and SR-MPQUIC are effective at avoiding packet losses by simultaneously leveraging the two interfaces. The priority data performance at the 99.9% level is identical for the two, but it should be recalled that FR-MPQUIC uses nearly double bandwidth to achieve the reliability as will be discussed later.

Figure G.6 summarizes the one-way latency of the priority data at the 99.9% level for different packet loss probabilities. It is shown that the FR-MPQUIC and SR-MPQUIC gains over SRTT-MPQUIC are consistent and appear even at very low loss rates of 0.1%. The latency improvement varies between 50 ms (at 0.1% loss rate) and 200 ms (at 1 % loss rate). Unexpectedly, SRTT-MPQUIC delay performance improves slightly under the presence of BD which we suspect is emulator related as the trend does not reoccur in other tests.

G.4. Results




Fig. G.5: Priority data latency CCDF, BD: 2 Mb/s, 2 interfaces (besides single path reference).




Fig. G.6: 99.9%-ile of the priority data latency for different packet loss probabilities, BD rates, and schedulers.

As seen in Figure G.6 the latency performance (e.g. the 99.9%-tile of the priority data latency) is similar between the FR-MPQUIC and SR-MPQUIC variants. To show the benefit of the SR-MPQUIC algorithm, the bandwidth efficiency of the different schedulers is compared in Table G.4. Bandwidth efficiency is defined as the ratio between goodput and throughput. FR-MPQUIC has a bandwidth efficiency of only 48% as it requires double the amount of bandwidth compared to SR-MPQUIC which has a bandwidth efficiency comparable to SRTT-MPQUIC (<4% point degradation).

Table G.4: Bandwidth efficiency 2 Mb/s BD, 0.1 % packet loss.

MP-QUIC	SR	SRTT	FR
Goodput	2.07 Mb/s	2.07 Mb/s	2.04 Mb/s
Throughput	2.25 Mb/s	2.16 Mb/s	4.27 Mb/s
Bandwidth Efficiency	92 %	96 %	48 %

G.4.2 Wired Domestic IEEE Internet model

Moving to the variable IEEE domestic wired internet delay model [19], Figure G.7 compares the performance of the different algorithms. In this model the connections are loss-free and have an uncorrelated gamma distributed delay with a minimum value of 7.5 ms. As a reference, we also include the ideal performance (i.e. the minimum latency achievable at every time instance) of two links using the IEEE model. It is observed that our system and protocols add some processing overhead of around 1-2 ms. However, the results show that the FR-MPQUIC and SR-QUIC methods are very robust to the amount of background data (BD: 0 Mb/s or 2 Mb/s) and perform very similarly to the ideal curve. Furthermore, they offer a significant latency reduction compared to SRTT-MPQUIC, which have similar latency to a single IEEE model link, from around 25 ms to 16-17 ms at the 10^{-3} -th percentile.

G.4.3 Wireless model - LTE highway drive test

Finally, results for the highway drive test are shown in Figure G.8 The BD rate has been reduced from 2 Mb/s to 1 Mb/s in order to operate closer to the bandwidth limit of the LTE network (limited at cell edge, high speed). For reference, the figure shows the LTE latencies of the two tested paths (Link 1 and 2) along with the resulting scheduler performances. The high-speed wireless environment poses significant challenges for mission critical services, with experienced 99.9% latency outage around 1 second for a single LTE link. SRTT-MPQUIC, prioritizing paths with shortest delay, can improve the performance with two interfaces, reducing the latency outage to around 700

G.4. Results




Fig. G.7: CCDF of priority data latency according to the IEEE wired domestic delay model.




Fig. G.8: Results for Highway Drive Test model, 1 Mb/s BD rate.

ms. As shown in previous results, FR-MPQUIC and SR-MPQUIC with two interfaces provide further latency improvement, e.g. the latency is reduced to 200 ms at the 10^{-3} -th percentile, a nearly 5x improvement as compared to the individual-link performance. Despite the obtained gains, there is still some gap to the final application requirements of 100 ms [17], meaning that at least 3-4 parallel LTE connections would be needed as shown in [5]. For 5G (or further optimized LTE for autonomous driving), significant latency improvements are expected suggesting that the proposed SR-MPQUIC should be able to achieve the requirements with only two interfaces.

For wireless deployments, the bandwidth efficiency gains of SR-MPQUIC over FR-MPQUIC are of high importance as the bandwidth aggregation are needed to ensure sufficiently high data rates at the cell edge/borders. While SRTT-MPQUIC or SR-MPQUIC only requires 1.1 Mb/s to be transmitted over the air interface, FR-MPQUIC requires 2 Mb/s which may be challenging to reliably achieve on the LTE network.

G.5 Conclusions

We introduced a new Selective Redundant MP-QUIC scheme suited for multi-UE connectivity solutions intended for wireless mission critical services in 5G era. Within a single connection, e.g. from network to an autonomous vehicle, the solution allows for selectively duplicating critical data flows such as drive control commands over multiple connections while simultaneously providing bandwidth aggregation benefits for normal priority data such as map updates etc. Compared to state-of-the-art, the new solution combines the reliability benefits of redundant MP-TCP with the applicability and efficiency benefits of SRTT MP-QUIC. The solution introduces two novel components, a stream priority-based scheduler and a new prioritized congestion algorithm. The new solution has been tested in different scenarios and is robust towards both packet losses and delay variations making it suitable for the harsh wireless environments. It achieves near-ideal reliability and same performance as a fully redundant scheme, but at a much better bandwidth efficiency supporting better the background data. Using drive test data, the new scheme outperforms the current SRTT MP-QUIC solution having three times reduction of the 99.9% guaranteed packet delay (from more than 700 ms to 220 ms) with less than 1 % reduction in overall bandwidth efficiency. Compared to a single-UE connection, the 99.9 % experienced data delay was reduced by nearly 5 times, showing significant promise of leveraging multi-UE techniques for wide area mission critical services.

References

- [1] 3GPP, "Study on enhancement of Ultra-Reliable Low-Latency Communication (URLLC) support in the 5g Core network (5gc)," Tech. Rep. TR 23.725, 2019.
- [2] A. Ravanshid, P. Rost, D. S. Michalopoulos, V. V. Phan, H. Bakker, D. Aziz, S. Tayade, H. D. Schotten, S. Wong, and O. Holland, "Multi-connectivity functional architectures in 5g," in *IEEE ICC Workshop*, May 2016, pp. 187–192.
- [3] J. J. Nielsen, R. Liu, and P. Popovski, "Optimized Interface Diversity for Ultra-Reliable Low Latency Communication (URLLC)," in *IEEE GLOBECOM 2017*, pp. 1–6.
- [4] C. Paasch, G. Detal, F. Duchene, C. Raiciu, and O. Bonaventure, "Exploring Mobile/WiFi Handover with Multipath TCP," in *ACM SIGCOMM Cellnet'12*, 2012.
- [5] M. Lauridsen, T. Kolding, G. Pocovi, and P. Mogensen, "Reducing handover outage for autonomous vehicles with lte hybrid access," in *2018 IEEE International Conference on Communications (ICC)*, 2018, pp. 1–6.
- [6] A. Frommgen, T. Erbschäuffer, A. Buchmann, T. Zimmermann, and K. Wehrle, "ReMP TCP: Low latency multipath TCP," in *2016 IEEE ICC*, pp. 1–7.
- [7] G. Yang, X. Lin, Y. Li, H. Cui, M. Xu, D. Wu, H. Ryden, and S. B. Redhwan, "A Telecom Perspective on the Internet of Drones: From LTE-Advanced to 5g," *CoRR*, vol. abs/1803.11048, 2018. [Online]. Available: <http://arxiv.org/abs/1803.11048>
- [8] A. Joseph, T. Li, Z. He, Y. Cui, and L. Zhang, "A Comparison between SCTP and QUIC," IETF, Internet-Draft draft-joseph-quic-comparison-quic-sctp-00, Mar. 2018, published: Working Draft. [Online]. Available: <http://www.ietf.org/internet-drafts/draft-joseph-quic-comparison-quic-sctp-00.txt>
- [9] C. Paasch and O. Bonaventure, "Multipath TCP," *Queue*, vol. 12, no. 2, pp. 40:40–40:51, Feb. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2578508.2591369>
- [10] F. Gont and A. Yourtchenko, "On the Implementation of the TCP Urgent Mechanism," RFC 6093, published: Internet Requests for Comments.
- [11] Q. D. Coninck and O. Bonaventure, "Multipath QUIC: Design and Evaluation," in *Conext'17*, Dec. 2017.
- [12] J. Iyengar and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport," IETF, Internet-Draft draft-ietf-quic-transport-14. [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-ietf-quic-transport-14>
- [13] T. Viernickel, A. Froemmgen, A. Rizk, B. Koldehofe, and R. Steinmetz, "Multipath QUIC: A Deployable Multipath Transport Protocol," in *2018 IEEE ICC*, pp. 1–7.
- [14] C. Paasch, S. Ferlin, O. Alay, and O. Bonaventure, "Experimental evaluation of multipath TCP schedulers," in *SIGCOMM CSWS '14*. ACM Press, pp. 27–32. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2630088.2631977>
- [15] G. Pocovi, T. Kolding, M. Lauridsen, R. Mogensen, C. Markmoller, and R. Jess-Williams, "Measurement Framework for Assessing Reliable Real-Time Capabilities of Wireless Networks," *IEEE Communications Magazine*, pp. 1–8, 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8450876/>

References

- [16] S. Hemminger, "Network Emulation with NetEm," p. 9.
- [17] 3GPP, "Service requirements for V2x services; Stage 1 (Release 15)," Tech. Rep. TS 22.185 V15.0.0 (2018-06), Jun.
- [18] —, "Study on Enhanced LTE Support for Aerial Vehicles (Release 15)," Tech. Rep. TR 36.777 V0.3.1 (2017-10).
- [19] Andrew Corlett, D. I. Pullin, and Stephen Sargood, "Statistics of One-Way Internet Packet Delays," IETF, Internet-Draft. [Online]. Available: <https://tools.ietf.org/html/draft-corlett-statistics-of-packet-delays-00>
- [20] M. Scharf and S. Kiesel, "NXG03-5: Head-of-line Blocking in TCP and SCTP: Analysis and Measurements."

Paper H

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

A. Fink, R. S. Mogensen, I. Rodriguez,
T. Kolding, A. Karstensen, and G. Pocovi.

The paper has been published in the
2022 IEEE Wireless Communications and Networking Conference (WCNC)

© 2022 IEEE

The layout has been revised.

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.

H.1 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 H.2 details our Wi-Fi multi-connectivity solutions. Section

H.3 introduces the experimental environment, the setups, and the different radio configurations tested for the multiple multi-connectivity schemes. Section H.4 presents the latency performance measurements results. Section H.5 contains the discussion of the results and highlights areas of potential improvements. Finally, section H.6 concludes the paper.

H.2 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. H.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.




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

H.2.1 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.

H.2.2 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. H.2 and H.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.

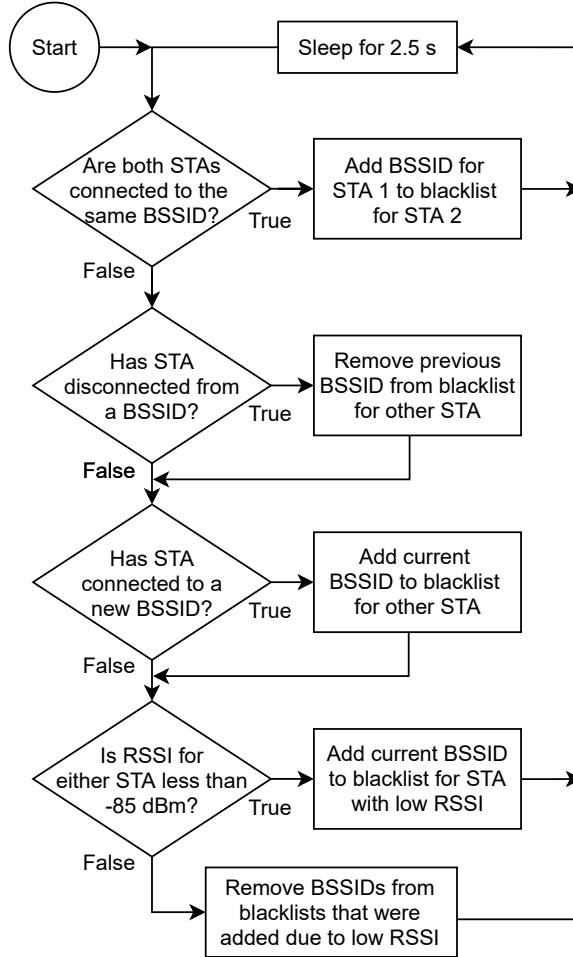


Fig. H.2: Flowchart for the packet duplication scheme.

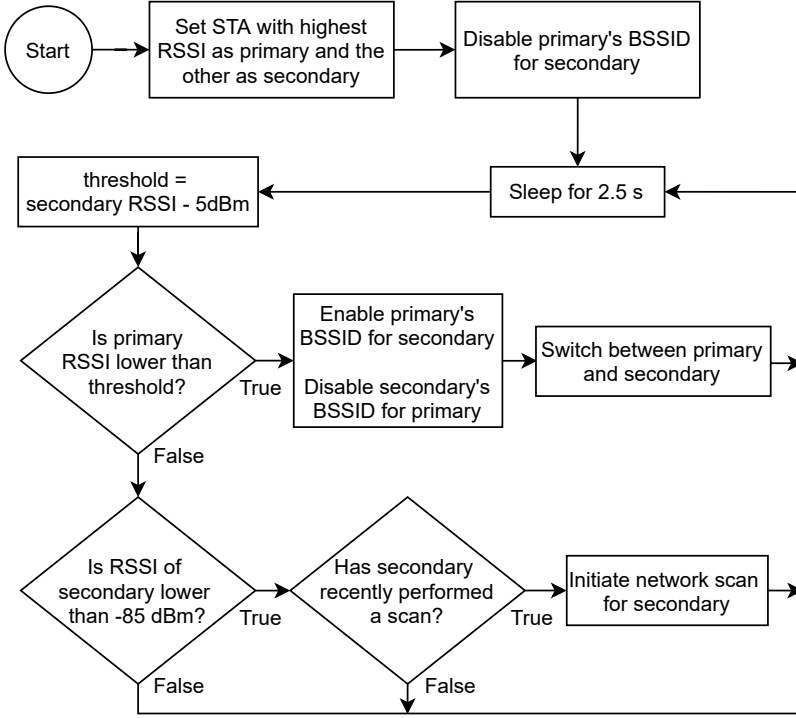


Fig. H.3: Flowchart for the best path scheduling scheme.

H.2.3 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.

H.3 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. H.4) is equipped with three ceiling-mounted CISCO MR36 Enterprise Wi-

H.3. Experimental Test Setup

Fi 6 APs [15] deployed throughout the lab as illustrated in Fig. H.5. In order to trigger the mobility aspects of our Wi-Fi solution performance evaluation, a MiR200 AMR (also shown in Fig. H.4) was used. The AMR was configured to follow a specific route through the lab as illustrated in Fig. H.5, carrying the implemented multi-STA device around at a maximum speed of 1.5 m/s.




Fig. H.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. H.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).

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. H.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.

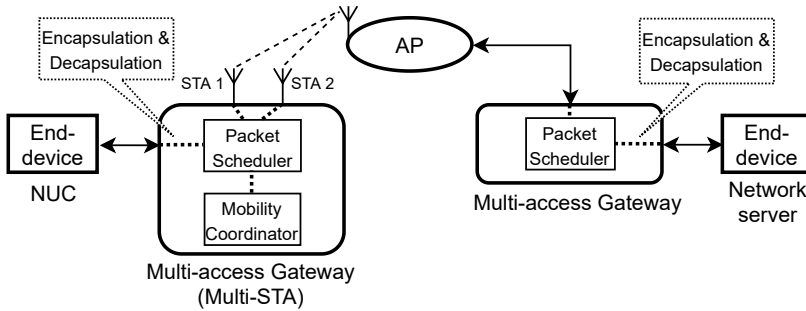


Fig. H.6: Test setup using multi-access gateways. Solid connections represent Ethernet connections, while dashed connections represent Wi-Fi.

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 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), 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 ~ -93 dBm).

H.4 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. H.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 constraints 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. H.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 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.

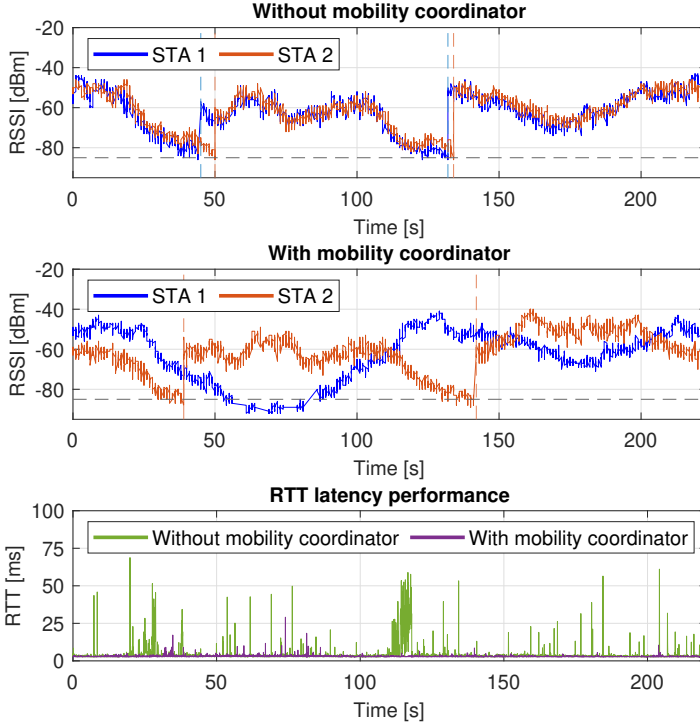


Fig. H.7: RSSI traces and overall Wi-Fi RTT latencies for two different PD schemes: with and without mobile coordinator. The -85 dBm threshold is highlighted with the horizontal line, and handovers are highlighted with vertical lines.

Fig. H.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 H.1. As a reference in both Fig. H.8 and Table H.1, 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.




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

Table H.1: Summary of RTT latency statistics and PER measurement results for the single- and multi-connectivity configurations operating over dedicated channels.

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

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. H.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.




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

H.5 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 devices, 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. H.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.

H.6 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, R. Mogensen, I. Rodriguez, T. E. Kolding, A. Karstensen, and G. Pocovi, "Empirical Performance Evaluation of Enterprise Wi-Fi for IIoT Applications Requiring Mobility," unpublished, submitted to European Wireless 2021.
- [2] E. J. Khatib, D. A. Wassie, G. Berardinelli, I. Rodriguez, and P. Mogensen, "Multi-Connectivity for Ultra-Reliable Communication in Industrial Scenarios," in *IEEE Vehicular Technology Conference (VTC2019-Spring)*, 2019, pp. 1–6.
- [3] M.-T. Suer, C. Thein, H. Tchouankem, and L. Wolf, "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, I. Leyva-Mayorga, and P. Popovski, "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, T. Kolding, G. Pocovi, and P. Mogensen, "Reducing Handover Outage for Autonomous Vehicles with LTE Hybrid Access," in *IEEE International Conference on Communications (ICC)*, 2018, pp. 1–6.
- [6] C. Raiciu, C. Paasch, S. Barre, A. Ford, M. Honda, F. Duchene, O. Bonaventure, and M. Handley, "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, D. Niculescu, and C. Raiciu, "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, Yusheng Ji, and S. Yamada, "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 and K. Chen, "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, C. Markmoller, T. K. Madsen, T. Kolding, G. Pocovi, and M. Lauridsen, "Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications," in *IEEE Vehicular Technology Conference (VTC2019-Spring)*, 2019, pp. 1–5.
- [11] H. Zhang, A. Elmokashfi, and P. Mohapatra, "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, I. Rodriguez, G. Berardinelli, A. Fink, R. Marcker, S. Markussen, T. Raunholt, T. Kolding, G. Pocovi, and S. Barbera, "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, C. Thein, H. Tchouankem, and L. Wolf, "Evaluation of Multi-Connectivity Schemes for URLLC Traffic over WiFi and LTE," in *IEEE Wireless Communications and Networking Conference (WCNC)*, 2020, pp. 1–7.

References

- [14] I. Rodriguez, R. S. Mogensen, A. Schjørring, M. Razzaghpour, R. Maldonado, G. Berardinelli, R. Adeogun, P. H. Christensen, P. Mogensen, O. Madsen, C. Møller, G. Pocovi, T. Kolding, C. Rosa, B. Jørgensen, and S. Barbera, "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/

Paper I

A Novel QoS-Aware Multi-Connectivity Scheme for Wireless IIoT

R. S. Mogensen, S. B. Damsgaard, I. Rodriguez, G. Berardinelli,
A. Fink, T. Kolding, and G. Pocovi.

The paper has been published in the
IEEE Access, vol. 10, pp. 104123-104134, October 2022.

© 2022 IEEE

The layout has been revised.

Abstract

Wireless technology is envisioned to be a major enabler of flexible industrial deployments, allowing agile installations and mobility of production elements. However, for industrial IoT (IIoT) use-cases, reliability and service availability remain as key concerns for wide adoption. To enhance wireless link reliability, we propose in this paper a Selective Duplication with QoS (SDQoS) technique, a cross-layer scheduling solution which leverages radio access technology metrics and transport layer metrics to schedule data transmission across one or more links. The framework is implemented as part of a multi-access gateway solution and its capability is demonstrated in a realistic two-hall industrial production environment for a multi-data flow autonomous mobile robot use-case, using Wi-Fi. The proposed QoS-aware solution shows a major improvement in preserving low latency and reliability for critical control data in the presence of large background traffic as compared to state-of-the-art solutions.

I.1 Introduction

The rapid development of the Industrial Internet of Things (IIoT) has led to a diverse set of new technological enablers and concepts that promise to improve flexibility, efficiency, and reliability of production setups in factories [1]. Among them, industrial mobile applications such as Autonomous Mobile Robots (AMRs), or cloud computing, which enables optimized control in a factory by migrating the ‘intelligence’ of Programmable Logic Controllers (PLCs) from the production shop floor to a centralized location, are becoming a reality [2]. The different industrial use cases and applications built on top of these new technological enablers, rely on the integration of communications, robotics, and edge-cloud; and are based on the exchange of data flows which, typically, exhibit a wide range of requirements in terms of e.g., latency, reliability, and throughput [3]. The industrial evolution has also motivated the parallel development of the industrial communication ecosystem which, in general, has evolved from the use of field-buses to Ethernet as the main control communication technology [4]. In this respect, wireless communication is now recognized as a major enabler of flexible industrial deployments, as it allows for agile installations, cable replacement, and mobility of production elements [5]. Here Wi-Fi and 5G New Radio (NR) are identified as main candidates for deployment in industrial scenarios and, while both support Industrial Internet-of-Things (IIoT) applications, in particular, effective reliability and service availability remain as key concerns [6].

As a potential solution, multi-connectivity is identified as an excellent candidate technique to enable the support of reliable communication in case of variable radio channel conditions [7], e.g. due to a harsh radio environment (as it is the case in industrial environments), or due to mobility, while eventually

preserving sufficiently high spectral efficiency. There are multiple approaches to multi-connectivity, which can be classified into two groups depending on the layer of the protocol stack in which the functionality is implemented: 1) technology specific (lower layers), or 2) network-agnostic (higher layers).

For 5G NR there is native multi-connectivity support at different layers. At the Physical Layer (PHY), NR supports non-coherent joint transmissions from Multiple Transmit and Receive Points (multi-TRP), while duplication is also possible at the Packet Data Convergence Protocol (PDCP) layer to utilize both a primary and secondary base station for transmission and reception [8]. Additionally, Dual Active Protocol Stack (DAPS) provides functionality to enable make-before-break handover [9]. A common characteristic of these approaches is that they require standard-compliant implementation of the multi-connectivity solution at both the User Equipment (UE) device side and network side [10]. Such features are generally not available in existing 5G NR networks as initial 5G NR deployments are mainly targeting traditional Mobile Broadband (MBB) applications. However, in future 5G NR industrial deployments, these specific multi-connectivity features could be enabled - provided that both network and user equipment (UE) vendors implement them.

For Wi-Fi there is no native multi-connectivity support, however there do exist proprietary technology-specific solutions which can provide make-before-break handover by utilizing multiple radio interfaces and Access Points (APs) in coordination, as well as improved reliability by employing a network controlled Time Domain Multiple Access (TDMA) scheduling scheme as a collision avoidance mechanism [11]. However, this requires specialized devices both at the network and device side and is, therefore, not well suited for environments requiring technology co-existence and multi-vendor interoperability.

Compared to the above techniques, which require network/device side specific features, network-agnostic multi-connectivity is performed at the network/transport layer. Therefore, Radio Access Technologies (RATs), such as 5G NR or Wi-Fi, are simply viewed as communication links which provide a data service with certain characteristics in terms of latency, bandwidth and reliability. This abstraction has multiple benefits, including fast time-to-market and the ability to combine multiple RATs easily. In this regard, this study elaborates on the network-agnostic multi-connectivity concept where multiple communication interfaces are exploited to enhance the communication quality in industrial environments, especially in mobility conditions. Previous works have highlighted the benefit of either duplicating packets or selecting the best interface for transmission [12]. However, quality of service (QoS) was disregarded, i.e. every packet was treated the same regardless of the criticality of the supported application. Multi-flow QoS differentiation on the transport layer has previously been investigated by the authors of [13] and our previous work in [14] for Quick User Datagram Protocol (QUIC) and

multi path-QUIC (MPQUIC). Such work illustrates the impact of leveraging the transport layer statistics such as congestion window estimates to correctly allocate sufficient resources to different flows using the same underlying communication channels. However, published solutions do not consider metrics of the underline radio technology, possibly leading to sub-optimal design and performance. This is especially important in e.g., Wi-Fi where the mobility is handled at the device side, as shown in our previous work [15].

In this paper, we propose a novel adaptive multi-connectivity scheme which exploits specific radio metrics of the underlining RATs to improve QoS performance at run-time. Besides looking at suitable metrics to improve the performance of the cross-layer hybrid multi-connectivity solution, we validate its performance by experimentation in a realistic IIoT scenario based on a real industrial use case: AMR with two different data flows (control link and video link), with different QoS requirements. For the experimental validation, a specific implementation of the novel multi-connectivity scheme based on Wi-Fi was done. However, it is important to note that the proposed connectivity solution is fully agnostic to the specific underline RATs, and can also be applied for multi-connectivity over other RATs, such as 5G NR, or across multiple different RATs. Further, in the considered implementation, we leverage QUIC as the transport layer protocol, as it already possesses a mechanism that allows for easy differentiation of traffic flows as well as for providing transport layer statistics. However, the presented multi-connectivity scheme is flexible enough to allow for a different choice of a different transport protocol, if required. As such the contributions of the paper are:

- Design of a novel two-step cross-layer link selection scheme, Selective Duplication with QoS (SDQoS), for network-agnostic multi-connectivity, aiming improving reliability of selected data links, based on specific QoS rules, especially in mobility scenarios.
- Implementation of the SDQoS a scheme into an industrial application-agnostic multi-access device with Ethernet traffic support, using Wi-Fi as the baseline RAT.
- Performance evaluation and validation of the SDQoS scheme operation in a real-world factory test environment considering a real industrial AMR use case and traffic flows.
- Comparison of the performance of SDQoS with respect to other state-of-the-art connectivity solutions, and discussion of the achieved gains.

The rest paper is structured as follows. Section I.2 presents the algorithms for the proposed novel cross-layer multi-connectivity link selection scheme, and its implementation-specific details considering an IIoT Ethernet-based gateway and Wi-Fi 6 as wireless technology, which are used in the subsequent

performance evaluations. Section I.3 describes the experimental setup and operational industrial environment, as well as the baseline algorithm parameter calibration and optimization. In Section I.4, the main performance results of our novel proposed QoS-aware multi-connectivity scheme are compared to those from other state-of-the-art connectivity schemes, and the main findings are analyzed; followed by conclusions and future outlook discussion in Section I.5.

I.2 QoS-Aware Link Selection Scheme

In this section, we introduce the Selective Duplication with QoS (SDQoS) scheme, a novel cross-layer multi-connectivity link selection scheme that is able to shape and steer traffic from distinct packet flows with separate QoS requirements over a shared set of network paths.

For its formal definition, we consider generally F packet flows ($J = \{J_1, J_2, \dots, J_F\}$) mapped over N network paths ($P = \{P_1, P_2, \dots, P_N\}$), where each of the flows has its own set of QoS policies. The operating principle of the SDQoS scheme, as shown in Figure I.1, consists of two stages:

1. Radio-aware link selection: the full set of N network paths are evaluated based on radio quality which leads to selecting a sub-set of M suitable paths for transmission ($I = \{I_1, I_2, \dots, I_M\}$).
2. Congestion-aware link selection: a further down-selection of the M interfaces based on transport layer statistics and a QoS policy is performed. K interfaces are selected for transmission of the F packet flows.




Fig. I.1: Overview of the proposed two-stage SDQoS scheme.

The procedure for the radio-aware link selection in stage 1 is presented in Algorithm 1. This algorithm sorts the network paths according to the value of an associated radio quality metric r_i , $i = 1, \dots, N$, and selects the best M paths. The sorting happens accordingly to a RAT-dependent function $f(\cdot)$ of the quality metric of the paths and a threshold δ , which introduces a margin in the sorting process. In practice, a value v_i is sorted to be larger than v_j if and only if $v_i > v_j + \delta$.

Then, the selected M paths undergo another selection to K interfaces by taking into account a set QoS parameters (q_i, BW_i) associated with each flow. Here, BW_i denotes an allocated bandwidth for the flow i.e. the priority and q_i a parameter determining optional replication of packets across multiple network paths to further improve reliability. The details of the selection process in stage 2 are presented in Algorithm 2, where high priority flows are critical and served with duplication over $K = 2$ interfaces, while low priority flows are less critical and only served on a single interface. Then, the size of the estimated bandwidth EBW is compared to the parameter BW_i of each flow to decide whether the packet is transmitted or not. The intended effect of BW_i for flow a F_i given a link capacity EBW is illustrated with an example in Figure I.2. In this simplified example, there are three flows F_1, F_2 and F_3 with associated parameters BW_1, BW_2 and BW_3 . While the link capacity EBW is above both thresholds, packet transmissions occur normally. However, when EBW decreases and goes below a threshold BW_i , the packet transmission for that specific flow will be stopped until EBW recovers and ensures enough capacity to re-establish the transmission of the flow. It should be noted that, when significant congestion is experienced ($q_i == high$) and duplication is introduced ($K \leftarrow 2$), a short delay could potentially be added to fine-tune the connection and avoid a more serious congestion.

In case multiple groups of paths belonging to different RATs are available, Algorithm 1 would be used to select candidate paths within each technology group and afterwards, and Algorithm 2 would be used to perform link selection based on transport layer mechanisms in a fully-agnostic manner to the underline RAT. Note that the radio-aware 'pre-filtering' of the communication interfaces simplifies the next QoS-aware scheduling procedure. This is meant as a first approach towards novel cross-layer schedulers. An optimal scheduler should eventually take into account both radio performance and QoS parameters when performing the scheduling decision. However, this may lead to a significantly larger complexity in the presence of several radio interfaces and technologies, and will be considered for future work.

Algorithm 1 Radio-aware link selection

```

1:  $I \leftarrow \{I_1, I_2, \dots, I_M\}$  ▷ Available radio interfaces
2:  $P \leftarrow \{p_1, p_2, \dots, p_N\}$  ▷ Observable paths by  $I$ 
3:  $R \leftarrow \{r_1, r_2, \dots, r_N\}$  ▷ Quality metrics for  $p_i \in P$ 
4:  $P' \leftarrow \text{SORTDESCENDING}(P, f(\cdot), \delta)$  ▷ Sort  $P$  with respect to quality mapped by some function  $f$ 
5: for  $m \leftarrow 1, M$  do
6:    $I_m \leftarrow P'_m$  ▷ Assign paths to interfaces
7: end for

```

Algorithm 2 Congestion-aware link selection - SDQoS

```

1:  $q_i, BW_i \leftarrow \text{GETQOSOF PACKET}(pkt)$ 
2: if  $q_i == \text{high}$  then
3:    $K \leftarrow 2$ 
4: else
5:    $K \leftarrow 1$ 
6: end if
7: for  $v \leftarrow 1, K$  do
8:   if  $EBW(I_v) > BW_i$  then
9:      $\text{TRANSMIT}(I_v, pkt)$ 
10:  end if
11: end for

```




Fig. I.2: Example of the impact of the available link capacity EBW on BW_i and flow F_j .

I.2.1 QoS-Aware Multi-Connectivity Scheme Implementation

In order to validate its design, the presented novel cross-layer QoS-aware multi-connectivity link selection scheme was implemented. Figure I.3 details the functional block diagram of the implementation, which is based on the wireless multi-access gateway (MAGW) with layer-2 tunneling capabilities, previously presented in [16]. This device allows for transparent end-to-end integration of Ethernet-based systems. In transmission, this is done by taking Ethernet packets as an input, processing them, and encapsulating their information into configurable transport protocol packets that are later sent for wireless transmission on selected interfaces. In reception, the process is opposite. The received wireless packets are processed, and their transport protocol payload is decapsulated to re-generate the original Ethernet-based frame information, which is later forwarded as an output.

A number of modifications and inputs are to be provided in order to accommodate the proposed two-step scheduling design to the current MAGW implementation. First, flow awareness needs to map Ethernet (or IP) traffic to a specific QoS rule (sets of BW_i , q_i). Second, the scheduling algorithm needs to receive an input of M interfaces, sorted according to the link quality (independently of the RAT). Finally, suitable transport layer statistics and flow parameters are needed to properly shape the outgoing traffic according to the specific BW_i and q_i constraints for the packet in question.

The specific implementation chosen for presentation in this paper is based on the following design choices:

1. For flow mapping, we leverage IEEE 802.1Q [17], which allows for traffic separation via the use of Virtual LAN (VLAN) IDs in the Ethernet header. Furthermore, this allows for full support of Ethernet and IP traffic flows.
2. For simplicity, we focus on a single RAT (IEEE 802.11, Wi-Fi). However, note that, while the study focuses only on a single RAT, we separate the radio-aware link selection part of from the rest of the tunnel implementation in order to maintain flexibility for future RAT additions (e.g., 5G).
3. Finally, we make use of QUIC [18] as transport layer protocol. This provides a simple way to obtain link capacity estimates in a similar manner to the stream-aware congestion algorithm from our previous work [14].

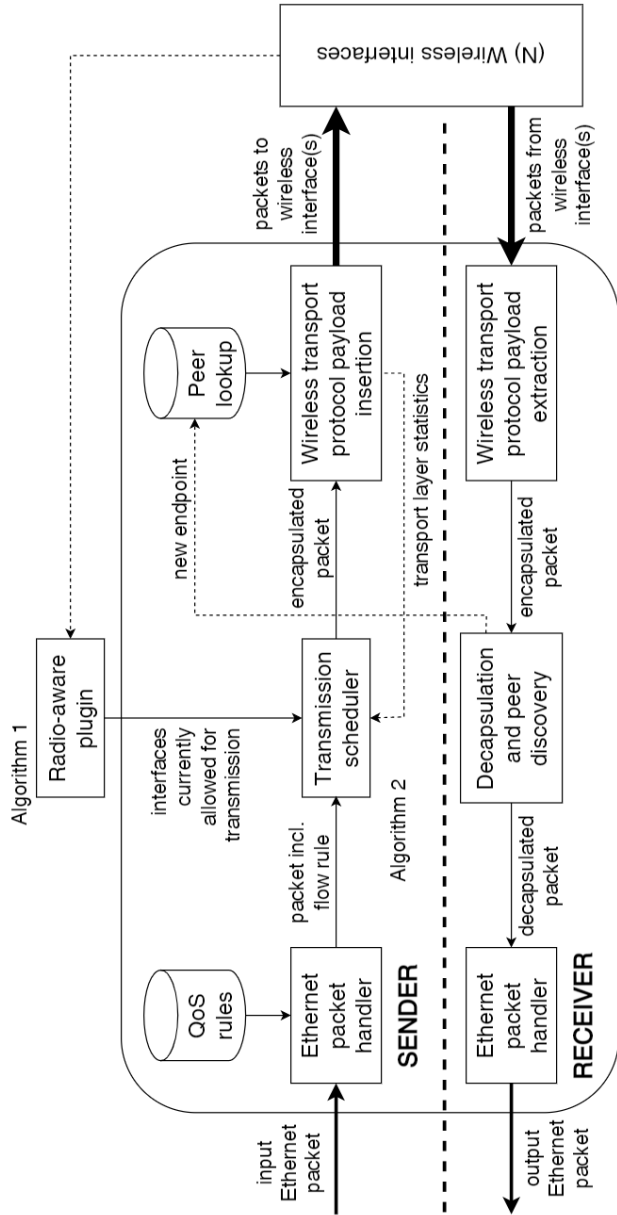


Fig. I.3: Functional block diagram description of the implemented MAGW with QoS-aware multi-connectivity capabilities.

Under those considerations, the functional blocks of the updated MAGW with QoS-aware multi-connectivity capabilities can be described as follows:

- **Radio-aware plugin:** provides a set of M sorted interfaces eligible for transmission according to Algorithm 1. As Wi-Fi is the selected RAT for the implementation shown in study, the selected radio quality metric r_i for sorting the available interfaces is the average beacon Received Signal Strength Indicator (RSSI) measured for a given service Access Point (AP). Therefore, the quality metric function $f(\cdot)$ in the implemented radio-aware link selection algorithm is simply defined as

$$f(RSSI, \theta) = \begin{cases} RSSI, & \text{if } RSSI > \theta \\ -\infty, & \text{if } RSSI \leq \theta \end{cases}$$

where θ is the value which triggers a handover to another AP. The radio-aware link selection procedure will then avoid including in the sorting procedure those paths who are about not to be eligible given the low signal quality. The implementation is based on wpa_supplicant [19] which also allows to perform the connection management. This includes initiating and updating network scans when the trigger level θ is hit, re-establishing connectivity in case of disconnections, maintaining mutual-exclusive connectivity for the modems, as well as polling the radio signal statistics.

- **QoS rules:** define the mapping of the VLAN IDs to each flow. This corresponds to the `GETQoSOF PACKET`, resulting in the q_i flag and the BW_i as per Algorithm 2.
- **Peer lookup:** contains the mapping of Medium Access Control (MAC) addresses behind a gateway (tunneled). The implemented look-up table relies on the Address Resolution Protocol (ARP) for automated network discovery, i.e., each endpoint MAC address gets assigned to a gateway ID.
- **Ethernet packet handler (sender)** is responsible for handling input Ethernet frames, extracting the 802.1Q header, and looking up the associated QoS rule.
- **Transmission scheduler** is responsible for scheduling an Ethernet frame according to Algorithm 2 based on the input received from the Ethernet packet handler (a QoS rule), the radio-aware plugin (sorted list of wireless interfaces allowed for current packet transmission), and protocol-specific statistics (depending on the configured transport protocol). Once a frame is scheduled, it is encapsulated and ready for

transmission. This would correspond to the entire cross-layer multi-connectivity link selection scheme, integrating the radio-aware and the congestion-aware link selection stages.

- **Wireless transport protocol payload insertion** takes care of setting the encapsulated information as payload over a selected transport protocol (layer-4) for wireless transmission, as well as of handling broadcast and unicast packets. This toggles the use of Algorithm 2 based on the configured QUIC transport layer protocol (where quinn-rs [24] is used). Once the packets are ready, they are forwarded to the wireless interfaces for actual transmission.
- **Wireless transport protocol payload extraction** performs reverse processing as compared to the wireless transport protocol payload insertion. It takes care of packet reception over the actual wireless interfaces, and decodes the transport layer payload to forward it to the next stage for decapsulation of the information.
- **Decapsulation and peer discovery** handles incoming data from the underlying wireless transport protocol and forwards the information for (re)constructing an Ethernet frame. It also updates the peer lookup table in case the source MAC has not been observed before, e.g., in case new devices are added to the network.
- **Ethernet packet handler (receiver)** is responsible for handling decapsulated received information and shaping it as Ethernet frames.

More details about the specific hardware and software components used in the implementation of the MAGW terminals with QoS-aware multi-connectivity capabilities are given in Table I.1.

Table I.1: Summary of hardware and software components used in the implementation of the MAGW with QoS-aware multi-connectivity capabilities.

Component	Hardware	Software
MAGW platform	Gateworks GW6404 [20] 4x OcteonTX @ 1.5 GHz 4GB DDR4 5x 1 Gb/s Ethernet	focal-newport [21] (Ubuntu, Linux 5.14)
Wi-Fi 6 interfaces	2x Intel AX200 rev. 1a [22]	wpa_supplicant 2.9 [19] Intel iwlwifi core 69 [23]

I.3 Experimental Validation

The experimental validation was performed at the AAU 5G Smart Production Lab, in Aalborg, Denmark. This lab, displayed in Figure I.4, is an industrial research testing facility composed of two small operational factory halls, expanding over a total area of 1250 m², resembling a real world factory, where researchers have access to a wide range of operational industrial-grade manufacturing and production equipment from different vendors including production line modules, robotic manipulators, AMRs, etc. [2].




Fig. I.4: Picture of the test setup in the AAU 5G Smart Production Lab, including the MAGW mounted on the AMR and the Wi-Fi 6 infrastructure.

I.3.1 Test description

A set of tests was conducted in order to evaluate the performance of the proposed QoS-aware multi-connectivity scheduling scheme for industrial applications. In our tests, we considered an AMR which needs to have two distinct traffic flows serviced via Wi-Fi network:

1. A critical control loop which handles the operations of the AMR from a centralized server location. This loops typically handles Priority Data (PD) such as robot status or mission information updates.
2. An uplink Best Effort (BE) stream of background data (BD); which depending on the AMR specific use could be, for example, video data (in case of surveillance or industrial process inspection robot) or robotic arm steering control information (in case of a mobile manipulator robot).

To capture the impact and performance gains of the proposed SDQoS scheme, it is benchmarked against the following 3 alternative state-of-the-art Wi-Fi based connectivity strategies:

- Single Path (SP) Wi-Fi connectivity with optimized handover; as presented in [15].
- Best Path Scheduling (BPS); as presented in [12]. This corresponds to a generalized execution of Algorithm 1 for the case $M = 2$, $K = 1$, i.e. there are two paths and there is no duplication.
- Selective Duplication (SD), i.e. duplicating the critical traffic and steering the background traffic according to the BPS. The SD procedure is shown in Algorithm 3, which is a simplified version of Algorithm 2, where the congestion control based QoS mechanism is missing, meaning that infinite capacity is assumed.

Algorithm 3 Congestion-aware link selection - SD

```

1:  $q_i \leftarrow \text{GETQOSOF PACKET}(pkt)$ 
2: if  $q_i == high$  then
3:    $K \leftarrow 2$ 
4: else
5:    $K \leftarrow 1$ 
6: end if
7: for  $v \leftarrow 1, K$  do
8:    $\text{TRANSMIT}(I_v, pkt)$ 
9: end for

```

It should be noted that, as the SP, BPS, and SD connectivity schemes do not make use of QoS indicators, their implementation makes use of User Datagram Protocol (UDP) as transport layer protocol; this is different from the proposed SDQoS scheme whose implementation is QUIC-based.

I.3.2 Setup and Measurement Procedures

Practically, testing was done with a MiR200 AMR [26] equipped with the MAGW with the implemented QoS-aware multi-connectivity scheme, which was set to navigate a predefined route with an approximated total length of 140 m in the factory test environment. Such route is marked as a black thick line in Figure I.5, which displays the layout of the factory test environment. The route is designed such that there are at least two APs at each position with a RSSI level above the handover threshold θ . Multiple iterations are taken in the route, in order to collect a sufficiently large number of measurements. The MiR200 is able to navigate between predefined checkpoints, using Light Detection and Ranging (LiDAR) technology combined with other onboard built-in sensors such as 3D cameras or proximity sensors, with a ± 5 cm accuracy. The use of such an AMR allows for automated execution of the measurement route. This ensures repeatability of the mobility pattern across the industrial scenario in the different measurement campaigns performed in the study. The factory floor is equipped with 3 CISCO Meraki MR36 Wi-Fi 6 APs [25], which are deployed at the positions also highlighted in Figure I.5. The APs are ceiling-mounted and provide full coverage to both factory halls. The test environment is fully radio-controlled, meaning that only devices and APs that are a part of the experiment generate traffic in the network, and that their spectral channel allocations can be fully organized. In this case, each of the APs are configured with 20 MHz non-overlapping channels in the 5.6 GHz region.

For the AMR data flows, the critical PD control loop is emulated via transmission of *ping* messages of 64 bytes between AMR and a centralized server location, transmitted in intervals of 50 ms. The BE BD stream is generated via *iperf3* transmission started at the AMR at a rate of 10 Mbps in uplink. Each of the flows is continuously monitored by the MAGW software in order to collect relevant performance statistics in terms of Round-Trip Time (RTT) latency, throughput, packet loss, and number of Wi-Fi handovers (number of switches between APs). For the critical control loop analysis based on PD, at least 100.000 samples are collected, while for the BE stream of BD, the amount of collected data in the same runtime is limited to 50.000 samples due to practical limitations of the tool [27]. Such a difference in measurement samples available for analysis between the two types of IIoT traffic types is not significant due to their different nature. While the PD can be considered of ultra-reliable type, insight into high reliability percentiles of the statistical distributions of the results is needed to assess its performance. However, the BD is BE and thus, lower statistical significance would suffice to understand its behavior. In any case, with the collected number of samples per test for both types of traffic, it is possible to get a statistical significance with a 95% confidence level close to the 99.99% (10^{-4}) and 99.9% (10^{-3}) reliability levels




Fig. I.5: Layout of the AAU 5G Smart Production Lab factory floor including location of the Wi-Fi 6 AP infrastructure (coloured circles), AMR test route (black thick line), and position of the interfering (INT) devices (coloured crosses).

for the PD and BD traffic, respectively [28].

The evaluation of the implemented novel SDQoS scheme and comparison with the SP, BPS, and SD strategies were done for both idle network and loaded network cases. In the idle network case, the only active device in the multi-AP Wi-Fi 6 network was the MAGW on the AMR with the two configured traffic flows. In the loaded network case, two additional active Interferer (INT) devices were deployed per AP. These were deployed co-located at the static positions denoted by the colored crosses in Figure I.5, and were configured to generate symmetric traffic, one by generating uplink traffic and the other by generating downlink traffic. Different interference levels (5, 10 and 25 Mb/s) were tested in order to gain further insights on the performance of the connectivity schemes under different network load conditions. Although the experiments presented in this study were performed with a single active AMR, the effect of having other active AMRs is implicit in the test cases considering interference.

I.3. Experimental Validation

Figure I.6 illustrates the full multi-connectivity test setup including the AMR with the two configured emulated traffic flows handled by the MAGW (PD flow between AMR controller and edge cloud server, and BD video flow from a camera sensor to the edge cloud server), the 3 Wi-Fi APs, and the interfering devices. Utilization of the implemented connectivity solutions require one MAGW at each end of the communication link. In this centralized IIoT case, the two MAGW are deployed as follows: one at the AMR (GW1 - device side) and one at the edge cloud server (GW2 - network side). Note that Device-to-Device (D2D) decentralized communication will also be easily enabled and configured by using the MAGW, just by deploying the two MAGW at device sides on different machines. The figure also depicts the 2 active (multi-connectivity) connections to different APs, and how the PD and BD flows are mapped over these connections via MAGW using VLANs.

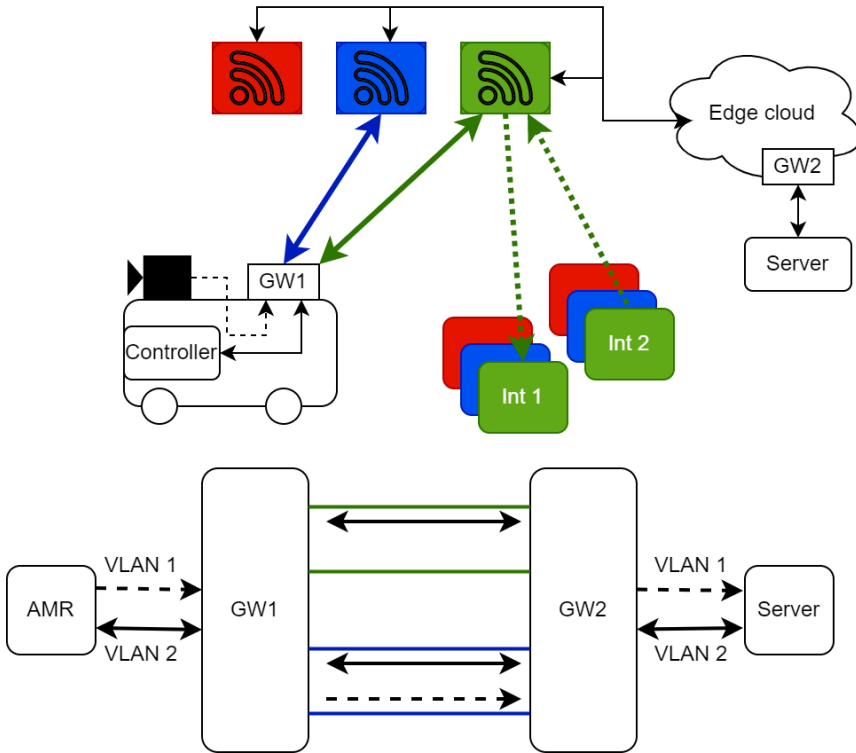


Fig. I.6: Simplified overview of the multi-connectivity setup including infrastructure and test elements, illustrating the critical PD (controller-server) and the BE BD data (camera-server) flows mapping over different VLAN connections (solid lines indicate for PD data and dashed lines for BE data).

I.3.3 Parameter optimization

As the sorting of the paths in Algorithm 1 depends on the threshold δ and the handover trigger level θ , we first performed a set of measurements to determine the optimal threshold in our industrial lab environment with respect to the two flow types. Note that both traffic flows are simultaneously active in each experiment. These measurements were performed using BPS with the cases of no interference (0 Mb/s) and 10 Mb/s (in each direction) for each AP. For each level of interference, we studied the impact of different threshold values, i.e. [0, 3, 9, 12] dB.

Figure I.7 displays the obtained results in terms of 99.9%-ile RTT and 99%-ile throughput, as well as loss rate, and number of switches, for different values of switching threshold, considering both no load and loaded cases. The RTT results show clearly the benefit of using a small value of the switching threshold, especially in case of high load. This is because a high threshold leads to 'stickiness' to a certain AP, which delays the switching to another AP with more favorable radio conditions. The presence of the BE BD flow in interfering scenarios and degraded radio conditions leads to the usage of more conservative transmission modes, therefore occupying more resources and possibly delaying the transmission of critical PD messages. Though a 0 dB threshold should offer the best performance as it results in instantaneous switching, it leads to a small throughput degradation and a significant packet loss rate, especially in the presence of interference. As it can be deduced from the number of switches results, this can be due to the emergence of ping-pong effects in the selection of the serving AP, that in our implementation may cause 'race conditions' and packet drops prior transmission.

Based on these results and observations, we set the δ threshold to 3 dB to ensure optimal performance of the connectivity schemes in our industrial lab test scenario, with reduced occurrence of switching events (and packet losses), and limited throughput impact, while still preventing excessive 'stickiness' to a given AP. Using this threshold configuration together with a handover trigger level, θ , of -80 dBm, we experience the link selection behavior described by the trace displayed in Figure I.8. In this trace, the MAGW is initially connected to two APs at a time (AP2 and AP3), perceiving different RSSI levels, with AP2 and AP3 acting as primary AP and secondary AP, respectively. When the weaker secondary link approaches the -80 dBm threshold, the MAGW triggers a handover operation, and AP1 becomes the secondary AP. As the GW is moving further from AP2, the RSSI drops and a switch of the primary AP role from AP2 to AP1 happens once the received RSSI from AP1 is higher than the RSSI from AP2 by a margin of 3 dB.

I.3. Experimental Validation




Fig. I.7: Summary statistics of the switching threshold optimization test.




Fig. I.8: Measurement test trace illustrating the behaviour of the link selection algorithm with the chosen parameters.

I.4 Results and Discussion

To compare the performance of the SDQoS scheme to the one from SP, BPS, and SD, results are presented in terms of complementary cumulative distribution function (CCDF) of RTT for the critical PD flow (Figure I.9), and throughput for the BD BE flow (Figure I.10), for all test configurations (idle network and load conditions).




Fig. I.9: CCDF for the PD (critical data flow) RTT for the novel SDQoS and the other scheduling schemes for different network load conditions.

As observed in Figure I.9, for the PD flow, SP offers the worst performance of all of the schemes even at no load. This is due to the impact of the Wi-Fi handover as well as excessive ‘stickiness’ to primary APs. Also, performance worsens visibly as load increases. BPS shows a visible gain with respect to SP thanks to the radio-aware link selection mechanism. When comparing it to the more advanced scheduling mechanisms, we observe that BPS offers similar performance to SD and SDQoS at 0 Mb/s extra network load (idle network

configuration) for the 99.99%-ile, but degrades quickly once additional network interference load is introduced. This is because BPS is not designed to leverage the additional network path for duplication. SD offers better or similar performance to the SDQoS scheme up to the 5 Mb/s network load. This is to be expected since the overhead of QUIC as compared to pure UDP introduces additional delay in the system. However, at increased network load (10 Mb/s) we observe a performance degradation of 26 ms, indicating that the capacity of the network is being exceeded as we can no longer serve the data flows fully. This performance difference increases up to approximately 100 ms further once the load is increased to 25 Mb/s. In general, our novel QoS-aware multi-connectivity scheme, SDQoS, offers consistent performance of around 57 ms for low-medium load levels (0, 10 Mb/s) and increases slightly to 85 ms in the presence of 25 Mb/s interfering load, indicating that, although a capacity issue arises, its effect is not very severe.




Fig. I.10: CCDF for the BD (best effort secondary data flow) server-side throughput for the novel SDQoS and the other scheduling schemes for different network load conditions.

While it is clear from the PD flow performance results that SDQoS offers better QoS (prioritization) when the network load is increased, this performance gain comes at a cost. Since the QoS mechanism deliberately drops incoming packets if available capacity is estimated to be lower than the reserved bandwidth for the PD flow, we expect there to be a clear difference in the BD data rate at high network load levels. Thus, Figure I.10 illustrates a clear performance throttling of the BE BD whose trend is consistent for both 0 and 5 Mb/s where it is around 7.9 Mb/s at the 99%-ile, whereas for SD it is around 9 Mb/s even though the PD RTT is similar. This suggests that the throttling mechanism may be too aggressive with the setting used for the test at these load levels. When increased to 10 Mb/s, the throttling is however at a similar level but with a performance gain for the priority data. At the 25 Mb/s level, SDQoS throughput decreases to 6.3 Mb/s as compared to SD which is at 8.5 Mb/s, but the tail of the PD flow is significantly better contained.

Summary statistics for achievable priority data RTT and effect on the background data rates are collected in Table I.2. The PD critical flow latency performance gains at the 99.99% reliability level achieved with a certain scheme with respect to standard SP operation are computed as ΔSP_{RTT} . For the novel SDQoS multi-connectivity, these gains range from 62.5% to 90% for idle network and loaded network with 25 Mb/s interference traffic, respectively. As previously addressed, these gains for the critical data flow come at the cost of a reduced performance for the secondary flow. ΔSP_{TP} indicates the throughput performance difference for the BD flow between the different connectivity schemes as compared to the one achieved in standard SP operation mode. With SDQoS, the BD flow is throttled by 22% in idle network conditions, while the impact is reduced to 12% in loaded network conditions in the presence of 25 Mb/s interference traffic, which clearly illustrates the benefits of using the proposed novel scheme as compared to less advanced schemes.

It is worth mentioning that the results related to the BPS, SD and SDQoS schemes are marginally affected by implementation-related issues caused by the Wi-Fi cards. It has been observed throughout the testing that sporadic disconnections can occur; especially when multiple Wi-Fi cards are active simultaneously [29]. This small issue is not related to the implementation of the schedulers itself, and does not impact its decision-making; but rather to firmware/driver issues, which affect mainly the performance of the non-priority data flow. This impact is evident from the fact that the SP scheduler offers slightly better BD throughput than BPS and SD; where the expectation would be that these should be similar. Supplementary debugging could be done to address this issue and further optimize the implementation. This is, however, beyond the scope of the paper, as it is expected that the performance would only improve for all three schemes if this issue was not present, while keeping the observed relative performance and interference tolerance trends unchanged.

Table I.2: Summary statistics for the PD latency and BD throughput flows evaluated for the different scheduling schemes and network conditions.

Traffic type		Priority data (PD) RTT latency			Background data (BD) throughput	
Connectivity Scheme	Interference load [Mb/s]	min [ms]	median [ms]	99.99%-ile [ms]	median [Mb/s]	99.99%-ile [Mb/s]
				(ΔSP_{RTT}) [ms]		(ΔSP_{TP}) [Mb/s]
Single Path (SP)	0	2.11	3.51	166 (-)	10	9.94 (-)
	5	1.99	3.81	305 (-)	10	9.90 (-)
	10	1.51	4.73	433 (-)	10	9.76 (-)
	25	1.85	7.52	883 (-)	10	7.17 (-)
Best Path Scheduling (BPS)	0	1.63	2.66	66 (-100)	10	9.94 (0)
	5	1.60	3.57	130 (-175)	10	9.85 (-0.05)
	10	1.69	3.98	176 (-257)	10	9.87 (+0.09)
	25	1.7	5.24	270 (-613)	10	9.55 (+2.38)
Selective Duplication (SD)	0	1.74	2.58	60 (-106)	9.98	9.05 (-0.89)
	5	1.63	2.61	61 (-244)	9.95	9.04 (-0.86)
	10	1.58	3.09	87 (-346)	9.95	8.97 (-0.79)
	25	1.58	5.05	198 (-685)	9.93	8.54 (+1.37)
Selective Duplication with QoS (SDQoS)	0	1.87	3.61	58 (-104)	9.94	7.69 (-2.15)
	5	2.05	3.94	56 (-249)	9.90	8.09 (-1.81)
	10	2.06	4.27	60 (-373)	9.82	7.87 (-1.89)
	25	2.07	5.06	85 (-798)	9.50	6.30 (-0.87)

I.5 Conclusions and future work

In this paper we have presented Selective Duplication with QoS (SDQoS), a novel network-agnostic cross-layer scheduling solution which leverages radio access technology metrics and transport layer metrics to schedule data transmission across one or more links. The presented two-step design first down-selects a number of eligible network paths and afterwards applies QoS differentiation on the flows that are to be served based on the estimated capacity of the chosen paths. We developed and integrated the scheme as a part of an operational industrial multi-access gateway (equipped with two Wi-Fi interfaces) and experimentally-validated its performance by running tests based on a realistic AMR industrial use case in a real-world production environment.

Experimental results show that for low load levels up to 5 Mb/s, the SDQoS scheme does not achieve any gain over simpler state-of-the-art multi-connectivity techniques when evaluated for priority data packet delay at the 99.99%-ile. However, when load increased to 10 Mb/s we start to observe a clear distinction between SDQoS and SD with a ~ 25 ms latency reduction, which increases up to ~ 110 ms at 25 Mb/s to 113 ms, at the cost of slightly throttled background data flow. Therefore, the proposed solution shows a major improvement in preserving low packet latency of priority data in the presence of large background traffic as compared to state-of-the-art solutions.

Future work includes the evaluation of the presented QoS-aware multi-connectivity scheduling scheme for scenarios with larger deployments and number of links, including more explicit testing with multiple active AMRs. We also aim at extending the presented solution to other (cellular) technologies such as 5G NR and investigating its impact on performance as compared to technology-specific multi-connectivity implementations. Along the same lines, we plan implementing and testing QoS-aware multi-RAT multi-connectivity solutions. Finally, we aim at designing more advanced schedulers which jointly combine the link selection and transport layer QoS, instead of having it separated as a two-step process.

Funding

This work was partially supported by the Spanish Ministry of Science and Innovation under Ramon y Cajal Fellowship number RYC-2020-030676-I.

References

- [1] E. Sisinni, A. Saifullah, S. Han, U. Jennehag, and M. Gidlund, "Industrial Internet of Things: Challenges, Opportunities, and Directions," in *IEEE Transactions on Industrial Informatics*, vol. 14, no. 11, pp. 4724-4734, Nov. 2018. doi: 10.1109/TII.2018.2852491.
- [2] I. Rodriguez, R.S. Mogensen, A. Schjørring, M. Razzaghpour, R. Maldonado, G. Berardinelli, R. Adeogun, P.H. Christensen, P. Mogensen, O. Madsen, C. Møller, G. Pocovi, T. Kolding, C. Rosa, B. Jørgensen, and S. Barbera, "5G Swarm Production: Advanced Industrial Manufacturing Concepts Enabled by Wireless Automation," in *IEEE Communications Magazine*, vol. 59, no. 1, pp. 48-54, Jan. 2021. doi: 10.1109/MCOM.001.2000560.
- [3] R. S. Mogensen, I. Rodriguez, G. Berardinelli, G. Pocovi, and T. Kolding, "Empirical IIoT Data Traffic Analysis and Comparison to 3GPP 5G Models," in *IEEE 94th Vehicular Technology Conference (VTC2021-Fall)*, 2021. doi: 10.1109/VTC2021-Fall52928.2021.9625319.
- [4] M. Wollschlaeger, T. Sauter and J. Jasperneite, "The Future of Industrial Communication: Automation Networks in the Era of the Internet of Things and Industry 4.0," *IEEE Industrial Electronics Magazine*, vol. 11, no. 1, pp. 17-27, Mar. 2017. doi: 10.1109/MIE.2017.2649104.
- [5] R. S. Mogensen, I. Rodriguez, G. Berardinelli, A. Fink, R. Marcker, S. A. Markussen, T. Raunholt, T. E. Kolding, G. Pocovi, and S. Barbera, "Implementation and Trial Evaluation of a Wireless Manufacturing Execution System for Industry 4.0," in *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*, 2019, pp. 1-7. doi: 10.1109/VTCFall.2019.8891231.
- [6] M. Gidlund, T. Lennvall, and J. Åkerberg, "Will 5G become yet another wireless technology for industrial automation?," in *2017 IEEE International Conference on Industrial Technology (ICIT)*, 2017, pp. 1319-1324, doi: 10.1109/ICIT.2017.7915554.
- [7] P. Popovski, C. Stefanović, J. Nielsen, E. de Carvalho, M. Angjelichinoski, K. F. Trillingsgaard, and A-S. Bana, "Wireless Access in Ultra-Reliable Low-Latency Communication (URLLC)," *IEEE Transactions on Communications*, vol. 67, no. 8, pp. 5783-5801, Aug. 2019. doi: 10.1109/TCOMM.2019.2914652.
- [8] M. Centenaro, D. Laselva, J. Steiner, K. Pedersen, and P. Mogensen, "System-Level Study of Data Duplication Enhancements for 5G Downlink URLLC," *IEEE Access*, vol. 8, pp. 565-578, 2020, doi: 10.1109/ACCESS.2019.2961658.
- [9] C. Pupiales, D. Laselva, Q. De Coninck, A. Jain, and I. Demirkol, "Multi-Connectivity in Mobile Networks: Challenges and Benefits," *IEEE Communications Magazine*, vol. 59, no. 11, pp. 116-122, Nov. 2021, doi: 10.1109/MCOM.111.2100049.
- [10] 5G Alliance for Connected Industries and Automation (5G-ACIA), White Paper: Industrial 5G Devices – Architecture and Capabilities. [Online] Available: <https://5g-acia.org/whitepapers/industrial-5g-devices/>. Accessed on: August 15, 2022.

References

- [11] IWLAN – the WLAN for challenging industrial applications, SIEMENS. [Online]. Available: <https://new.siemens.com/global/en/products/automation/industrial-communication/industrial-wireless-lan.html>. Accessed on: August 15, 2022.
- [12] A. Fink, R. S. Mogensen, I. Rodriguez, T. Kolding, A. Karstensen, and G. Pocovi, "Radio-Aware Multi-Connectivity Solutions based on Layer-4 Scheduling for Wi-Fi in IIoT Scenarios," in IEEE Wireless Communications and Networking Conference, 2022.
- [13] F. Chiariotti, A. A. Deshpande, M. Giordani, K. Antonakoglou, T. Mahmoodi, and A. Zanella, "QUIC-EST: A QUIC-Enabled Scheduling and Transmission Scheme to Maximize VoI with Correlated Data Flows," IEEE Communications Magazine, vol. 59, no. 4, pp. 30-36, Apr. 2021, doi: 10.1109/MCOM.001.2000876.
- [14] R. S. Mogensen, C. Markmoller, T. K. Madsen, T. Kolding, G. Pocovi, and M. Lauridsen, "Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications," in IEEE 89th Vehicular Technology Conference (VTC2019-Spring), 2019. doi: 10.1109/VTCSpring.2019.8746482.
- [15] A. Fink, R. S. Mogensen, I. Rodriguez, T. Kolding, A. Karstensen, and G. Pocovi, "Empirical Performance Evaluation of EnterpriseWi-Fi for IIoT Applications Requiring Mobility," in 26th European Wireless Conference, 2021.
- [16] I. Rodriguez, R. S. Mogensen, A. Fink, T. Raunholt, S. Markussen, P. H. Christensen, G. Berardinelli, P. Mogensen, C. Schou, and O. Madsen, "An Experimental Framework for 5G Wireless System Integration into Industry 4.0 Applications," Energies, vol. 14, no. 15, 4444, 2021. doi: 10.3390/en14154444.
- [17] IEEE 802.1Q-2018: IEEE Standard for Local and Metropolitan Area Networks–Bridges and Bridged Networks, IEEE Standards. [Online]. Available: <https://standards.ieee.org/ieee/802.1Q/6844/>. Accessed on: August 15, 2022.
- [18] IETF RFC 9000 QUIC: A UDP-Based Multiplexed and Secure Transport. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc9000.html>. Accessed on: August 15, 2022.
- [19] Linux WPA Supplicant (IEEE 802.1X, WPA, WPA2, RSN, IEEE 802.11i). [Online]. Available: <https://w1.fi/wpasupplicant/>. Accessed on: August 15, 2022.
- [20] Gateworks Newport GW6404 Rugged & Industrial Single Board Computer. [Online]. Available: <https://www.gateworks.com/products/industrial-single-board-computers/octeon-tx-single-board-computers-gateworks-newport/gw6400-single-board-computer/>. Accessed on: August 15, 2022.
- [21] Gateworks - Ubuntu on Newport. [Online]. Available: <http://trac.gateworks.com/wiki/newport/ubuntu>. Accessed on: August 15, 2022.
- [22] Intel Wi-Fi 6 AX200. [Online]. Available: <https://www.intel.com/content/www/us/en/products/docs/wireless/wi-fi-6-ax200-module-brief.html>. Accessed on: August 15, 2022.

References

- [23] Intel backport-iwlwifi Core 69 Wi-Fi driver and firmware update. [Online]. Available: <https://github.com/intel/backport-iwlwifi>. Accessed on: August 15, 2022.
- [24] D. Ochtman, B. Saunders, and J.-C. Begue, "Quinn: Async-friendly QUIC implementation in Rust," GitHub Repository. [Online]. Available: <https://github.com/quinn-rs/quinn>. Accessed on: August 15, 2022.
- [25] CISCO Meraki MR36 - High Performance 802.11ax Wireless. [Online]. Available: https://documentation.meraki.com/MR/MR_Overview_and_Specifications/MR36_Datasheet. Accessed on: August 15, 2022.
- [26] MiR200, Mobile Industrial Robots. [Online]. Available: <https://www.mobile-industrial-robots.com/>. Accessed on: August 15, 2022.
- [27] iPerf - The ultimate speed test tool for TCP, UDP and SCTP. [Online]. Available: <https://iperf.fr/>. Accessed on: August 15, 2022.
- [28] L. D. Brown, T. T. Cai, and A. Dasgupta, "Confidence intervals for a binomial proportion and asymptotic expansions," *The Annals of Statistics*, vol. 30, no. 1, pp. 160-201, Feb. 2002, doi: 10.1214/aos/1015362189.
- [29] Bug 203709 - iwlwifi: 8260: frequently disconnects since Linux 5.1 "No beacon heard and the time event is over already" - WIFI-25906, Kernel.org Bugzilla. [Online]. Available: https://bugzilla.kernel.org/show_bug.cgi?id=203709. Accessed on: August 15, 2022.

ISSN (online): 2446-1628
ISBN (online): 978-87-7573-767-3

AALBORG UNIVERSITY PRESS