D4.4

REPORT ON VALIDATION OF 5G USE CASES IN THE FACTORY

The 5G-SMART project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no 857008.

Report on Validation of 5G Use Cases in the Factory

Grant agreement number: Project title: Project acronym: Project website: Programme:	857008 5G Smart Manufacturing 5G-SMART www.5gsmart.eu H2020-ICT-2018-3
Deliverable type: Deliverable reference number: Contributing workpackages: Dissemination level: Due date: Actual submission date:	Report (R) D16 WP4 Public 2022-05-31 2022-05-31
Responsible organization: Editor(s): Version number: Status:	Bosch Davit Harutyunyan V1.0 Final
Short abstract:	This deliverable introduces the validation scenarios defined for the Cloud-based mobile robotics and TSN/Industrial LAN over 5G use cases and presents their test results, including the encountered issues and learned lessons.
Keywords:	Industrial LAN over 5G, Controller-to-controller communication, Machine-to-server communication, Mobile robots, AGVs, 5G standalone (SA), cloud control, AGV collaboration

Contributor(s):	Davit Harutyunyan (editor, Bosch) David Ginthoer (Bosch) Peter Buseck (Bosch) Sven Boje (Bosch) Andreas Hohner (Bosch) Hans-Juergen Decker (Bosch) Norbert Reider (ERI-HU) Gabor Nemeth (ERI-HU) Sándor Rácz (ERI-HU) Attila Vidács (BME) Gábor Fehér (BME) Markosz Maliosz (BME) Dhruvin Patel (ERI-DE) Leefke Grosjean (ERI-SE) Joachim Sachs (ERI-SE) Erik Nacken (ERI-DE) Detlef Hoheisel (ERI-DE)

Disclaimer

This work has been performed in the framework of the H2020 project 5G-SMART co-funded by the EU. This information reflects the consortium's view, but the consortium is not liable for any use that may be made of any of the information contained therein.

This deliverable has been submitted to the EU commission, but it has not been reviewed and it has not been accepted by the EU commission yet.

Executive summary

This report is the fourth and the last deliverable of the activities pertaining to the trial and validation of 5G at the Bosch semiconductor factory in Reutlingen. It describes in detail the validation scenarios and their test results for two 5G-SMART use cases trialed at the Bosch semiconductor factory in Reutlingen, namely *Cloud-based mobile robotics* and *Time Sensitive Networking (TSN)/Industrial LAN (I-LAN) over 5G*. We first present the methodology adopted to validate and evaluate the use cases. We then introduce the industry goals set by the factory management for each of the use cases in order to find out if and how much each industry goal is affected by the use of the 5G mobile network technology instead of a WiFi and wired Ethernet technology, which are currently used at the factory by the applications considered in these use cases.

A brief summary of the use case is provided regarding the Cloud-based mobile robotics use case. For additional information, we refer the reader to 5G-SMART's Deliverable D1.1 [5GS20-D11] for more details on the use case and its requirements and 5G-SMART's Deliverable D4.3 [5GS20-D43] for the implementation details on the hardware and software architectures of the mobile robots in this use case. The industry goals for this use case are then discussed, specifying why they are essential for the factory. This is followed by introducing the Key Performance Indicators (KPIs) used to evaluate the defined industry goals. Nine validation scenarios are then presented along with their validation results, discussion, and conclusion.

As for the TSN/I-LAN over 5G use case, we first introduce applications considered in this use case. Specifically, we describe the controller-to-controller (C2C) communication, including its operating mode and the message flow between the controllers and the machine-to-server communication, referring the reader to 5G-SMART's Deliverable D1.1 [5GS20-D11] for more details. We then specify the KPIs used to evaluate the defined industry goals, followed by the definition of those KPIs and their calculation methodology. While the calculation of most of the KPIs requires experimental testing, some others, such as annual cost savings in setting up C2C communication and machine-to-server communication over the 5G network, require no experiments and are calculated based on already available information. After presenting the validation KPIs, we describe six validation scenarios considered in the TSN/I-LAN over 5G use case, including their validation results, discussion, and conclusion. Finally, we summarize the entire trial at the factory and provide concluding remarks.

All in all, it has been validated that 5G is capable of supporting the traffic of industrial applications. Moreover, it has been demonstrated that 5G wireless connectivity has a positive impact on most of the industry goals defined for the use cases with the conclusion being that the 5G technology is ready to be used for the considered industrial applications.

Contents

1 Introduction

One of the main goals of the trial and validation of 5G at the Bosch semiconductor factory in Reutlingen is to validate and evaluate two use cases, namely *Cloud-based mobile robotics* and *Time Sensitive Networking (TSN)/Industrial LAN (I-LAN) over 5G*. Cloud-based mobile robotics use case focuses on validation and evaluation of collaborative control of mobile robot or Automated Guided Vehicles (AGVs) from the factory cloud over the 5G network. Currently, the AGVs used at the factory are not aware of each other, which sometimes causes coordination issues (e.g., it may take some time for AGVs to bypass each other). Moreover, self-collected and limited knowlede of individual AGVs about their sourroundings may result in a non-optimal route selection, which in turn may reduce the efficieny of AGVs. A 5G-enhanced commercial AGV and a research AGV are employed in this use case, respectivelly, with partly and fully offloaded control functions to the factory cloud. Moving the control logic into the cloud benefits from scaling of the workload when changing the tasks for the AGVs, ease of maintenance of the control software and improved resiliency to software and hardware failures. Furthermore, decoupling the control logic from the AGV enables innovative control solutions such as collaboration between individual AGVs by, e.g., facilitating the creation and sharing of up-to-date common maps, which in turn can optimize the route selection of the AGVs.

TSN/I-LAN over 5G use case focuses on investigating and validating the applicability of 5G for transporting the traffic of I-LAN applications. This use case includes Controller-to-Controller (C2C) communication as well as machine-to-server communiaction. Nowadays, due to the stringent requirements of the industrial applications, all operational I-LANs are realized based on fixed (wired) communication networks. Limited flexibility for setting up new production lines or for restructuring an existing production line, as well as complex and costly maintenance, are major drawbacks of the wired I-LAN realizations. In particular, this can be an issue in view of the recent trends for making the industrial environments as flexible as possible, e.g., smart factories of the future in the context of Industry 4.0. Introducing 5G comes with the potential of reducing the cables and connectors wear and tear, for the mobile machines/controllers, resulting in reduced maintenance costs. Additionally, replacing the cables for communications between controllers and machines with 5G communications results in a greater flexibility for implementation and adaptation of the industrial manufacturing infrastructure.

Both use cases have been tested in various validation scenarios. Moreover, several Key Performance Indicators (KPIs) have been defined and used to evaluate the industry goals defined for these use cases and the feasibility of running this use cases over the 5G network.

1.1 Objective of the report

This report aims to describe the validation scenarios defined for the Cloud-based mobile robotics and Time Sensitive Networking (TSN)/Industrial-LAN (I-LAN) over 5G use cases and their test results, including the encountered issues and learned lessons. Besides, this report introduces the industry goals set for these use cases and describes the validation KPIs calculated in various validation scenarios to evaluate those industry goals.

1.2 Relation to other documents in 5G-SMART

This report takes input from 5G-SMART's Deliverable D1.1 [5GS20-D11], where the use cases trialed at the Bosch semiconductor factory in Reutlingen are presented along with their requirements and Key Performance Indicators (KPIs). It then builds on top of 5G-SMART's Deliverable D4.3 [5GS20-D43], which reports on the implementation and integration status of the Cloud-based mobile robotics and TSN/I-LAN over 5G use cases, also describing their end-to-end integrated system architecture. Specifically, this report introduces the validation scenarios and presents the testing results of these scenarios for both use cases. It also describes the industry goals set by the factory management for these use cases and specifies the KPIs used to evaluate those goals. Finally, this deliverable also takes input from 5G-SMART's Deliverable D4.1 [5GS20-D41], which reports on the 5G system deployment at the semiconductor factory in Reutlingen.

1.3 Structure of the report

This report is composed of five sections. After the introduction, Section 2 details the validation and evaluation methodology, industry goals for the use cases, and network characteristics at the trial site. Section 3 introduces the Cloud-based mobile robotics use case, and the KPIs used to evaluate its industry goals. It also presents the validation scenarios and results of this use case. This is followed by providing the relevant information for the TSN/I-LAN over 5G use case in Section 4. Finally, the report is summarized in Section 5.

2 Demonstration, validation, and evaluation of 5G capabilities

2.1 Validation and evaluation methodology

The validation methodology adopted for the Cloud-based mobile robotics and TSN/I-LAN over 5G use cases is illustrated in Figure 1. Based on the daily operation of the applications considered in these use cases, a set of industry goals has been defined by the management of the semiconductor factory in Reutlingen, as introduced in Section 2.2. These industry goals are in line with the goals of the use case analysis for business value creation defined in 5G-SMART's Deliverable D1.2 [5GS21-D12]. They have been evaluated based on a number of KPIs defined for these use cases, such as C2C communication efficiency, packet loss, Round Trip Time (RTT), power consumption, reaction time, and successful handover rate, as introduced in Section 3.3 and Section 4.3. These KPIs have been calculated in various validation scenarios defined per use case. A common map, cloud reliability, and cloud safety are some of the validation scenarios for the Cloud-based mobile robotics use case. In contrast, TSN/I-LAN over 5G use case considers validation scenarios such as C2C communication over the 5G network with and without background traffic in the network. The details of the validation scenarios and their results are presented in Section 3 and Section 4. All validation scenarios have then been tested over a Release 15-based standalone and private 5G network deployed in the production area of the semiconductor factory in Reutlingen. The details of the deployed 5G network characteristics are presented in Section 2.3.

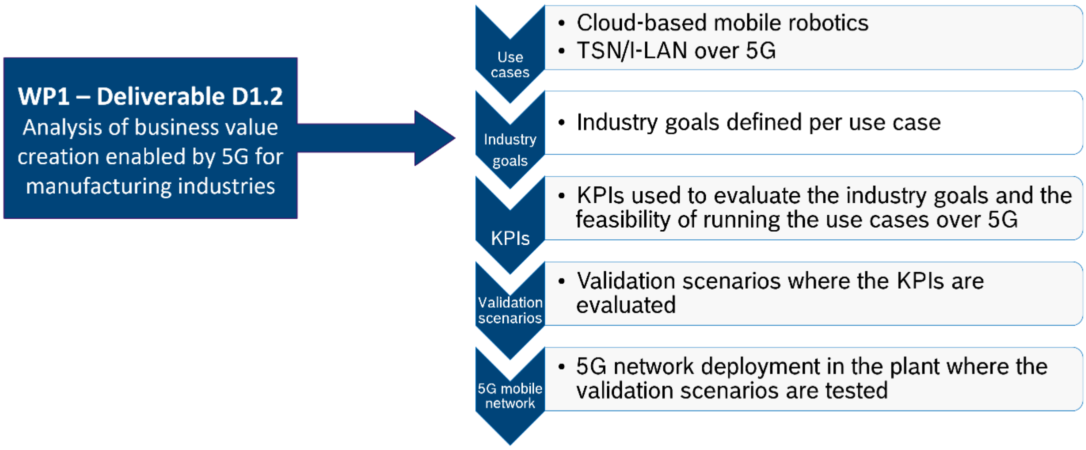

Figure 1: Validation methodology of the Cloud-based mobile robotics and TSN/I-LAN over 5G use cases

2.2 Industry goals

The overarching goal of investigating the Cloud-based mobile robotics and TSN/I-LAN use cases over 5G at the semiconductor factory in Reutlingen is to determine if these use cases are feasible to be run over 5G, as well as if and how much the execution of these use cases over the 5G network improves or deteriorates the industry goals set by the factory management. This section defines the industry goals (shown in Figure 2) for our use cases. In contrast, in Section 3.2 and Section 4.2, we specify the KPIs used to quantify and evaluate those industry goals, respectively, in the Cloud-based mobile robotics and TSN/I-LAN over 5G use cases. It is important to mention that the definition of the industry goals are sometimes different between the considered use cases. Moreover, not all the industry goals are targeted by both use cases (as shown in Table 1). Some of the goals are unaffected by implementing the use cases over the 5G network instead of realizing them over the current communication system. This includes wired Ethernet networks to enable C2C and machine-to-server communication, and WiFi networks to manage mobile robots on the shopfloor.

Figure 2: Industry goals of the Cloud-based mobile robotics and TSN/I-LAN over 5G use cases

Use cases	Industry goals					
	Flexibility	Mobility	Productivity	Quality	Safety	Sustainability
Cloud-base mobile robotics	x	x	x	x	x	x
TSN/Industrial LAN over 5G	x		x	x		x

Table 1: Industry goals per use case

2.2.1 Flexibility

Flexibility is one of the most important industry goals for the considered use cases. In the TSN/I-LAN over 5G use case, the C2C communication and machine-to-server communication are traditionally enabled via a wired Ethernet network, which incurs high maintenance and cabling cost if the network layout changes frequently. Moreover, it makes the customization and modification of production cells cumbersome. More specifically, when relocating the controllers or the industrial machines, one must make sure that proper cabling and network sockets are installed in the new location of these devices, which is a time-consuming and costly process. Additionally, for example, it is not feasible to do so via cable when retrofitting the machines with additional sensors. Especially in the exemplary semiconductor plant in Reutlingen, it is often impossible to lay additional cables across the factory due to safety and space limitations. Thus, one of the advantages of using the 5G mobile network communication for these applications in the TSN/I-LAN over 5G use case is the high level of flexibility when it comes to customization and modification the production cells, allowing them to be quickly restructured with a minimum engineering effort and changeover time.

To realize efficient transportation in smart factory scenarios with a large fleet of AGVs, there is a need for a high degree of modularity, programmability, and flexibility. Thus, our overarching industry goal is to bring dynamicity and flexibility into the operation of AGVs in the production environments with the application of the 5G technology. For instance, a manufacturing environment with a substantial number of AGVs requires dynamic adaptation and customization of AGV's operations due to changes in the environment, requirements, or operation policies. In that case, executing required software updates and configurations for AGVs individually should be avoided, since it can be resource-intensive and time-consuming. 5G enables time-critical communication with, e.g., bounded latencies and enables to shift the local intelligence from the individual AGVs and realize it in a cloud computing environment at the edge of the network infrastructure. Additionally, a co-location of software from several AGVs in the same processing environment is made possible, as well as enriching the control with data that otherwise would be inconvenient or infeasible on individual devices (e.g., shared maps, operational/usage data sets).

Further objectives that contribute to increased flexibility are to enable new use case development and accelerate innovations in the area of automated transportation in the factories of the future by, e.g., enabling the physical platform and control intelligence to evolve separately from each other. This can be achieved by decoupling the control intelligence from the robot hardware platform. Thus, interactions and collaborations among individual AGVs and control logic (also from other industrial devices) are facilitated based on other sensory inputs.

2.2.2 Mobility

Mobility is essential for mobile robots doing transportation in the factory. The size of the semiconductor factory requires mobility in wireless communication for seamless operation of the connected services. Thus, our goal is to validate that the performance of the 5G mobility capabilities support cloud-based robot control without any considerable negative effect on robot motion. Mobility should support a large number of AGVs and provide a seamless connection service for each robot.

The communication between external systems (e.g., fleet management) and individual AGVs over Wi-Fi might become inefficient and, in some cases, infeasible, e.g., operation of AGVs in large indoor/outdoor areas of the Bosch semiconductor factory has proven infeasible due to the limited reliability of Wi-Fi in the factory and hand-over delays. It is a target to increase mobility capabilities on the shopfloor with ubiquitous coverage and reliable network operation. A seamless mobility service is a requirement for the proper functioning of the AGVs across the whole mobility area. A seamless service across the whole mobility area is also important for machines (such as the Rudolph F30 used in the TSN/I-LAN use case) since their position may change inside the factory in case of rearrangement of the production.

2.2.3 Productivity

There are various factors affecting the plant's productivity, and changing the existing cable-based communication for the industrial machines into a 5G-based communication can increase the plant's productivity. One of such factors is the downtime of production machines. Since some of the production machines have moving stations, modules, and external sensors, laying Unshielded Twisted Pair (UTP) cables to connect the machines to the central control unit might damage existing cable connections, thereby resulting in unplanned downtime. Therefore, replacing these connections with 5G-communication would reduce downtimes. Another factor impacting the plant's productivity is the operational speed of each production machine. Each machine is controlled by one or multiple control applications. The industry goal “productivity” for the TSN/I-LAN over 5G use case refers to the productivity of the control application in the C2C communication. We define the productivity of the control application based on the rate the programmable logic controller (PLC) can process commands, which depends on the network characteristics (e.g., packet transmission delay, RTT) between any two machines exchanging control commands using C2C communication. While the 5G network communication may increase the productivity of the plant due to, for example, a shorter C2C communication setup time and reduced downtime, it is not expected to increase the productivity of the control application in the C2C communication due to higher packet transmission delay and RTT compared to the wired Ethernet network.

In the Cloud-based mobile robotics use case, productivity is aimed to be improved by increasing transportation efficiency. Efficiency can be directly increased, e.g., shorter mission execution times, by optimized route selection and real-time replanning based on the shared map and its continuous update from the robots. Furthermore, due to improved software management (e.g., redundant software execution for increased reliability), downtimes can be reduced. The total cost of ownership (TCO) is also expected to lower due to the reduction in maintenance complexity.

2.2.4 Quality

Since the communication technology and medium are changed for the applications considered in the TSN/I-LAN over 5G use case, it is of paramount importance for the plant to estimate the communication quality, which is expressed in terms of packet losses, between the controllers as well as between the industrial machine and its backend server to be compared with that estimated over the wired Ethernet network.

For the Cloud-based mobile robotics use case, the goal is to provide high-quality mobile robot operation that requires high quality robot hardware and control software architectures and a high quality execution environment and network connection in terms of latency, reliability, and connection lost. For example, keeping the quality of robot motion high means avoiding sudden stops, glitches in the movements and realizing agile reactions to unexpected events that may occur in the environments of the AGVs.

2.2.5 Safety

When it comes to transportation by mobile robots, safety is one of the top-priority goals. Mobile robots need to move freely across the whole production area, which is usually an environment occupied by humans (workers, visitors, etc.). Collisions of AGVs with humans can cause serious injury, especially in cases with heavy and fast mobile robots. Safety mechanisms have to be considered. Typically, mobile robots rely on safety-certified hardware components, such as safety laser scanners that can detect objects around the robot and categorize them into different safety zones based on the distance from the robot. Usually, there are yellow and red zones identified, and the scanners provide reliable trigger output that can be used as input to a safety PLC or directly to a servo motor to slow down or stop the robot platform.

The overall aim is to provide the transportation service without any collisions, emergency stop and interruption. We implemented certified HW-based safety mechanisms that run onboard. Besides that, we investigated a software-based safety solution called cloud safety that can be implemented in the factory cloud next to the control stack and predict potential collision cases and stop the robot platform in time. Investigating the possibility to execute robot safety functions from the cloud environment over 5G is another important industry goal for the project.

2.2.6 Sustainability

Sustainability is yet another important industry goal for the factory. The TSN/I-LAN over 5G use case is expressed in terms of the annual cost savings which can be achieved as a result of replacing the wired communication between the industrial machines (e.g., Rudolph F30) and their servers. There are various factors, such as reduced cabling, network sockets, and maintenance effort, attributing to the savings in the capital and operational expenses for setting up and maintaining the industrial machines used in the TSN/I-LAN over 5G use case.

Also, for the Cloud-based mobile robotics use case, it is expected to increase transportation's energy- and cost-efficiency. It is achieved by introducing real-time collaboration among the AGVs interacting with the environment to execute a challenging task. The availability and the real-time update of the common map enable the AGVs to optimize their trajectories by avoiding blocked routes and thereby saving significant travel time and distance. If the mobile robots can reduce mission execution time, the number of robots needed for a set of transportation tasks may decrease reducing the transportation cost. Also, the floor space needed for transportation paths is reduced, and the probability of vehicles blocking each other (and reducing task execution efficiency) is lowered. Furthermore, the simplified robot HW architecture such as the Research AGV can lower equipment costs.

2.3 Network characteristics at the trial site

Several steps have been undertaken to determine a suitable 5G network deployment in the factory's production area. Initially, the deployment constraints have been derived considering the trial site constraints, including coverage, performance, and other integration aspects. These constraints have then been considered in the 5G deployment plan as well as in the 5G network solution design. Finally, a radio planning activity was conducted in order to provide a suitable location for placing radio antennas to ensure full coverage and performance across the shopfloor, and the 5G network was deployed.

The 5G mobile network deployed at the Reutlingen trial site is a dedicated, on-premise, and standalone (SA) 5G network based on 3GPP Release 15 specification. It is operated as an isolated non-public network (NPN), which means that all network components, such as the Radio Access Network (RAN) and the 5G core network, required to operate the 5G system are located in the semiconductor factory. Thus, the non-public network is separate from any public network and, therefore, both the control and user data traffic stay within the defined premises. The RAN of the 5G network deployment, as shown in Figure 3 is realized with the Ericsson Radio Dot System (RDS) [RDS], which is specifically designed to provide a flexible and cost-effective architecture for indoor deployments.

The 5G network operates in the 3.7-3.8 GHz spectrum, and its deployment covers around 8000 sqm area with a total of 16 RDs to provide good coverage and high performance over the entire factory shopfloor. Thus, 100 MHz bandwidth is used to test and validate the use cases. The deployment design has taken into account the challenging manufacturing environment and in particular the location of all potential AGV routes and industrial machines. The deployment is composed of two cells to enable handover investigations. Each of these two cells is composed of eight Radio Dots (RDs), which are mounted on the ceiling and on the walls in the production area according to the radio planning results to provide radio coverage in the specified area. More details about the network deployment can be found in the 5G-SMART's Deliverable D4.1 [5GS20-D41]. Basic coverage and throughput tests have been performed prior to use case execution.

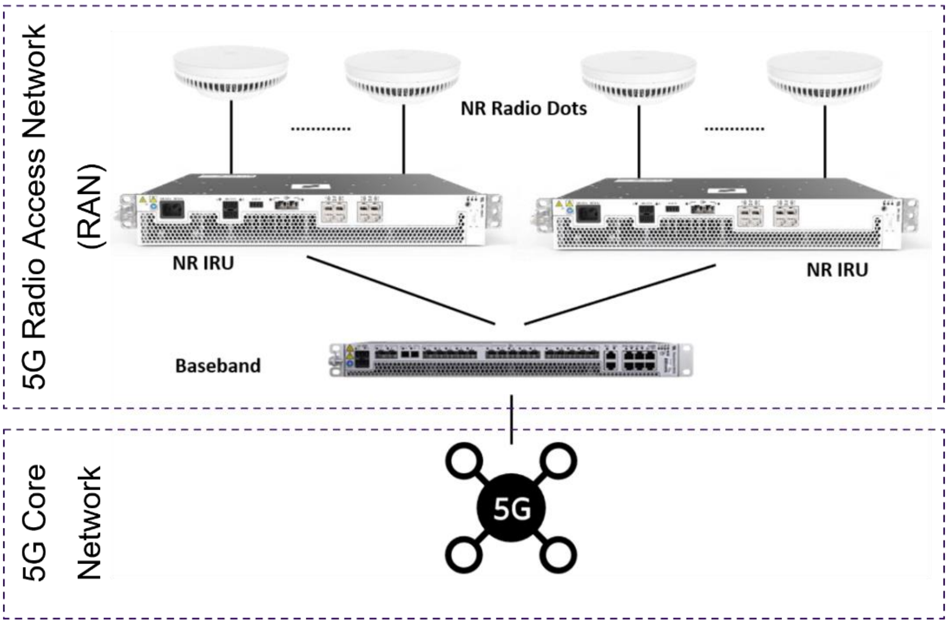

Figure 3: The 2-cell 5G standalone RAN deployment at the Reutlingen trial site, Source: Ericsson. The RAN consists of three main components: Baseband Unit (BBU), Indoor Radio Unit (IRU), and Radio Dot (RD). Each RD is connected to an IRU and is in charge of providing radio coverage, while the BBU, which is connected to the 5G core, is responsible for the baseband processing of radio waves.

3 Cloud-based mobile robotics

3.1 Introduction

This use case focuses on the feasibility, flexibility, and performance of wirelessly controlled Automated Guided Vehicles (AGVs) on a manufacturing shop-floor equipped with 5G technology. In this section, only a short summary is given since a detailed description of this use case, and its requirements can be found in 5G-SMART's Deliverable D1.1 [5GS20-D11]. The AGVs used in this work do not rely on any guidance infrastructure such as colored or magnetic stripes on the floor or any other markers. They are fully autonomous mobile robots (due to consistency with earlier deliverables, the terminology of AGVs will be used throughout this document). Besides the need for low-latency and reliable radio connectivity provided by 5G as an enabling technology, one novelty of this use case is the possibility to decouple the closed-loop control of the robot from the robot's embedded system and place it into an edge cloud execution environment (i.e., a factory cloud, please see our terminology document in [5GS20-CT]) while sustaining the KPIs, like sufficiently low execution latency, dynamic and flexible behavior, and adequate fault-tolerance. Moving the control logic into the cloud benefits from ease of maintenance of the control software and improved resiliency to software (SW) and hardware (HW) failures. Furthermore, decoupling the control logic from the AGV enables innovative control solutions such as collaboration between individual AGVs by, e.g., facilitating the creation and sharing of up-to-date common maps. For instance, object detection and navigation capabilities of an AGV can enhance the route selection for other AGVs in real-time, i.e., one AGV detects an obstacle, and the other one reacts by finding another path to the destination.

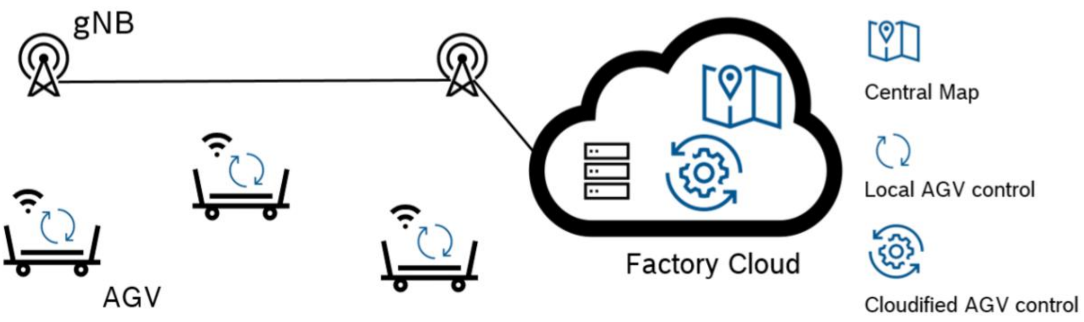

Figure 4: High-level illustration of the Cloud-based mobile robotics use case

The high-level visualization of the Cloud-based mobile robotics use case is illustrated in Figure 4. In the final realization, we implemented a hybrid solution where the use case is realized by using two AGVs, a research one and a commercial one, to demonstrate different aspects of AGV cloudification. It means that a suitable research platform (called *Research AGV*) is used to show how we can leverage 5G and cloud technologies to enable novel collaborative control solutions based on the cloud-native realization of the AGV control. The certified commercial AGV platform is connected over the 5G system to show the benefit of collaborative knowledge collected in the factory cloud (e.g., using the common map for AGV trajectory planning). We refer to this *as 5G-enhanced commercial AGV*. The APIs of the commercial platforms are typically closed, but we could still change the connection technology to 5G, as well as apply customized, cloud-based high-level control instead of their legacy software packages. On the other hand, commercial platforms are better suited for testing in real factory environment as they have gone through all required certifications and can be tested without deeper risk assessments of planned test cases.

As the final choice, we selected a MiR 100 as the commercial platform and a custom AGV based on the HEBI Mobile Base [HEBI] platform as the research AGV. We have a detailed description of the selection of robot platforms, necessary sensors, and other equipment and the implementation of the hardware and software architectures in 5G-SMART Deliverable D4.3 [5GS20-D43]. In Figure 5, we illustrate the functional architectures of the state-of-the-art commercial AGV (left side), the 5G-enhanced commercial AGV (middle), and our custom-built Research AGV (right side). As the figure shows, more and more components are offloaded from the hardware device to the local factory cloud as we move from the commercial AGV towards the Research AGV. The latter one has almost all software blocks implemented in the cloud except the drivers of the motors and sensors since they are attached to physical components. Even the safety functions are realized in the cloud, but they are also kept onboard to be in line with the safety regulations of the factory.

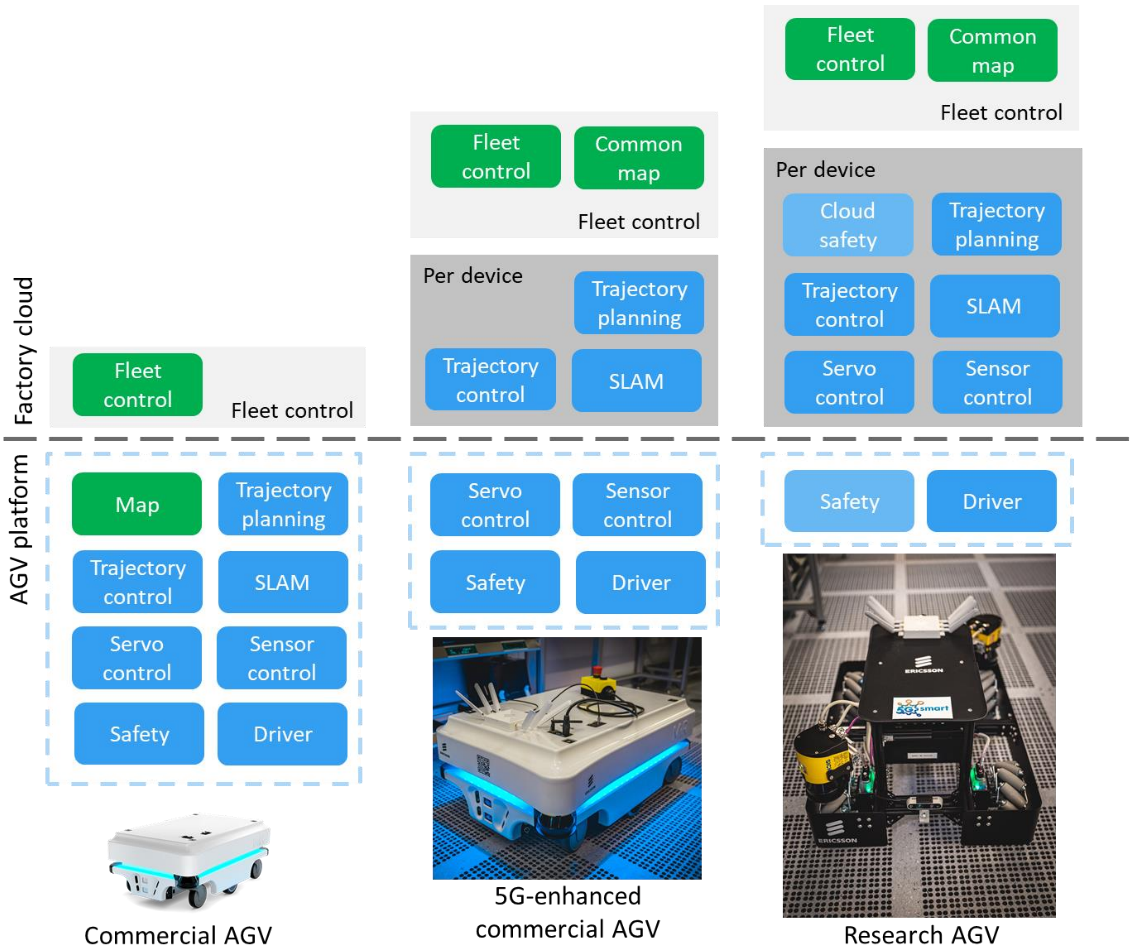

Figure 5: The functional architecture of the state-of-the-art commercial AGV (left side) compared to the architectures of the 5G-enhanced commercial AGV (middle) and the custom-built Research AGV (right side) applied in the final use-case implementation and deployed in the factory

3.2 Industry goals

In this section, we discuss all KPIs that have been proposed and evaluated for the Cloud-based mobile robotics use case and how they relate to our industry goals introduced in Section 2.2, while Section 3.3 details the definition of the KPIs and the methodology used to calculate those KPIs.

Flexibility in transportation with mobile robots is a key enabler to increasing efficiency. We introduced several KPIs to quantify the benefit of flexibility in relation to production, such as measuring the time needed to execute a mission compared to when no coordination is implemented among the robots. We measure how quickly an AGV can react to changes noticed by the other AGV. Furthermore, we have put sudden obstacles in the way of the robots during motion to test how they handle unexpected events that can easily happen in real scenarios as well. During the whole test period, we measured the number of emergency stops, lost network connections, and alarms since they can negatively affect productivity.

Quality is measured primarily by network related KPIs (latency, duration of connection loss, handover). However, safety aspects are captured in measuring the time and distance needed for an AGV to stop when a potential collision is anticipated. This reflects the performance of our cloud safety feature without activating the onboard safety mechanisms.

Sustainability aspects are measured in multiple KPIs. We provide an in-depth analysis of the battery lifetime and power consumption KPIs, which compare the values with and without cloud control. In addition, we measure the gain of robot collaboration which significantly contributes to energy-efficiency, since it can reduce travelled distance and mission execution times by a large extent.

3.3 Validation KPIs

We defined 16 KPIs (some of them even have multiple variants or sub-KPIs) to validate the capabilities and performance of the 5G technology in a real factory environment. We summarize the KPIs defined for this use case in Table 2, which are in priority order defined by Bosch showing the most important KPIs at the beginning of the table.

KPI name	Industry goal	KPI short definition
Collisions (#/h)	Safety, Mobility	Number of AGV collisions per hour with stationary or mobile obstacles (humans, other AGVs, trolleys).
Emergency stops (#/h)	Productivity, Safety	Number of AGV emergency stops per hour.
Lost connections (#/h)	Productivity, Safety	Average number of lost connections per hour, which results in a stop of the AGV.
Average duration of a lost connection (ms)	Productivity, Quality	Average duration of a lost connection.
Alarms (#/h)	Productivity, Safety	Number of alarms due to a lost network connectivity (for e.g., duration of 1s) per hour.
Battery lifetime (h)	Sustainability, Productivity	Battery lifetime of the research AGV with and without offloaded control stack.
Power consumption (W)	Productivity	Power consumption of the research AGV with and without offloaded control stack.
Successful handover rate	Quality, Mobility	Number of successful handovers (i.e., without visually observable stop) divided by the total number of handover events between the two 5G radio cells.

Reaction time (ms)	Productivity, Flexibility	Average time needed to update the common map (obstacles in the common map seen by another AGV) and change the route of an AGV accordingly (different options proposed and described below).
Relative time to reach the destination (s)	Productivity, Flexibility	Average relative time taken by an AGV to reach its destination in an environment with obstacles and multiple paths.
Interactions (#/h)	Flexibility	Number of interactions between the operator and the AGVs per hour.
Interaction duration (s)	Flexibility	Average duration of an interaction between the operator and an AGV.
Stopping distance (cm)	Safety	Average distance the AGVs travel due to the delay between an AGV detecting an obstacle and stoping (different options proposed and described below).
Stopping time (s)	Safety	Average time for an AGV to detect an obstacle and stop.
Round-trip time (RTT) end-to-end (ms)	Quality	Average RTT measured end-to-end (between the robot and the factory cloud) by using ping messages.
Control takeover time (ms)	Flexibility, Safety	Average time for the secondary control stack to take over the control in case of the Research AGV.

Table 2: Summary of the KPIs of the Cloud-based mobile robotics use case

3.3.1 Evaluation

As shown in Table 2: Summary of the KPIs of the Cloud-based mobile robotics use case KPIs are also assigned to Industry goals as discussed in Section 2.2. Below we describe each KPI in more detail and disclose their measured values where the KPI is independent of the validation scenarios. Otherwise, the KPI results are detailed in the sections of the corresponding validation scenarios.

Collisions (#/h): 0

We count how many times any of the two AGVs collides with any object during one hour of the validation tests on the shopfloor. We run tests during a week resulting in **27 hours** of net operation time. We note that the tests were conducted inside the production area with humans and other robots, so there were opportunity for such events to happen. Furthermore, testing means the execution of the nine validation scenarios defined by Bosch (described in Sections 3.4 – 3.12) including obstacle avoidance and AGV route intersection. During the 27 hours of testing, there were zero collisions. We run safety functions in the cloud (in case of the Research AGV), as well as on the AGVs, which can activate immediately if, for some reason, navigation fails to replan the trajectory to avoid potential collision.

Emergency stops (#/h): 0

We had only one emergency stop of the Research AGV due to a human mistake. The charging cable of the Research AGV was misplaced after a recharge, and it slowly moved into the safety zone of the AGV and stopped the whole platform. We needed to approach the platform and place the cable correctly into the cable support. After that, the robot continued its active trajectory. We note that the robot platforms faced extremely high vibration on the perforated tiles of the floor where the air is sucked continuously to maintain the cleanroom environment. We discussed this with Bosch and decided not to report this case as a valid emergency stop case. Therefore, this sample is not counted in the emergency stops KPI, so we ended up with zero emergency stops for the whole 27 hours.

Lost connections (#/h): 0

This KPI counts the average number of lost connections per hour, which results in a stop of the AGV. We defined a threshold for the connection loss to be 100 ms. This number is sufficiently low when it comes to using human visual senses to judge if there was a visually observable stop of the AGV. This means that a short glitch of up to 100 ms in the motion of the robot is not considered to be significant from the industrial application perspective. Furthermore, this number is also aligned to the measured end-to-end round-trip time (RTT) of the system (see RTT KPI below). If the measured delay of the ping message reaches or exceeds 100 ms, we increase our lost connection counter (the ping messages are sent on a millisecond rate). During the whole testing period (27 hours) there was no such lost connection reported by the measurement script in the database, nor did we experience glitches in the motion of the robots.

Average duration of a lost connection (ms): 0 ms

The idea behind this KPI was to also measure the duration of the connection losses in case they occur. This is also measured by the measurement script that we used to calculate the Lost connection KPI. Since there was no connection loss during the time of the validation tests, as we explained above, the duration is considered to be 0 ms.

Alarms (#/h): 0

This KPI intends to measure the number of alarms due to a lost network connectivity per hour. We define an alarm when the duration of the connectivity loss exceeds 1 second. These means that 10 consecutive lost connection events are defined as an alarm. The reasoning behind is that 1 second outage in the connection can have significant impact on the robot motion, which should be reported and analyzed in more detail. Since we did not measure any connection loss above 100 ms, the number of alarms during 27 hours of operation was 0.

Battery lifetime (h): 1h:40m (cloud control) vs. 1h:38m (onboard control)

This industrial KPI measures the battery lifetime of the research AGV with and without offloaded control stack. It is interesting to examine if there is any difference in battery lifetime of the research platform when all control components are running onboard versus in the cloud. We may use more CPU resources (and thereby more power) onboard since we must execute all control components on the integrated mini-PC of the robot. On the other hand, less traffic is transmitted over the 5G modem, because no low-level commands or sensor data leave the platform in this case, thus, the power consumption of the 5G modem may decrease.

When the control stack runs in the cloud, the situation is the opposite. We do not use the mini-PC of the robot platform for the control. Still, all sensory data and per-servo, low-level motion control information are transmitted over the 5G network, thereby, the onboard computational requirement is minimized. Still, the 5G modem is used more extensively. Therefore, it is interesting to measure what these two different deployment options mean in battery lifetime.

Since we implemented the application based on Docker containers, the full control logic could be deployed on the Linux-based mini-PC of the research AGV with only minor configuration changes. We used only the battery integrated into the central box of the research AGV (external batteries can also be added for extended uptime). A script was used during the tests to run the AGV continuously back and forth between two locations without any stop until the battery fully depletes. We reached 1 hour 40 minutes battery lifetime on average when we used full cloud control while running all components onboard resulted in 1 hour 38 minutes uptime.

We can conclude that it is not a significant difference, so there might not be notable instantaneous power saving achieved. However, we note that during the planning of the system, the power consumption was not a design criterion, so we did not pay particular attention to that when selecting the HW components of the robot platform. One example is the use of the onboard computer. We could have easily used a constrained device instead of a full-featured mini-PC of Intel NUC, such as a Raspberry Pi or an Orange Pi with very low power consumption. Furthermore, we did not use the camera input during these measurements, since it was not required by the demo scenarios at the Bosch factory, which would increase the power consumption of the mini-PC some extent and may result in larger difference in battery lifetime of the two deployment options.

We also note that we can still achieve significant battery lifetime savings by offloading the control stack of the AGV in the factory scenario, since the overall travelled time and distance are significantly reduced by using the common map in the cloud. Thus the energy requirement of executing a mission is significantly reduced. We measured this effect in a separate KPI called “Relative time to reach the destination” that will be detailed below.

The Power consumption KPI is also evaluated below to understand better the measured battery lifetime values of the Research AGV.

Power consumption (W): 86.5W (cloud control) vs. 87.7W (onboard control)

The Power consumption KPI is introduced to evaluate the power consumption of the Research AGV and to separate the measured values per component. We used a Texas Instruments INA226 ultra-precise current/voltage/power monitoring board that we integrated into the Research AGV. All measurements were executed for 15 minutes collecting instantaneous power consumption values in every 100 ms. In Table 3: Power consumption of the research AGV when the full control is deployed onboard or in the cloud, we illustrate the results of the power consumption measurements.

Scenario	Avg. power consumption (W)	Control and idle diff. (W)	Onboard and cloud control diff. (power saved by cloud control in W)
Idle robot	83.58		
Onboard control	87.7	4.12	
(after moving)	113.36	29.78	
Cloud control	86.5	3.01	1.2
(after moving)	112.5	27.1	0.86

Table 3: Power consumption of the research AGV when the full control is deployed onboard or in the cloud

Idle robot means that we power on all HW components such as the two SICK safety laser scanners, the MikroTik RouterBoard network switch, the 5G modem, the HEBI servo motors, and the Intel NUC mini-PC. That is, the robot is up and running and ready to be controlled. In this state, the average power consumption was 83.58W (the voltage of the system is 36V). When the control stack is deployed and executed onboard, the power consumption increased by 4.12W ended up with 87.7W in total. In another measurement, we removed the control from the onboard mini-PC of the Research AGV and executed the full control stack in the cloud, which resulted in total power consumption of 86.5W on average. Compared to the onboard control case, the difference is only 1.2W in average, which yields to a minor 2 minutes difference in the battery lifetime (1h:40m and 1h:38m with onboard and cloud control, respectively).

We experienced an interesting phenomenon during the measurements (also shown in Table 3 in the "after moving" rows). The power consumption of the system is different after the first trajectory is executed. After that, the total power demand of the Research AGV increases significantly, which is most probably due to the servos being in a different state than the initial one, but without having the servo firmware implementation it is not possible to fully validate this hypothesis. We experienced that when a trajectory finished, all servos are braked (like executing 0.0 rad/sec velocity commands), that may cause the increased power consumption. The difference in the average between the onboard and the cloud case is roughly the same (0.86 W).

Based on the above results, it is interesting to investigate the power consumption for each hardware component that is relevant and can be measured or calculated separately.

Per component	Power consumption (W)	Difference in W (without Sick sensors + MikroTik)
Sick sensors + MikroTik	41.53 (avg)	
Mini-PC idle	48.94 (avg)	7.41
full load	105.8 (avg)	64.27
5G modem idle	50.27 (avg)	8.74
full load (UL heavy)	57.9 (avg)	16.37
Servo motors	86.5 – 310	Up to 223.5

Table 4: Power consumption of the research AGV per HW component

We connected the power measurement board to the batteries and powered on the HW components one-by-one and the results are shown in Table 4. The SICK sensors and MikroTik switch components are directly connected to the batteries in order not to operate the AGV without safety, so they cannot be separated and switched off. They are the baseline during the measurements, so they also contribute to the total power consumption when other HW elements are tested. The per component consumption can be calculated by deducting this value from the total power consumption measured in different cases. The baseline scenario (when the SICK sensors and the MikroTik RouterBoard are switched on) consumes 41.53 W on average. Then, the mini-PC was also powered on, resulting in a total consumption of 48.94 W. This means that the mini-PC needs 7.41W in idle state as calculated in the "Difference in W" column of Table 4. We started a dummy process using the full capacity of the CPU which increased the power consumption by 64.27 W (105.8 W in total). After this measurement, switching off the mini-PC and powering on the 5G modem resulted in 50.27 W in total, which means 8.74 W consumption in idle mode of the modem. To test the maximum consumption of the modem, we started a speed test to a speed-test server installed locally onto the edge cloud servers connected to the 5G system directly. This traffic used the full bandwidth in uplink, as well as in the downlink, but the power consumption was measured during the period of the uplink traffic delivery since the traffic pattern of the use case is also uplink heavy. In this case, the 5G modem consumption was 16.37 W (57.9 W in total system).

The last row of Table 4 shows the power consumption of the case when all HW components are used, including the servo motors. Since the consumption depends heavily on the actual movement of the platform, these values have high variability up to 223.5 W in our actual test scenario that we implemented for the measurements (going back and forth between two locations on the track).

Power consumption KPI conclusions

As we already mentioned in connection with the Battery lifetime KPI, there is no significant battery lifetime and instantaneous power savings achieved with cloud control in this particular realization of the research AGV (which was not optimized for energy-efficiency). By quantifying the different parts of the Power consumption KPI, we can estimate how much power could be saved in this setup. In case of onboard control, a more complex application may eat up all the computational resources of the mini-PC increasing the power consumption by ~64 W. We can then assume an idle 5G modem consumption, since basically all data remain onboard. In the cloud control case, we may increase the consumption of the modem to the maximum in the worst case resulting in only ~16 W power increase, so there is some room left for decreased power consumption (~48 W) with cloud control in this realization.

Furthermore, as highlighted above, the four servo motors are responsible for ~80% of the total power consumption, which limits the theoretically achievable instantaneous power savings. However, the cloud-based control can still be much more efficient in battery lifetime than onboard control since we can significantly reduce the travelled distance and time, as we illustrate later in the common map-related validation scenario and its corresponding KPIs in Section 3.8. Also, the overall weight of the cloud-controlled platform is significantly less (with the same loading capabilities), which can further increase battery lifetime. For instance, the weight of our robots are 40 kg and 77 kg for the Research and the Commercial AGV, respectively.

Successful handover rate

This KPI measures the number of successful handovers between the two 5G cells inside the cleanroom divided by the total number of handovers between those two cells. We define a handover event as successful if there is no visually observable robot stop. We proposed two options. One is to run the whole control stack in the cloud, including navigation and localization components to steer the robot. We applied joystick control in the cloud in the second option, disabling the navigation and localization components. Section 3.10 explains the two options mentioned above as two validation scenarios that we used to measure this KPI.

Reaction time (ms)

This KPI is defined to evaluate the efficiency of cloud control over 5G by measuring the average time needed to update the common map with obstacles seen by the other AGV and change the route of the other robot accordingly. We proposed two versions of this KPI where the *"Reaction time – end-to-end"* measures the time taken from putting an object in the way of a robot until the other one reacts to that, i.e., it includes the full cycle end-to-end from sensing till the other robots execute its reaction. In the second version, called *"Reaction time - processing"*, we measure the time taken between updating the common map (i.e., obstacles introduced in the common map by another AGV) and changing the route of an AGV accordingly. More specifically, we measure the time from the common map update until the new control command (as a reaction to the map update) is sent to the other AGV. The sensing and command execution parts are not included in this version. It focuses on the processing part happening in the cloud. Both KPIs are evaluated in detail in the validation scenario *"Common map without alternative route"* and described in Section 3.8.2.

Relative time to reach the destination (s)

The average time taken by an AGV to reach its destination in an environment with obstacles and multiple paths is measured by this KPI. We compare the values with and without using the common map to measure the efficiency improvement we can achieve in the corresponding validation scenario described in Section 3.8. We would like to evaluate the gain in terms of travelled time of the robots in a particular mission while they use the common map feature compared to the state-of-the-art solutions when coordination is not available. A similar KPI could be defined for the travelled distance with and without the common map, but here we focus on the time as requested by Bosch. Since the actual values of this KPI depend on the validation scenario, we disclose the KPI results in Section 3.8 after the description of the scenario.

Interactions (#/h): 0

The Interactions KPI counts the number of interactions with the AGVs, i.e., how many times an operator needs to approach the robot which is not able to move otherwise. As we mentioned at the Emergency stops KPI, once we needed to approach the Research AGV due to a misplaced charging cable that caused emergency stop, which would mean one interaction during the 27 hours of testing. However, as discussed with Bosch, we ignore human mistakes, and therefore the number of interactions is reported as zero.

Interaction duration (s): 0

Interaction duration measures the average duration of interactions between the operator and any of the AGVs. Similar to the previous KPI, disregarding human-caused errors, we have not experienced any interaction thus the duration was 0 second. Just for the record, we spent 19 seconds at the Research AGV to realize and correct the problem of the misplaced charging cable being in the safety zone of the robot.

Stopping distance (cm)

Stopping distance is one of the safety-related KPIs. It measures the average distance the AGV travels from the moment an object is put in the way of the robot until the robot platform completely stops (without extra load carried). We defined a specific validation scenario for the evaluation that is defined in Section 3.11.2.

Stopping time (s)

This KPI is also connected to safety, as well as to the previous KPI, since it measures the time the AGV needs to fully stop from the moment we throw an object in the way of the robot. This KPI is evaluated by executing the same validation scenario as used for the Stopping time KPI in Section 3.11.2.

RTT end-to-end (ms): 18.3 ms

We measured the average RTT using the ping tool, which sends a request message to one of the servo motors in every 100 ms from a monitoring process running in the same cloud environment next to the application. We selected 100 ms to have manageble size of measurements for the 27 hours of testing. We note that we used ping messages on millisecond rate to detect connection loss in the network (see the Connection loss KPI), so here we loaded only one RTT value per 100 ms in the measurement database and used those values for the evaluation of the end-to-end RTT. The reply message of the servo motor was sent back to the same monitoring process where the delay is calculated. We note that the measured 18.3 ms includes the two-way delay of all components end-to-end, from the cloud-based application to the robot servo motor and back. This means that not only the delay of the network is represented, but the latency of the cloud part, and the packet handling delays of the 5G modem and the servo motor. In the implementation of the use case, we applied IP-IP tunneling between the cloud server and the mini-PCs of the AGVs to allow multiple endpoints with separate IP addresses to be connected. The encapsulation and decapsulation of a packet add 2.1 ms extra latency, which is also included in the measured value.

Control takeover time (ms)

Control takeover time shows how fast the secondary control stack can take over the control in case the primary control instance fails. We run two instances from the whole control stack for both AGVs as detailed in 5G-SMART's Deliverable D4.3 [5GS20-D43]. We measured the takeover time by intentionally killing the active control stack of the Research AGV. The system detects the failure and switches to the secondary instance, and it becomes the active controller of the robot. We measure the time from the moment we kill the active instance until the first servo feedback message arrives to the motion controller that belongs to the secondary controller validating that the secondary control instance is the active control stack. We implemented the servo control such that the servo motor injects the ID of the active controller into the feedback message. For instance, controller ID #0 in a feedback message means that the control command was generated by motion controller #0, thereby we can determine the active control instance. The details of the validation scenarios and the results of this KPI can be found in Section 3.11.

3.4 Validation scenario 1: Map creation

3.4.1 Description of the scenario

In this scenario we show how AGVs scan the environment and create a map that is uploaded in the cloud and used as the basis of the common map. There is no KPI defined directly to the map creation scenario, but there are other KPIs validating the usage of the common map, such as the *Reaction time* and *Relative time to reach the destination*. The idea is to show the feasibility and validate that the robots can create the map by themselves without using existing maps or floorplans.

3.4.2 Validation results

We show a snapshot of the map creation process in Figure 6: Creating the map of the environment where the MiR AGV scans one of the corridors of the cleanroom and creates a map in real-time using its lidar measurements. A video clip of the map creation can be watched here as part of the demo video we recorded during the validation. Black, white, and gray pixels mean object seen, no object (free space), and unknown area, respectively. After scanning the whole area, the map is uploaded into a map server that resided in the cloud. This is a two-dimensional map based on the 2D laser scans, but we add several layers to this map to constitute a full-featured map, including dynamic obstacles, costmap layers, as well as 3D information from the depth camera, as illustrated in Figure 7. Costmap layers are constructed based on sensory data (laser scan, camera, static map). Costmap assigns "cost values" to each pixel (location) of a map that is used by the planning and navigation algorithms that try to find a minimum cost trajectory on the (cost)map.

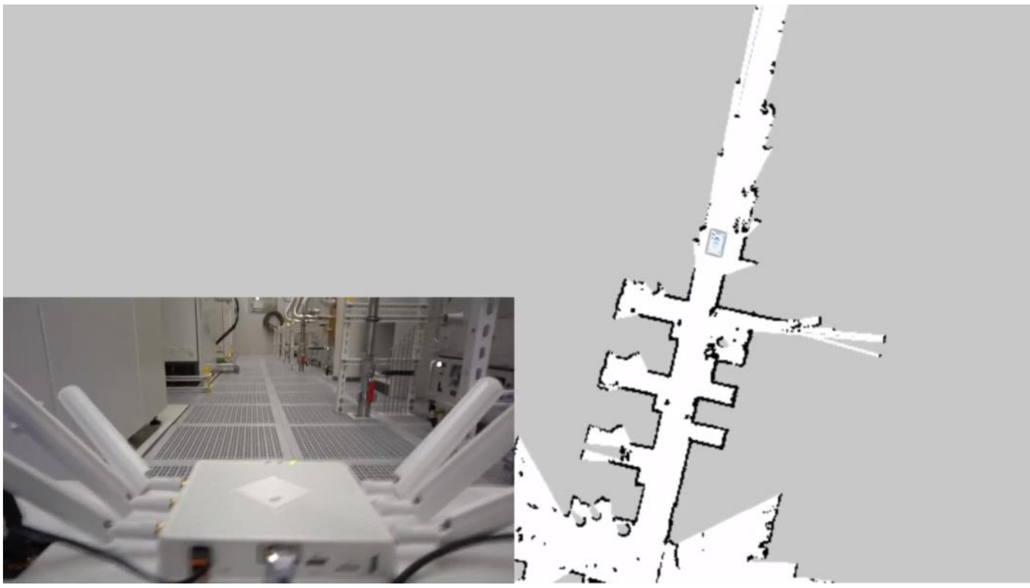

Figure 6: Creating the map of the environment

Both AGVs continuously update the common map in real-time, and they navigate based on that, if any robot discovers an obstacle, it adds that to the map, and the other AGV can use this information in real-time. We note that our fleet manager is aware of the exact robot positions and thus we can handle the other robot as a distinguished obstacle. For instance, we can neglect adding the other AGV as obstacle to the common map.

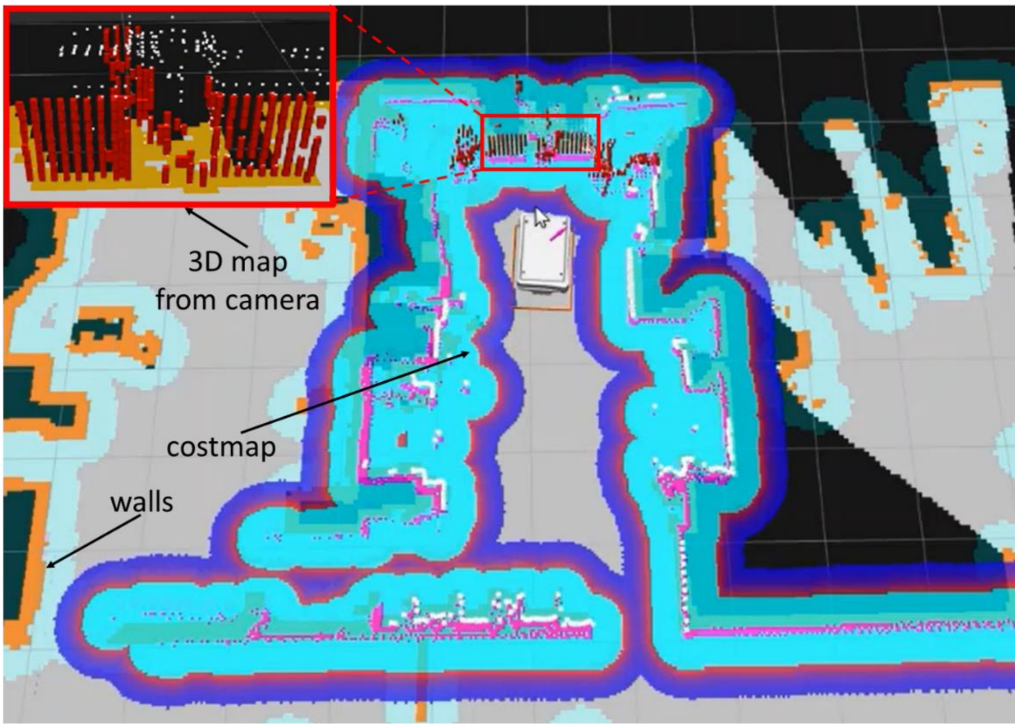

Figure 7: Illustration of the map with additional layers on top of the static map

3.4.3 Discussion

The conclusion of this validation scenario is that it is feasible to create a map of the environment by the robots in real-time, which, extended with additional information, can be used during navigation.

3.5 Validation scenario 2: Obstacle avoidance

Obstacle avoidance scenario shows how responsive the navigation of the AGVs is to sudden dynamic obstacles such as humans. We want to challenge the system since the number of collisions and emergency stops are the top priority KPIs for Bosch plant management. We illustrate a snapshot of the validation process in Figure 8. In many trials, the AGVs replan their trajectories and react to the obstacle in an agile way without having emergency stops or collisions, as can be seen in the obstacle avoidance part of the demo video.

Figure 8: Illustration of obstacle avoidance when human suddenly steps in front of the robot

3.6 Validation scenario 3: Route intersection

In this scenario, the trajectories of the AGVs cross each other. It is of interest how the robots handle this situation when they arrive to the route intersection at the same time. Due to the centralized control in the cloud, our fleet manager is aware of the exact location of each robot and can solve this problem in a coordinated manner. We define a restricted zone around the route intersection, and a robot can enter the zone only if there is no other robot inside. This means that the AGV that arrives first at the intersection has priority, and the other one must wait until the first one leaves the restricted area. A snapshot of the validation is illustrated in Figure 9. One execution of this scenario is also available in the demo video.

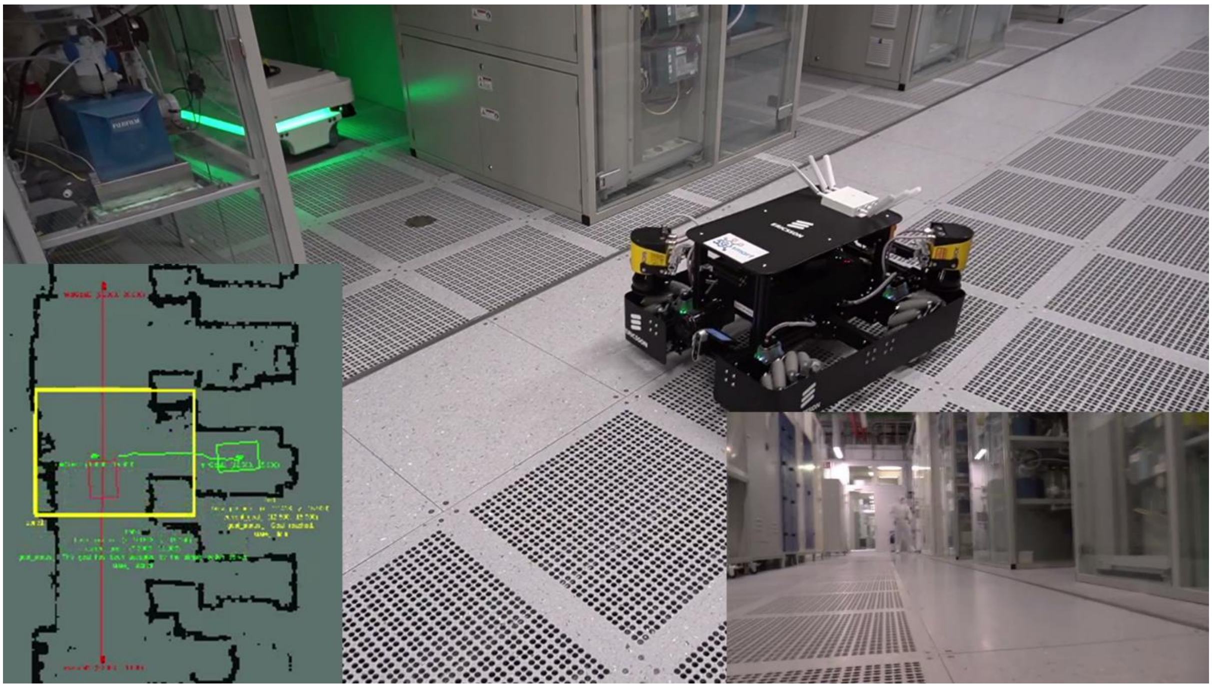

Figure 9: Snapshot from the validation of the route intersection scenario. The Research AGV (red rectangle on the map) entered the restricted zone (yellow rectangle) first, and MiR AGV (green rectangle) waits until the zone becomes empty again

3.7 Validation scenario 4: Navigating on narrow corridor

Another frequent problem in semiconductor factories is that the corridors are narrow, walled by large machines, and the AGVs must maneuver in this environment. In this validation scenario, we would like to show that coordination can help to solve problems when two AGVs must pass each other on a narrow corridor. We implemented a solution that detects this issue and gives priority to the commercial AGV and the Research AGV is instructed to move away from its way. Later, the Research AGV also continues its current trajectory. We note that the Research AGV maneuvers much better due to its mecanum wheels, therefore, it can quickly move out of the way of the commercial robot. That is the reason why we prioritized the MiR.

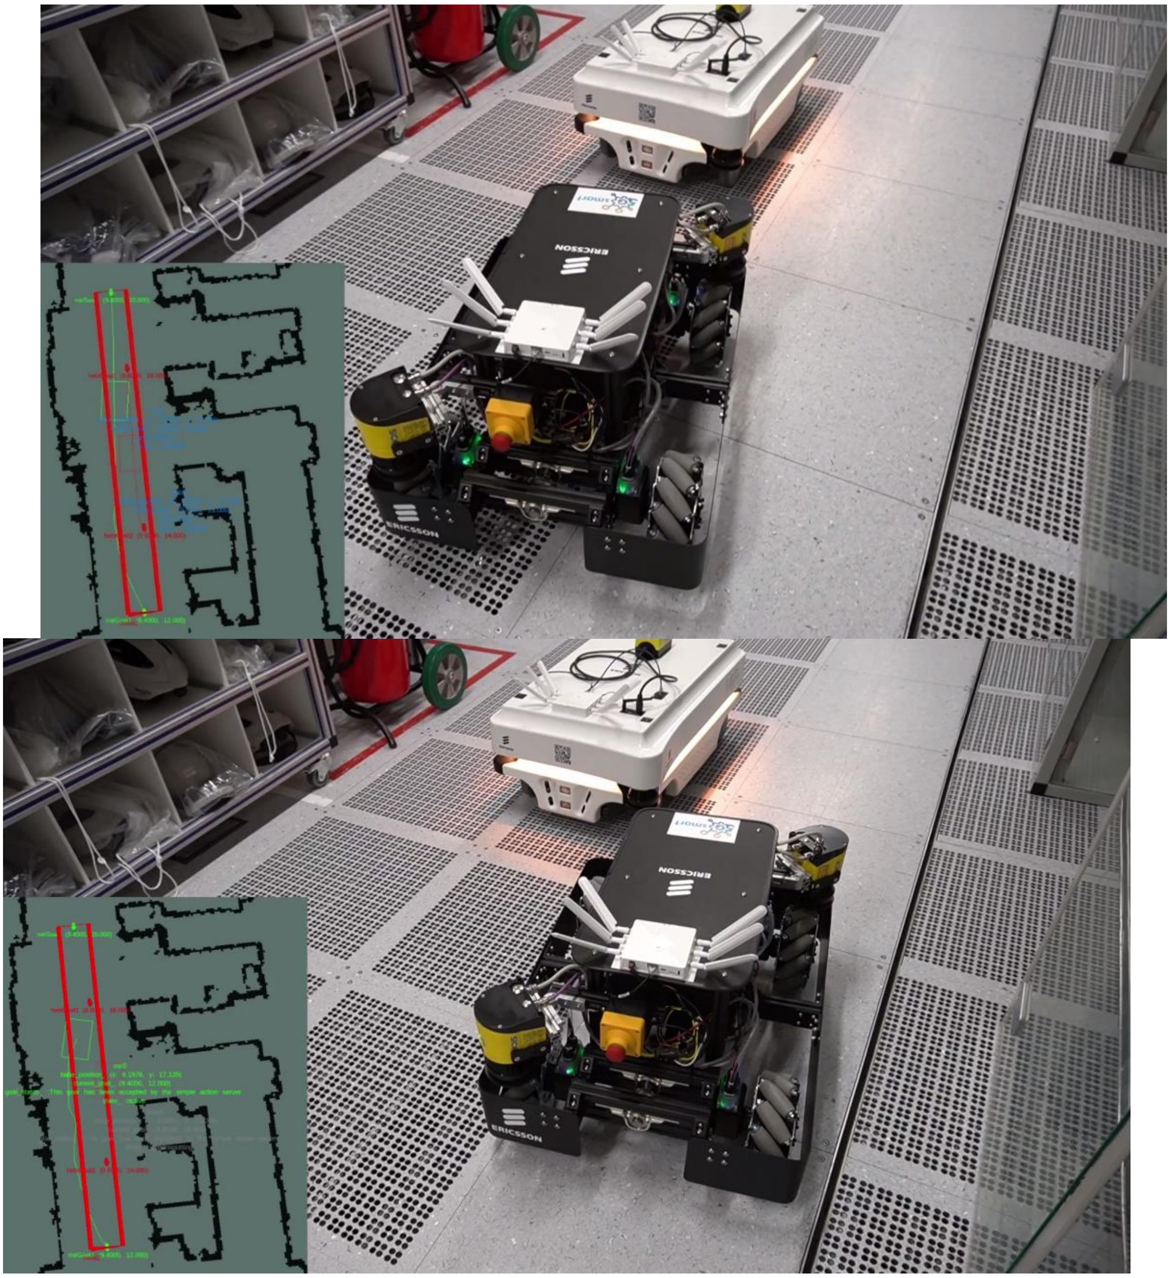

Figure 10: Snapshots of navigation on narrow corridors when AGVs pass each other. The Research AGV moves to the right to let the other AGV move. After the commercial robot left, the Research AGV continues its trajectory to the original destination.

The fleet control checks the distance between the two robots. If the central points of the AGVs are closer than 2 meters, the trajectory of the Research AGV is modified to move away from the way of the Commercial AGV as much as possible, so it has enough space to pass and continue its journey to the destination. A related video clip from the demo video can be watched here. Without this feature, both AGVs see the other one as a (static) obstacle and try to bypass that by running an obstacle-avoidance strategy locally on-board. It often ends up in moving left and right in an uncoordinated way making the bypass rather lengthy, or even unsuccessful after several trials.

3.8 Validation scenario 5: Using the common map

3.8.1 Description of the scenario

Using the common map is one of the most important validation scenarios to show the benefit of cloud-based control over 5G. The goal is to quantify how much we can save in mission execution time when the two AGVs share information about the environment in real-time via the common map in the cloud. We need to run the AGVs in an area where there is at least one alternative route to a target destination. Since such an area was not available at the time of the validation, we modified the map a bit to have two different routes (only on the map) to a location as you can see in Figure 11. We defined two starting points and two destinations for the two AGVs in the cleanroom as illustrated in the right side of Figure 11. Thereby, the AGVs can find alternative routes to the destinations in case the shortest paths are detected as blocked. In Figure 11, we illustrate the planned trajectory per AGV being the shortest path to the target locations by green and red lines for the Commercial and the Research AGVs, respectively. Since this route does not exist on the real floorplan, i.e., it is physically blocked, both AGVs have to replan. When the common map feature is not used, both AGVs start executing the planned trajectory until they detect the blocked route and then replan to find the alternative path. This is inefficient since both AGVs must get into the “dead-end road” and only then replan the trajectory to go back and follow on the longer way to the destination.

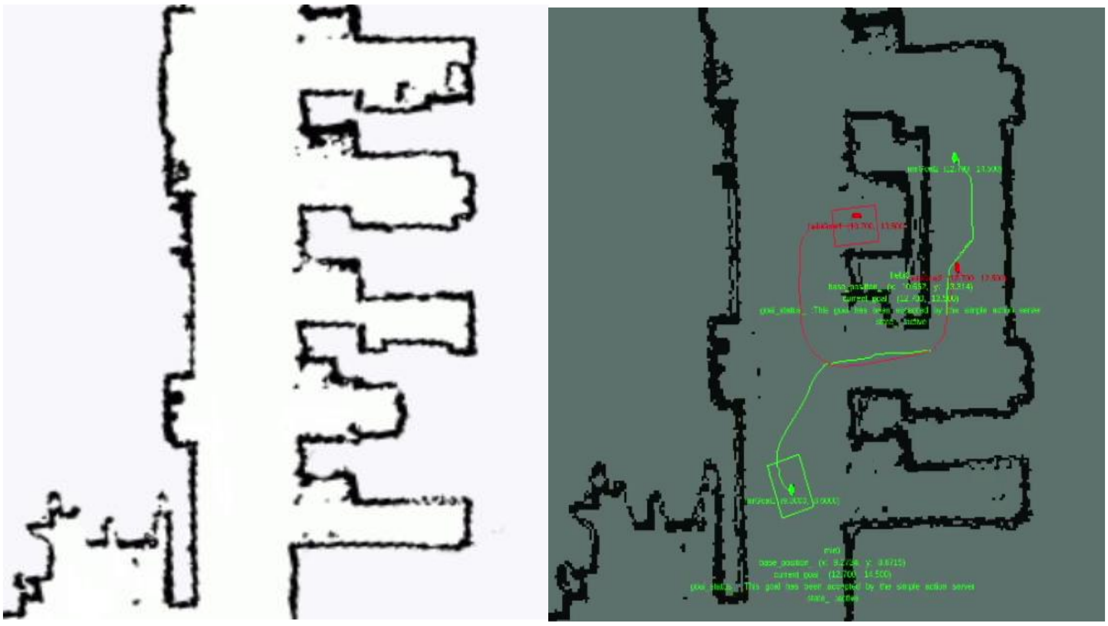

Figure 11: Left: original map segment. Right: modified map segment with two routes from the starting points towards the destinations that are denoted by green and red dots for the Commercial and the Research AGVs, respectively.

When we apply our common map solution, only one AGV needs to detect the obstacle. When doing so, it immediately notifies the other one by updating the common map in real-time, so the other robot can react and act according to the new situation. In this case, the Commercial AGV notices the obstacle first and updates the common map so the Research AGV can replan without going into the blocked side-corridor as shown in Figure 12, and thereby achieving significant improvement in mission execution time. The corresponding video clip can be found here.

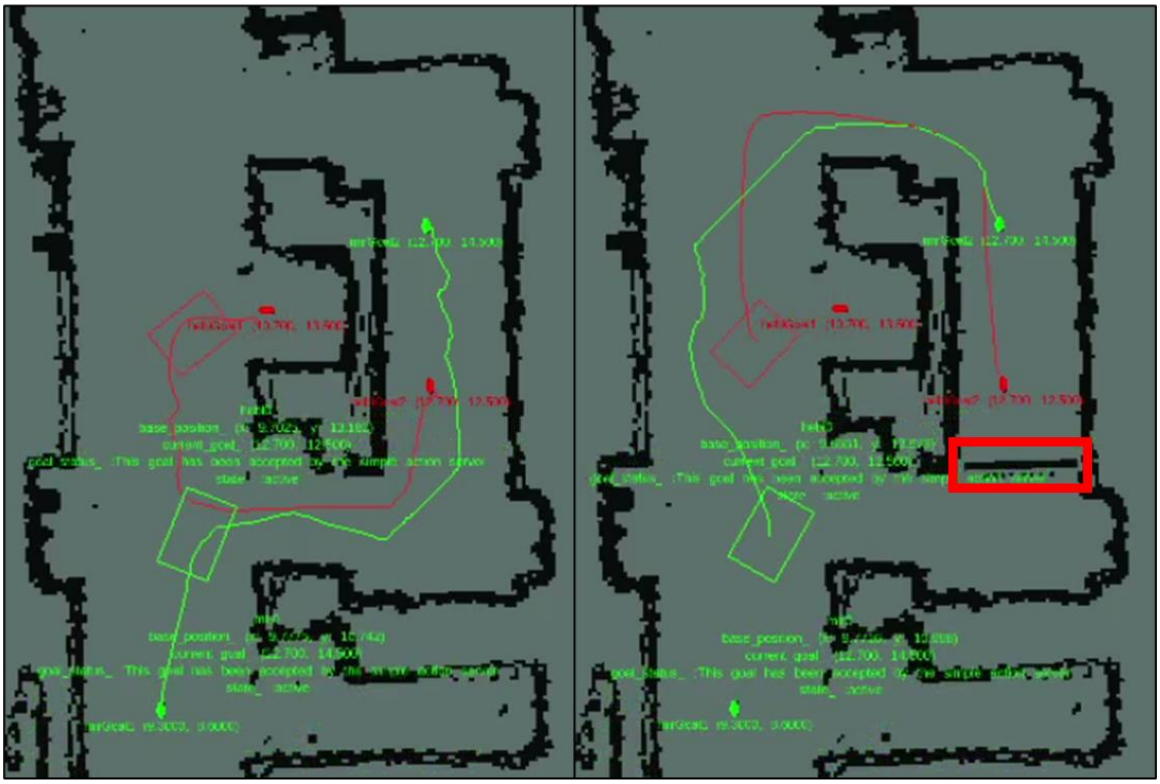

Figure 12: Illustration of using the common map. The Research AGV replans its trajectory based on the update from the Commercial AGV (blocked route highlighted by red rectangle), thus achieving significant improvement in mission execution time.

3.8.2 Validation results

We introduced the KPI “Relative time to reach the destination” in Section 3.3. This KPI is defined to measure the average gain in this particular scenario discussed above. We ran 10 executions of this scenario and averaged the results with and without the common map feature as listed in Table .

Relative time to reach the destination (s) Research AGV	
With common map	**Without common map**
28	39
30	40
25	44
27	43
38	46
31	38
27	40
27	48
26	45
28	39
28.7	**42.2**

Table 5: Validation results of the KPI "Relative time to reach the destination"

3.8.3 Discussion

Based on the average results, we calculate the average gain (1- 28.7/42.2) = 0.32. It means that, on average, we could speed up the mission execution time by 32% in this particular setup. We note that the gain largely depends on the layout of the validation scenario. Furthermore, the more is the number of robots using the common map, the higher is the gain.

During the development of the use case, we built a digital twin of the whole industrial application as illustrated in Figure 13. In this environment, we can test functionalities and even measure KPIs, since the hardware components are accurately modelled, and the same map is used that we recorded in the real factory. The graphical interface is also identical to what we used during the real executions.

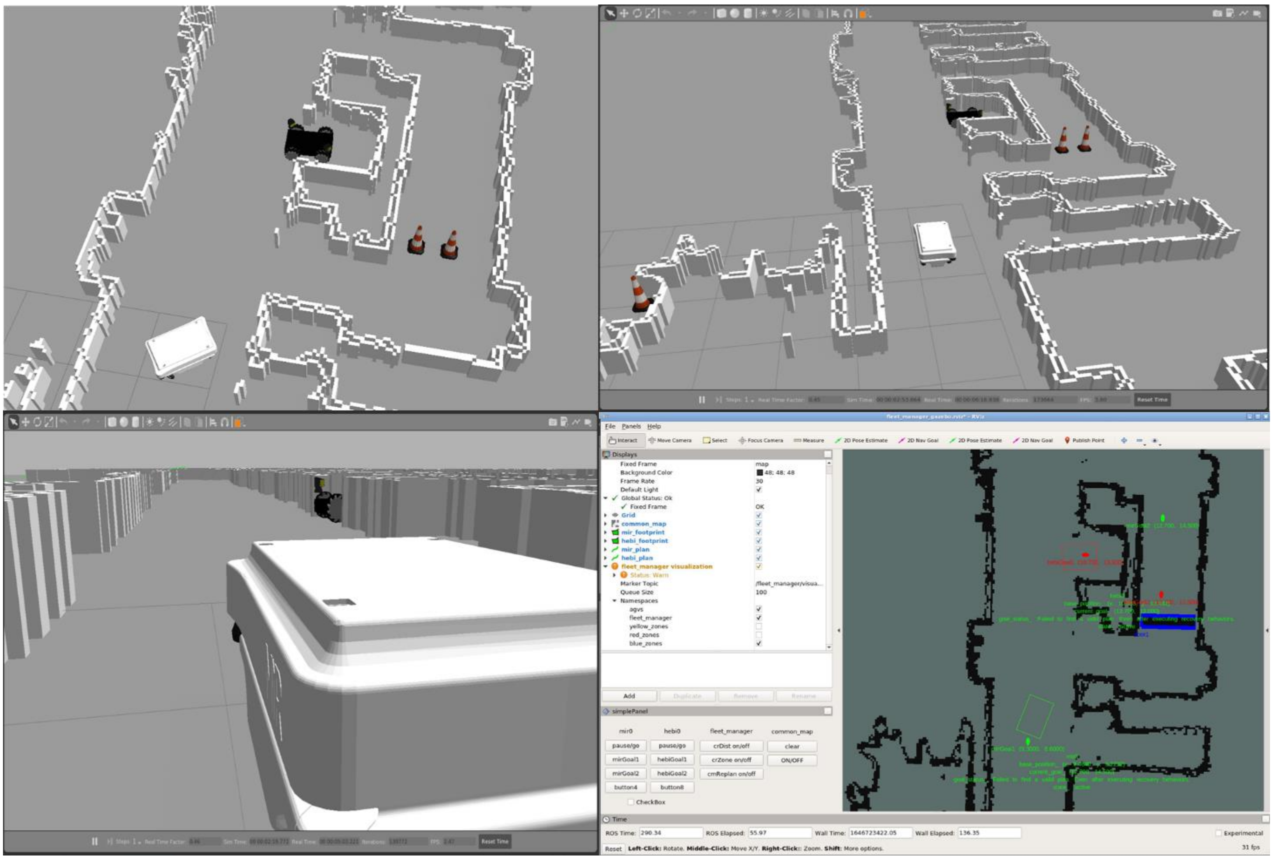

Figure 13: Snapshots from the digital twin of the AGV use case that we used during the implementation of the application.

Using the above-mentioned digital twin application, we run some simulations to evaluate the Relative time to reach the destination KPI. The average gain calculated the same way as in the case of real executions was measured to 33.8%, which is very close to what we measured with the real robots in the real factory.

3.9 Validation scenario 6: Common map without alternative route

3.9.1 Description of the scenario

This validation scenario is similar to the previous one, with the difference that now there is only one route to the destination. We block this path and let the Research AGV sense this change while we send a mission to the MiR via the blocked route. We measure the time needed from the moment the Research AGV detects the obstacle and updates the common map until the MiR reacts to the map change and stops. The validation scenario is illustrated in Figure 14.

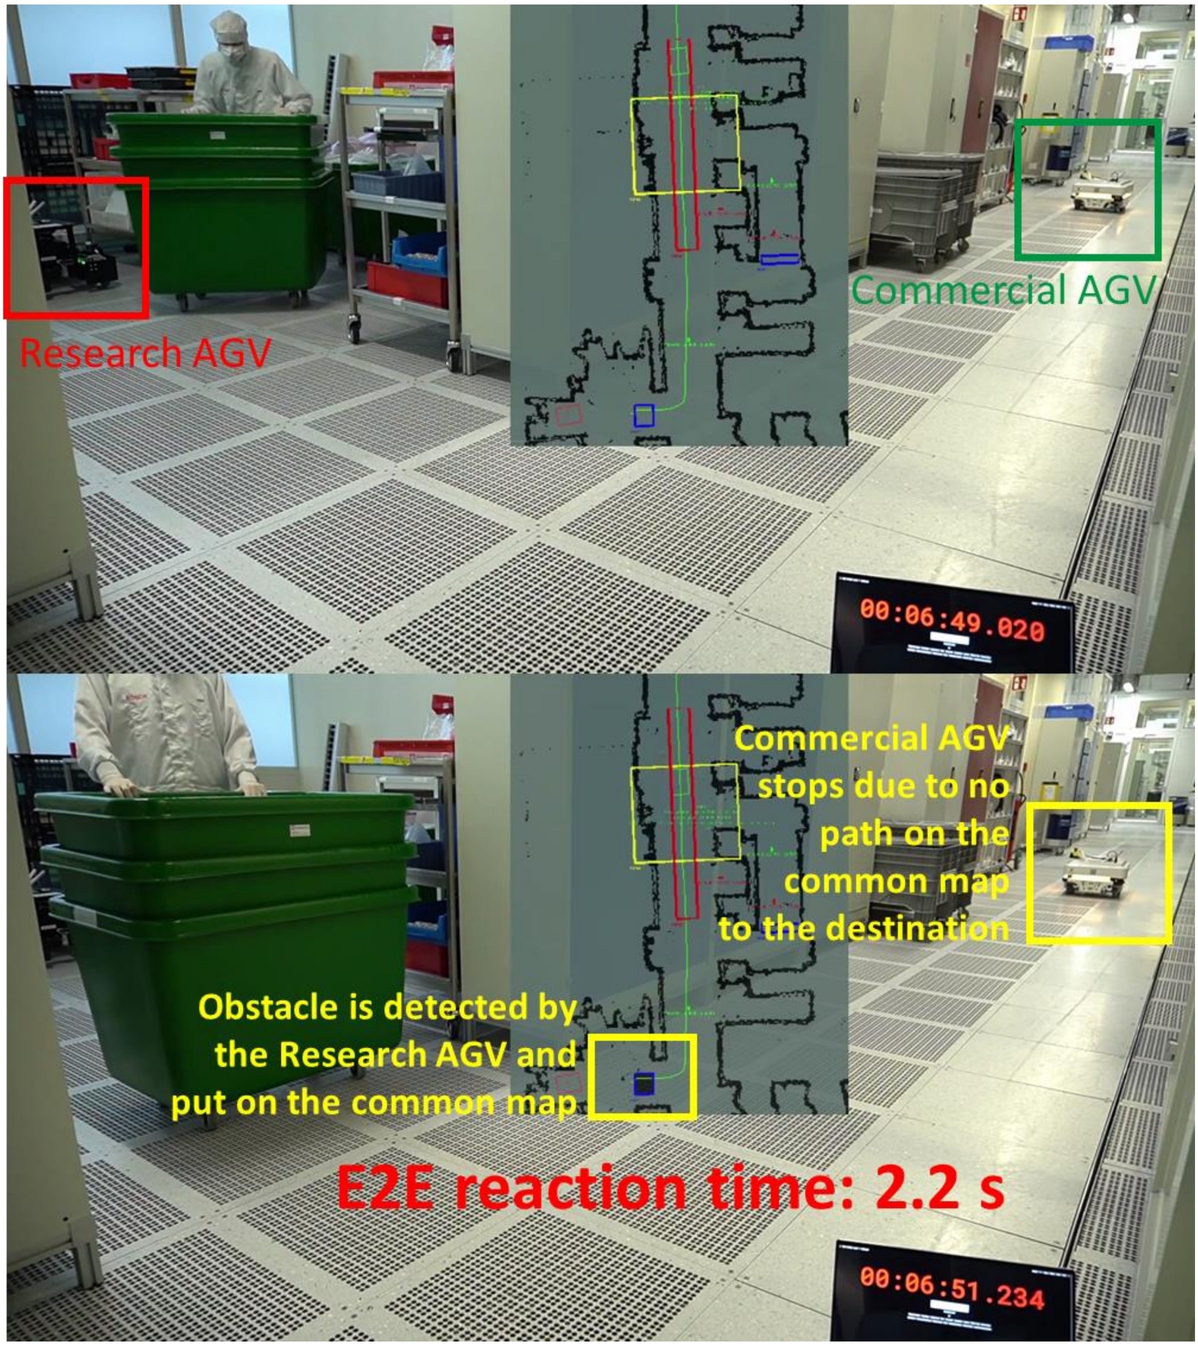

Figure 14: Illustration of the common map use when there is only one route to the destination of the MiR AGV. Since we block that route by a container, the MiR stops right after the common map is updated by the Research AGV which noticed the container. The reaction time end-to-end KPI had the value of 2.2 seconds in this particular execution.

3.9.2 Validation results

The Reaction time KPI is measured during the executions of this scenario. We defined two options as mentioned in Section 3.3. The *Reaction time – end-to-end* measures the full period from sensing till the other robot executes its reaction. The *Reaction time – processing* measures the time between the map update from the Research AGV until the controller sends the stop command to the MiR. We show the results of both options in 17 executions of this validation scenario in Table 6, including the average time in the last row.

Reaction time KPI	
Processing (ms)	**End-to-end (s)**
560	4.37
605	3.38
560	4.13
641	3.77
640	2.82
638	3.3
543	4.48
600	3.84
666	4.97
519	3.31
605	3.8
596	2.21
574	4.15
516	2.91
572	5.02
563	5.11
508	3.41
582.7 ms	**3.8 s**

Table 6: Results of the *Reaction time – end-to-end* and *Reaction time – processing* KPIs measured in 17 execution of the validation scenario *Common map without alternative route*

3.9.3 Discussion

The results show that if the common map is already updated by the Research AGV then we need only 582.7 ms on average to send a new velocity command to the MiR to stop. If we look at the full cycle from sensing the object and updating the common map, as well as waiting for the stop command to be executed, then we need 3.8 seconds on average. The sensing of the obstacle by the Research AGV takes half a second, while most of the time is consumed by the map update.

The full common map update procedure can be broken down into the following steps: each robot builds its local map based on the received common map plus its sensory data (lidar and optional camera). The local maps are updated and published periodically at the rate of 1.15 Hz (870 ms). Meanwhile, the robots localize and position themselves on their local map, and their footprint position is also published periodically at a rate of 2 Hz (500 ms). The fleet manager constructs and distributes the merged common map to the navigation components of the robots. To do so, the fleet manager waits for the local map updates and published robot positions from each robot, then processes and merges the map layers into the updated common map, considering the newly recorded changes (obstacles) detected by any of the robots. The common map is published at a rate of 1 Hz (1 sec). Finally, when the robots' navigation stacks receive the newly updated common map, they use it as the base layer to build up their own local maps incorporating their recent sensor data. The updated maps are used by robot navigation to carry on the path planning and execution, re-planning the local and/or global navigation plan if necessary. It should be noted that the local and common map publishing rates can be set according to the application's needs and capabilities. Higher publish rates result in shorter reaction times but require more processing power and data transmission rate and volumes. We also note that the above mentioned map processing happens in the cloud, it only requires the sensor input (lidar and optionally camera) from the AGVs. The data volume can range from several 100 kbit/s up to ~20 Mbit/s depending on the configuration of the camera stream.

3.10 Validation scenario 7: Operation during handover

3.10.1 Description of the scenario

In this scenario, we evaluate the performance of the 5G handover from the application point of view. As proposed by Bosch, the validation is done via human observation of the movement of the AGVs during handovers. We defined a KPI as introduced in Section 3.3 to calculate the successful handover rate as counting the number of successful handovers divided by the total number of handovers. We agreed to define a handover as successful if no long, i.e., visually observable stop is experienced in the motion of the robots. First, we measured the handover locations in the factory as illustrated in Figure 15.

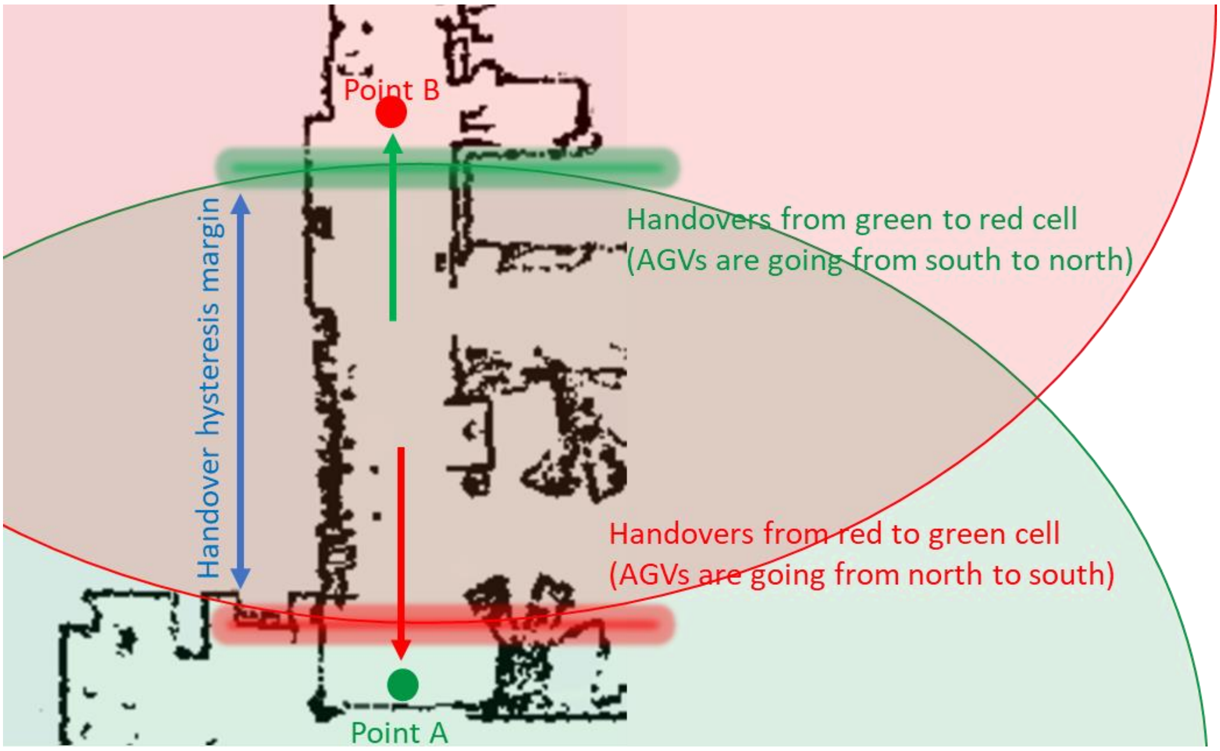

Figure 15: Locations of handovers and target destination points for the AGVs in the factory.

We moved the AGVs programmatically between two points, Point A in the green cell and B in the red cell as also shown in Figure 15 and carefully observed the handover executions (as well as recorded on video). The evaluation process is illustrated in Figure 16.

Figure 16: Snapshots of the evaluation of AGV operation when handovers happen during the motion of the robots (the Physical Cell ID (PCI) changed from 79 to 80).

3.10.2 Validation results

We proposed and evaluated two different control setups during the validation. In the first one, we run the full control stack from the cloud as usual. In the second setup, we applied joystick control for steering the robot, the navigation and localization components were disabled. This means that the joystick application was also running in the cloud and connected to the motion control component (resided in the cloud). In this setup we could evaluate the motion of the robots without the effects of navigation and localization, i.e., we could observe how the robots move when only the motion control steers them where the speed and the direction of the AGVs are determined by the joystick. Thereby, we could focus on the network performance since it is a more controlled environment. For instance, using joystick control, if a robot decreased its speed or stopped during handover, we could exclude the imperfections of the localization or the trajectory control as the cause of the problem. We altogether executed (and recorded on video) 31 handovers and we did not experience any considerable glitch in the movement of the robots in any of the two setups, thus the successful handover rate is evaluated to be 1.

3.10.3 Discussion

The main reason why we introduced the second validation setup with the joystick control was that the non-ideal behavior of the localization and navigation components can impact the smooth motion of the robot platform. For instance, if there are many dynamic objects on the path, the accuracy of the localization may decrease and thus the navigation may adjust the speed of the platform (typically decrease). We have seen such behavior several times when dynamic objects appeared on the route of the robots which were not on the map. We excluded those effects by using the joystick control to only see the effect of network performance only on the motion of the Research AGV during handovers. We did not experience any glitch during handovers in any control setup (full cloud-deployed control, as well as joystick control).

3.11 Validation scenario 8: Cloud reliability

3.11.1 Description of the scenario

The scenario *Cloud reliability* is defined to showcase the reliability feature of our cloud solution. As discussed in Deliverable D4.3 [5GS20-D430], we run two instances from the full control applications of both AGVs (only the fleet manager remained in single instance). The idea is to have a primary instance as the active controller, and a secondary instance is kept as a hot standby, which receives the necessary information to build its own state and can be switched into operation when the active controller fails. We would like to measure how the switching of the control instances performs. In this validation scenario, we evaluate the *Control takeover time* KPI. We run the Research AGV between two points and kill the primary controller instance, as illustrated in Figure 17.

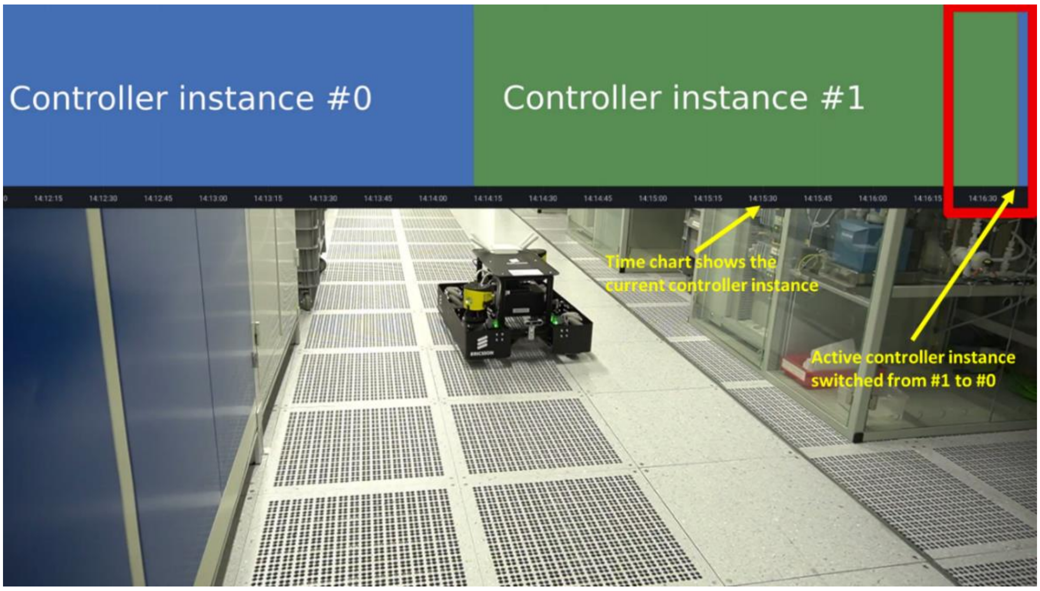

Figure 17: Snapshot of the evaluation process of the *Cloud reliability* scenario where we measured the values of the *Control takeover time KPI.*

3.11.2 Validation results

In Figure 18, we illustrate the histogram of the measured values of the *Control takeover time KPI* ranging from 240 to 290 ms. The effect of the control switching can also be observed in the demo video where we can see that the control takeover does not have any significant impact on the motion of the robot.

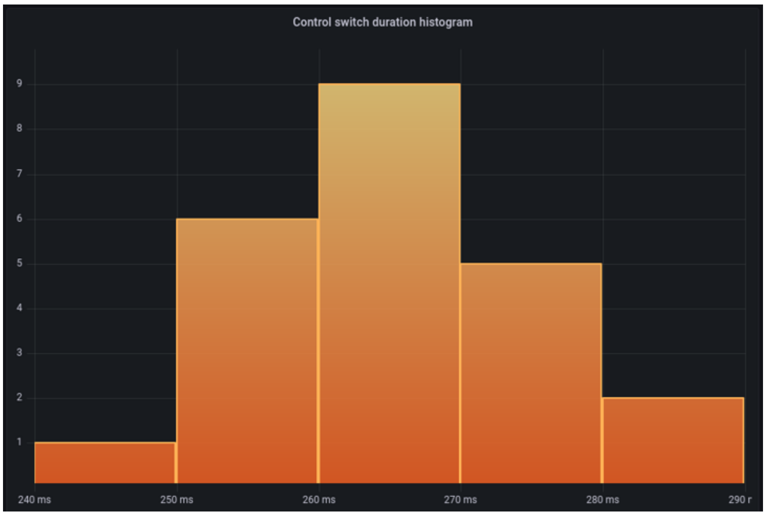

Figure 18: Histogram of the control takeover time KPI values.

3.11.3 Discussion

We note that these results include the delay contribution from all components. For instance, the servo motors of the Research AGV do not accept control commands for 100 ms from different IP address and port number than the active controlling instance. This means that the new controller is able to control the servo motors only after 100 ms. This timeout is also included in the measured values. So, these numbers represent the end-to-end view and show what performance the robot may get while doing failover.

3.12 Validation scenario 9: Cloud safety

3.12.1 Description of the scenario

We evaluate the performance of our cloud safety solution that is described in Deliverable D4.3 [5GS20-D43]. We implemented a safety solution in the cloud which stops the robots in case a collision is anticipated. We run the Research AGV (but the solution works for both AGVs) and evaluate the *Stopping time* and *Stopping distance* KPIs to quantify how much time and distance the AGV travels from the moment an object is put in the way until it stops completely as illustrated in Figure 19 (the corresponding video clip can be found here).

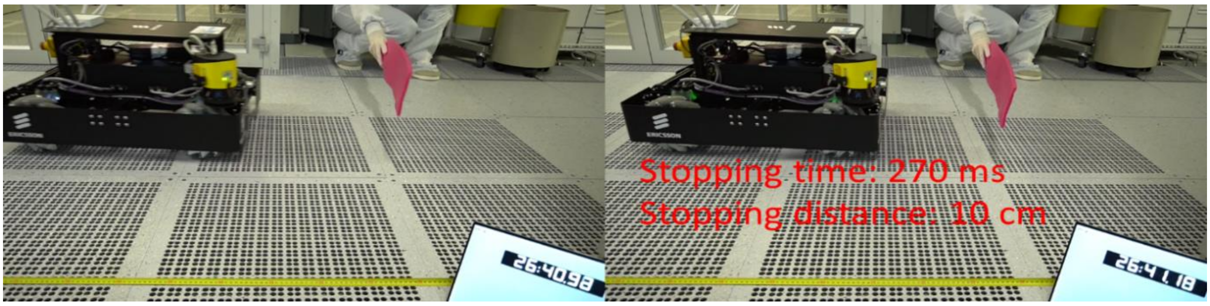

Figure 19: Snapshots from the validation of the Cloud safety scenario.

3.12.2 Validation results

In Table 7, we disclose the measured values of 10 executions.

Safety KPIs	
Stopping time (ms)	**Stopping distance (cm)**
300	14
270	11
310	12
360	16
290	13
350	14
270	10
380	18
290	12
370	18
319	**13.8**

Table 7: Results of the Stopping time and Stopping distance KPIs

3.12.3 Discussion

We note that these measurements depend largely on the speed of the robots, of course. During the validation tests, the speed of the AGVs was around 0.9-1.0 m/s. This speed is typically the maximum speed of many commercial AGVs. However, AGVs in factories are rarely used at full speed in such an environment, so these values may significantly decrease using lower maximum velocity of the robots.

3.13 Conclusion

Based on the above-detailed results, we can conclude that the validation tests of the Cloud-based mobile robotics use case were overall successful. Altogether 16 KPIs were evaluated, some of the refined and further subdivided during the measurements to better reflect the most important aspects of the evaluation. The performance of the 5G network is shown to be good for all validation scenarios of the use case as it is reflected also in the network-related KPIs. No connection loss, inadequate latency values or any other negative effects of the network on the performance of the AGVs have been identified during the tests. Overall the expectations from Bosch regarding AGV performacne via 5G were met.

The control application has been developed, including a separate fleet controller that instructed the robots when coordination was needed. The performance of the implementation was good enough according to Bosch, although the parameters can be further optimized to reduce end-to-end reaction time or map update latencies. The cloud-based execution of the logic over 5G showed significant gains in the travelled time (mission execution time), which involves significant energy savings and efficiency improvements, since the AGVs execute the missions faster, travel less per mission and thereby increase their battery lifetime. Furthermore, reduction in fleet size is another implication and can be considered as a cost saver. The reliability of the implementation was further enhanced with the execution of multiple control instances and low control takeover delays. A Kubernetes-based cloud environment served as the execution platform of the Docker containerized software components. We have not experienced any negative effect of the cloud environment on the performance of the robots. However, we realized that the cloud platform needs careful configuration to support moving data in and out of the cloud environment (e.g., towards or from the servo motors, sensors, etc.), especially when multiple instances are running from the control stack distributed to multiple physical servers.

We realized a rearchitected mobile robot in the form of the Research AGV where no computation is kept onboard, with decoupled HW and SW architectures. This made the whole HW construction much simpler and lighter, facilitating decreased maintenance cost. Overall, the implementation of the full control logic in a cloud environment over 5G using standard IT tools, protocols and programming languages enabled us to create a wireless friendly implementation that is more flexible and thereby can decrease the total cost of ownership. We showed that coordinated and real-time interaction among robot devices can considerably increase the efficiency of transportation while maintaining the desired KPIs.

4 TSN/Industrial LAN over 5G

4.1 Introduction

The integration of industrial communication with wireless 5G technology has been recognized as a key topic to establish 5G in the manufacturing verticals. Significant efforts have been made to characterize requirements from the industrial domains and enable Industrial Ethernet support in 5G, e.g. in the 5G-ACIA [5GACIA-IE]. The integration of 5G with TSN, an upcoming set of standards that will likely bring convergence in the broad field of Industrial Ethernet technologies, has been specified in the 3GPP Rel. 16. The integration of TSN and 5G has been investigated in detail in the 5G-SMART's Deliverable D5.1 [5GS20-D51] and Deliverable D5.3 [5GS21-D53]. Since the beginning of the project, the focus of the TSN/I-LAN over 5G use case has been to develop a proof-of-concept scenario to test and validate industrial use cases. Given the novelty and pre-commercial state of both technologies, i.e. 5G and TSN, a major challenge was to obtain the necessary hardware for the trials. As a result, industrial controllers and networking components supporting TSN features (e.g., TSN switches) are not widely available yet, which leads to the situation that TSN is currently not used in production as of today, making it difficult to find existing manufacturing use cases. Furthermore, 5G systems support TSN features only from Rel-16, which was not fully finalized and hence were not commercially available during the start of the 5G deployment in the semiconductor factory. Therefore, our 5G network deployment is based on Rel-15 products, which limits our possibilities for testing TSN features over the 5G system. Accordingly, we have designed the use case implementations with these limitations in mind.

The high-level architecture of the TSN/I-LAN use case is shown in Figure 20. This use case encompasses two types of industrial communication. The first type of industrial communication is between PLCs, while the second one is between the industrial machine Rudolph F30 and its backend server. Both industrial communication applications are briefly described in the subsequent sections, while more detailed description and the requirements of these applications can be found in 5G-SMART's Deliverable D1.1 [5GS20-D11].

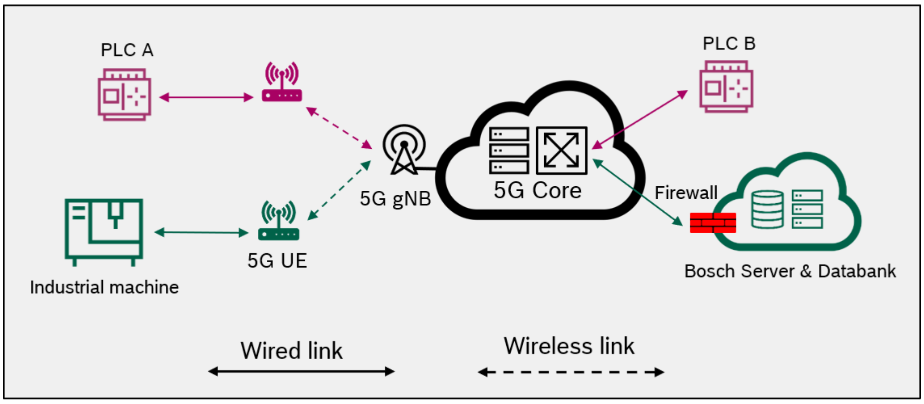

Figure 20: High level architecture of the TSN / Industrial LAN over 5G use case.

4.1.1 Controller to controller (C2C) communication

C2C communication refers to the communication between industrial controllers. It is required when, e.g., multiple controllers are executing coordinated tasks. A common application is a so-called line controller which monitors and controls different machines and actuators along a production line.

In the semi-conductor production in Reutlingen plant, each wafer undergoes multiple process steps on individual machines. The line controller tracks all production steps of each wafer and monitors relevant parameters. These parameters can include, for example, an identification number per wafer, temperature, and consumption of energy, pressurized air, coolant, and other chemical precursors in each machine. The monitoring of this data allows conclusions to be to drawn about quality and costs per wafer. Furthermore, it is possible to do predictive maintenance and react to abnormal behavior of the production machines. Collecting these parameters wirelessly by replacing the currently cabled connection between machines and the line controller does not directly pose any real-time requirements on the network. However, the delay and reliability of the network can influence the efficiency of production. Before a wafer enters the next machine, the line controller determines its ID, approves the processing, and collects some status data after completion. Some machines process wafers every few seconds and if the communication between line controller and machine takes significantly longer in a wireless setup compared to the wired setup, the utilization of the machines is reduced. Hence, the network is required to provide high availability and low latency to maximize production efficiency.

In our use case, we employ a simplified proof-of-concept setup of the above-described application. We use two types of industrial control units from Bosch Rexroth, the S20-ETH-BK [Eth] bus coupler and the IndraControl XM22 [Indra], as shown in Figure 21. The XM22 is a programmable embedded control platform that represents the line controller. The S20-ETH-BK is a simple coupling device that provides the sensor input from the industrial machine over the Ethernet interface to the XM22. In our use case, we connected four modules on the machine that provide data on energy consumption, as well as a single pressure sensor that measures the usage of compressed air. The logic of this test application resides in the XM22, which continuously requests the data from the S20-ETH-BK in predefined cycles. In a realistic setup, there is one S20-ETH-BK on each individual machine which tracks parameters depending on the process. The XM22 is programmed to sequentially request this data from each machine and store this data for post-processing. We are using this setup to better understand the structure of data packets and time behavior of the communicating controllers. Furthermore, we can draw conclusions on the networking delays, which is an important metric if such a line controller should be connected wirelessly in future deployments.

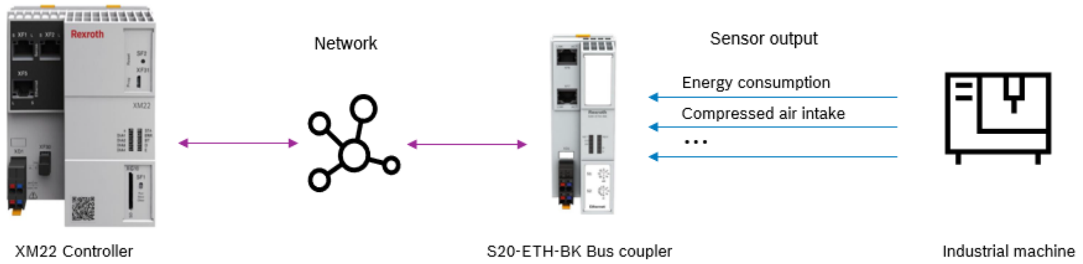

Figure 21: Test setup with Industrial control unit IndraControl XM22 and S20-ETH-BK bus coupler.

Operation mode of the C2C communication

The simplified operation mode of the considered use case is shown in Figure 22. In this example, we have one PLC A (XM22) and multiple PLC B, C, ... (S20-ETH-BK). PLC A establishes a TCP communication with another PLC and starts exchanging data. After one second, the TCP connection is terminated and a new connection to the next PLC on a different machine in a predefined order is initiated to exchange the latest data sets. For the test scenario, we only used two PLCs and analyzed the C2C communication between them. Nonetheless, the PLC connection is terminated and re-established to the same PLC each second to mimic the same behavior. The data exchanged between the PLCs can vary depending on what type of machine the S20-ETH-BK PLC is connected to. As mentioned before, we are tracking energy consumption from 4 different modules and the readings from a single pressure sensor. The S20-ETH-BK remains passive during operation and only sends out data packets upon request from the XM22. The request and response-type communication is performed over Modbus/TCP as shown in Figure 23. Modbus is a well-established fieldbus protocol in the industry based on a client/server architecture and it is defined in IEC 61158. Modbus is commonly used in control and monitoring devices to exchange I/O information and register data. Modbus/TCP uses TCP/IP packets to wrap and transmit the Modbus messages. It can be observed that the C2C communication does not expose a strictly consistent traffic pattern. For the energy measurements, the PLC A sends read requests to PLC B which are then answered by a read response (a). For the I/O sensor measurements, the PLC A is required to send a write register request, followed by a read request (b). As a result, the traffic pattern exhibits not a single fixed periodicity, but multiple different periodicities. This makes it more challenging to e.g. reserve resources in 5G which is sometimes achieved via semi-persistent scheduling.. The cycle time t_{cycle} of the application is set to 10 ms throughout the test cases in the following sections.

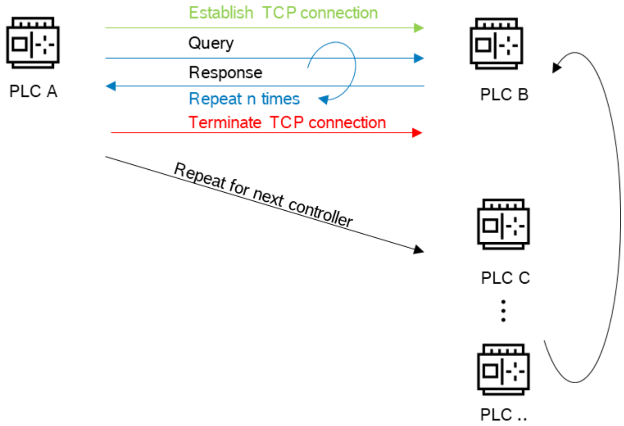

Figure 22: TCP/Modbus communication between PLCs.

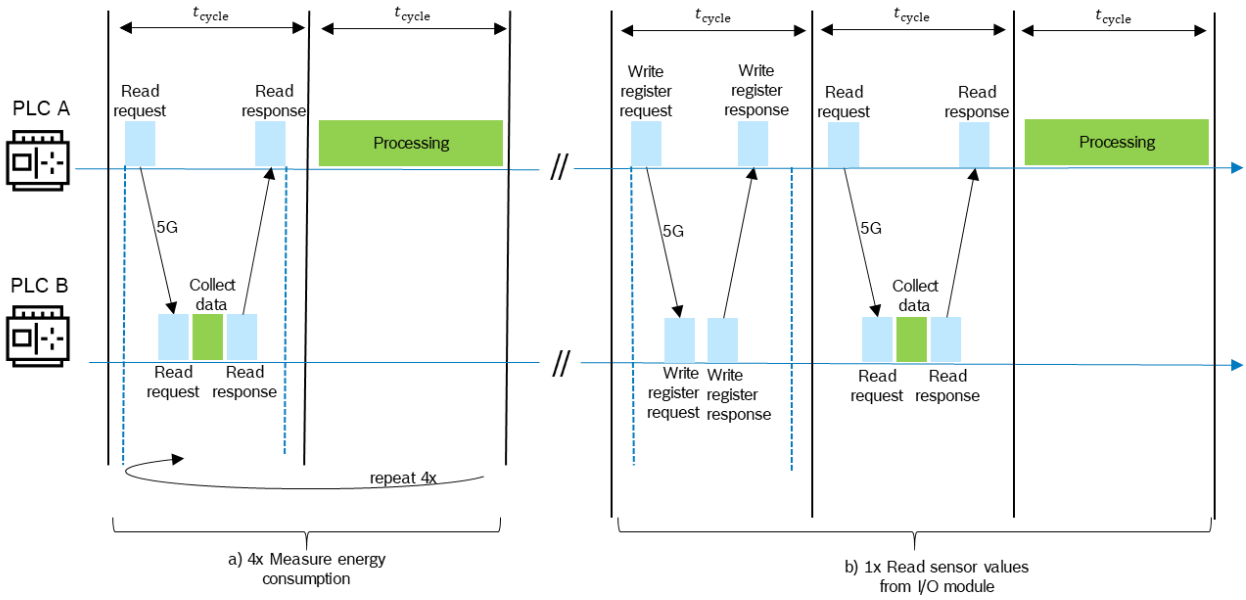

Figure 23: C2C communication between PLC A and PLC B using Modbus/TCP for energy measurements (a) and sensor readings (b).

A snapshot of the exchanged data between PLC A and PLC B for the energy measurements (a) and the sensor readings (b) is shown in Figure 24. Since the protocol is based on TCP, we can observe additional TCP acknowledgment packets on top of the aforementioned Modbus packets.

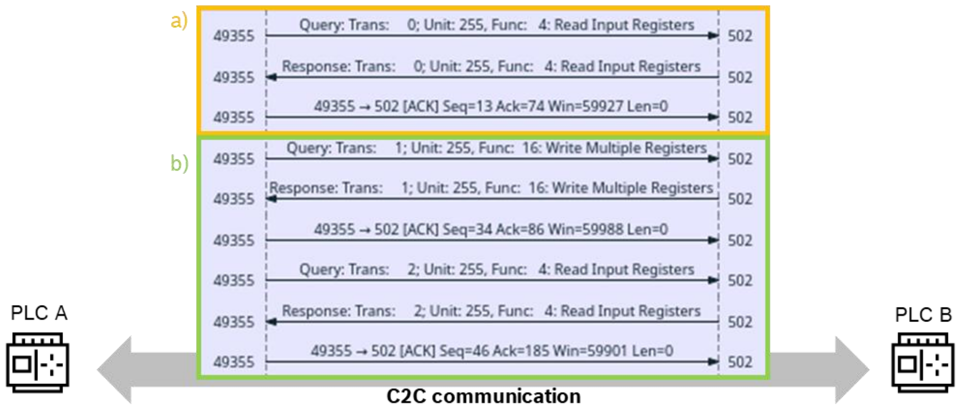

Figure 24: Message flow between PLC A and PLC B for energy measurements (a) and sensor readings (b).

4.1.2 Industrial machine to backend server communication

As a part of the TSN/I-LAN over 5G use case, we have a communication between the industrial machine Rudolph F30 (Figure 25 (a)) and its backend server, which is located in the Bosch network. This industrial machine is used to perform optical inspection of wafers in different steps of the production process. The result of the inspection are wafer coordinates of anomalies and abnormalities (defects), including a review and photos of this review (Figure 25 (b)). A single Rudolph F30 production equipment inspects around one hundred wafers per hour and the generated data is loaded into the backend server over a TCP/IP communication for automated classification of the detected defects. The detailed process description of the wafer inspection can be found in 5G-SMART's Deliverable D4.1 [5GS20-D41].

The original purpose of using the communication between Rudolph F30 and its backend server was twofold. First, we wanted to validate this machine-to-server communication over 5G. Second, we wanted to use this communication as background traffic for the C2C communications. Since the amount of data generated by Rudolph F30 does not vary significantly during its operation, it has been decided to use the iPerf3 tool as a source of background traffic. This is because it allows us to dynamically change the network load and test the performance of the C2C communication as well as the communication between Rudolph F30 under different network load scenarios. The details of conducted tests are described in the validation scenarios in Section 4.

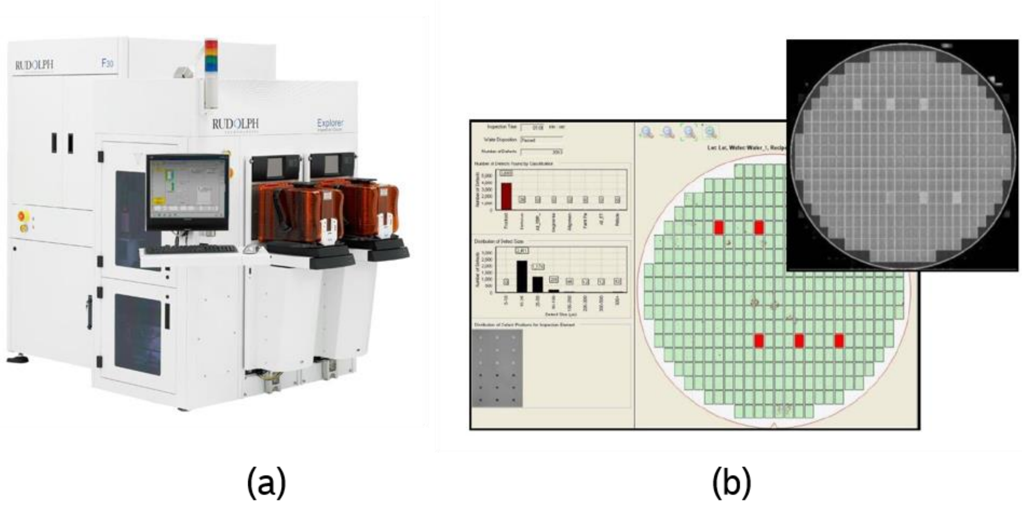

Figure 25: (a) Rudolph F30, (b) Output of the visual inspection. The pictures are properties of Bosch.

4.2 Industry goals

This section details the KPIs that have been used to evaluate the industry goals defined for the TSN/I-LAN over 5G use case, while the subsequent section provides the definition of the KPIs and the methodology used to calculate those KPIs.

Due to various reasons such as increasing production demand of different types of semiconductor products in the plant, there is often the need for setting up new industrial machines or relocating the existing ones. Given that currently the industrial communication in the plant is based on a wired Ethernet technology, both relocation as well as commissioning of industrial machines are considered to be time-consuming processes, which require a significant amount of preparatory engineering work in order to make sure that the industrial machine can be successfully operated in its target location. Therefore, it is important for the plant to reduce as much as possible the engineering effort and changeover time, thereby increasing the flexibility in commissioning or relocating industrial machines. This is achieved by replacing the wired Ethernet communication between these industrial machines with the 5G mobile communication. In the TSN/I-LAN over 5G use case, the industry goal *flexibility* is expressed and evaluated in terms of the commissioning time reduction and relocation time reduction for industrial machines (e.g., Rudolph F30). These are some of the KPIs of this use case (defined in Section 4.3), which can be achieved thanks to using the 5G mobile network technology instead of the wired Ethernet technology.

Currently, deployment, operation, and maintenance of the industrial machines in the plant incurs high capital and operational cost due to various reasons such as the need for proper cabling between the industrial machines as well as the need for installing network switches and network sockets to enable end-to-end communication between those industrial machines. Therefore, another important goal for the plant is to curtail both the capital and operational cost of the cables, connectors wear and tear, network switches and network sockets. This is reflected in the industry goal *sustainability*, which is expressed and evaluated in terms of annual cost savings KPI defined for this use case.

The *productivity* of the control application in the C2C communication is one of the important industry goals in this use case. A number of KPIs such as uplink (UL) and downlink (DL) transmission delay in the 5G network, the RTT, packet transfer interval as well as the C2C communication efficiency are used to evaluate the productivity of the control application. While employing a Release-15-based 5G network for the industrial applications considered in the TSN/I-LAN over 5G use case improves the industry goals such as flexibility, mobility, and sustainability, it does not deliver the same level of productivity for the control application in the C2C communication over the 5G network compared to wired Ethernet communication. This is mostly due to much higher packet transmission delay in the 5G network compared to the wired Ethernet network, as detailed in the subsequent sections.

Changing the communication technology and medium for the C2C communication and the communication between the industrial machine Rudolph F30 and its backend server may have a significant impact onto the *quality of the communication*. The communication quality has a direct impact onto the performance of both applications considered in the TSN/I-LAN over 5G use case and is expressed in terms the KPIs such as the number of packet drops.

4.3 Validation KPIs

In this section we introduce the KPIs that have been defined for this use case for both C2C communication as well as the industrial machine Rudolph F30 and its backend server communication. We also present the results of the KPIs that are not specific to any validation scenario described for this use case.

4.3.1 Validation KPIs for the C2C communication

The validation KPIs are summarized in Table 8. We use the KPIs to compare performance of the C2C use case in the wired and different wireless scenarios.

KPI name	KPI category	KPI short definition
Packet loss	Quality	Rate of packet loss between two PLCs.
Transmission delay	Productivity	The latency a packet experiences when transmitted over the wireless network. In the wireless network we distinguish between uplink latency and downlink latency.
Round trip time	Productivity	The round-trip time is the combined latency between sending a query and receiving a response. The RTT can be based only on the network delay, or additionally include the processing time between query and response.
Transfer interval time	Productivity	The transfer interval refers to the time between two transmissions. The C2C use case requires a response message before the next packet is sent out. Hence, the round-trip time provides a lower bound on the

		minimal feasible transfer interval time and thus also the cycle time of the application.
C2C communication efficiency	Productivity	The network delay directly impacts the minimum transfer interval time and hence the maximum rate at which packets can be exchanged over the network. The C2C communication efficiency is defined by the ratio of total number of read response transmissions in the 5G setup, compared to the total number of read response transmissions in the wired best-case scenario within the same time period.

Table 8: Validation KPIs for the C2C communication.

To measure the KPIs, we rely on Ethernet taps from Profitap as shown in Figure 26. It supports precise hardware timestamping and a pulse per second (PPS) signal for timestamp synchronization between two taps. The taps are placed close to the endpoints, i.e. the PLCs, to measure the end-to-end performance of the network in between. The network over which the C2C traffic is routed can be a wired Ethernet network, or a wireless 5G network. The taps capture each packet between the PLCs and add a timestamp when it was captured. Both tap A and tap B are connected over a PPS signal which allows a precise synchronization between timestamps captured on either tap. From the packet traces, all KPIs displayed in Table 8 can be derived or calculated as detailed in the following.

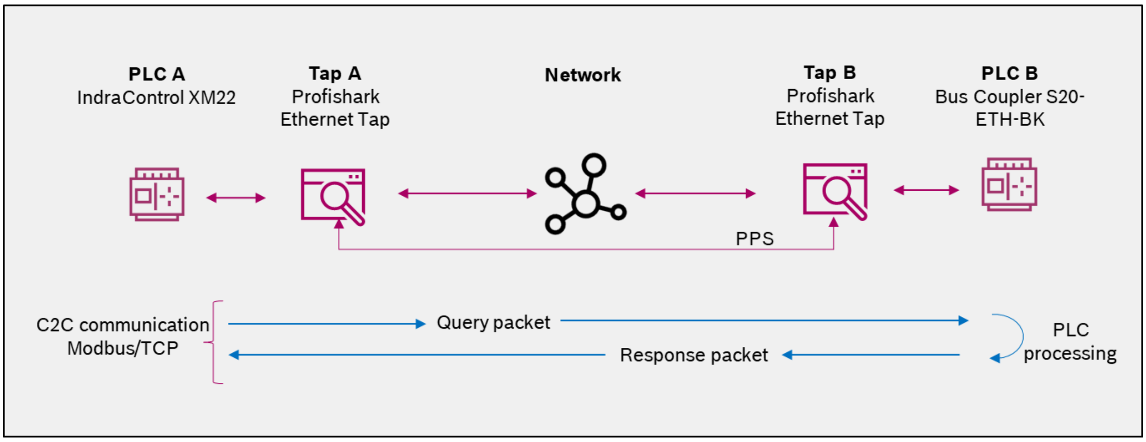

Figure 26: Basic testing setup with PLC endpoints and ProfiShark 1G+ Ethernet TAPs.

Packet Loss

Each packet sent out by one controller must arrive at the other controller and vice versa. The packet loss is hence determined by comparing the packet IDs pid captured on tap A and tap B. If any packet ID is only captured on one tap, a packet must have been lost inside the network. We only trace the end-to-end packet losses, 5G-internal processes such as potential retransmissions remain unrecognized by the taps. Throughout the validation scenarios, we have not observed any packet losses on application level, unless we intentionally send more data over the network than its bandwidth allows for. Hence, we can state that we observed a 100% availability of the network during the tests.

Transmission delay

Each packet with ID pid is captured on the first tap A with timestamp T_{pid}^{A} and on the second tap B with timestamp T_{pid}^{B}. The latency can be determined by taking the difference of both timestamps $T_{pid}^{B} - T_{pid}^{A}$. In case of a 5G network, we distinguish between uplink latency $t_{\mathrm{5G\_UL}}$ and downlink latency $t_{\mathrm{5G\_DL}}$. The latency of each packet is measured end-to-end from controller to controller, hence including 5G core, transport and UE latency. The taps are connected over a PPS signal which allows for a synchronization accuracy of up to 10ns according to the manufacturer [Profitap].

Round-trip time

The round-trip time calculation is based on the fact that queries and responses carry the same ID in Modbus/TCP. The round-trip time including the PLC processing time can hence be determined by taking the timestamp of the outgoing query from PLC A on tap A with ID pid, i.e. $T_{pid}^{A,\mathrm{query}}$, and the timestamp of the incoming response with the same ID pid, i.e. $T_{pid}^{A,\mathrm{response}}$(cf. Figure 24).

$$t_{\mathrm{RTT}} = t_{\mathrm{5G\_UL}} + t_{\mathrm{PLC}} + t_{\mathrm{5G\_DL}} = T_{pid}^{A,\mathrm{response}} - T_{pid}^{A,\mathrm{query}}$$

The 5G RTT without processing time is calculated by adding the one-way latencies with matching IDs in either direction. In case the query is sent by PLC A, the 5G RTT is calculated by

$$t_{\mathrm{5G\_RTT}} = t_{\mathrm{5G\_UL}} + t_{\mathrm{5G\_DL}} = T_{pid}^{B,\mathrm{query}} - T_{pid}^{A,\mathrm{query}} + T_{pid}^{A,\mathrm{response}} - T_{pid}^{B,\mathrm{response}}$$

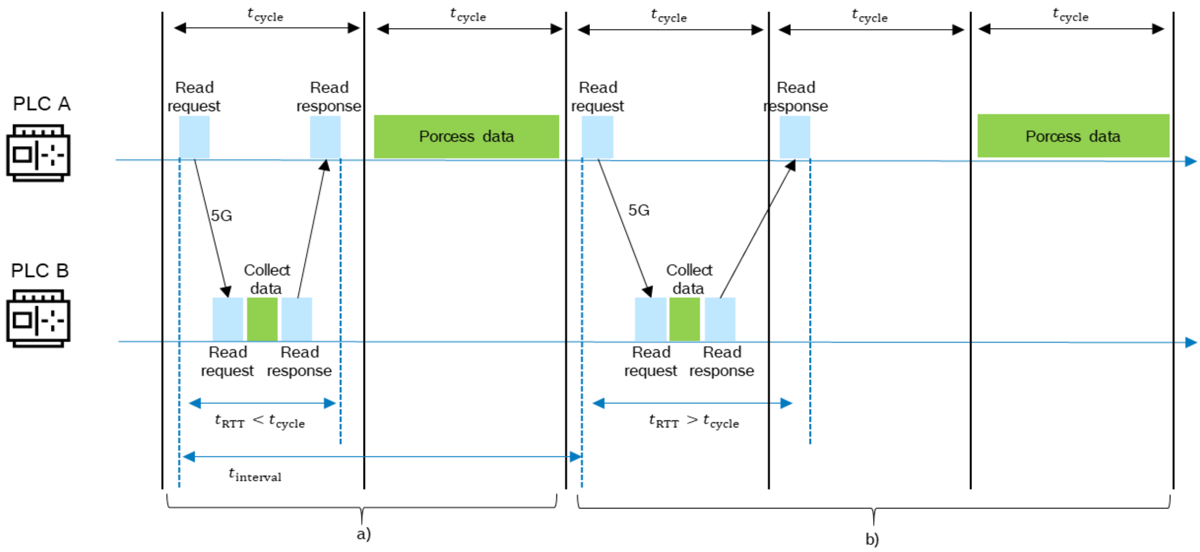

Figure 27: Dependency of round-trip time, transfer interval time, and cycle time.

Transfer interval time

The transfer interval time is the time between any two transmissions in one direction. This value depends on the application cycle time t_{cycle}, and the actual sequence of tasks programmed in the PLC. The dependency between transfer interval time and cycle time can be seen in Figure 27. For the read commands, the PLC can either collect the data from the remote PLC or process the received data per cycle. In the ideal case (a) where $t_{\text{RTT}} < t_{\text{cycle}}$, this leads to $t_{\text{interval}} = 2t_{\text{cycle}}$. However, when the data acquisition cannot be completed within one t_{cycle}, i.e. $t_{\text{RTT}} > t_{\text{cycle}}$, the processing of the data can only start a cycle later (b), hence $t_{\text{interval}} = 3t_{\text{cycle}}$. In our particular use case, network delays are hence tolerable and lead to longer idle times and hence reduced efficiency of the PLC. The PLC would only abort after a time-out of multiple cycles without successful communication. In other time-critical applications, it is however also common that a single missed cycle could lead already to abnormal behavior of the application. We can verify the transfer interval time by determining the time delta between any two timestamps.

C2C communication efficiency

The efficiency is a calculated value that indicates how efficiently data can be exchanged between two controllers. Since a new data packet can only be requested from the PLC after the last packet has been received successfully, the RTT effectively limits the packet rate, as shown in Figure 27. Hence with increasing latency on the 5G network, we should observe a decreasing number of packets transmitted within the same time frame. The C2C communication efficiency is defined by the ratio of total number of read response transmissions over 5G compared to the total number of read response transmissions in the wired best-case scenario within the same time period, i.e.

$$e = \frac{\#\ \text{Read response packets over 5G}}{\#\ \text{Read response packets over Ethernet}}$$

4.3.2 Validation KPIs for industrial machine to its backend server communication

The validation KPIs for the industrial machine Rudolph F30 to its backend server communication are summarized in Table 9.

KPI name	KPI category	KPI short definition
Commissioning time reduction [h]	Flexibility	The time difference between the commissioning time of an industrial machine (e.g., Rudolph F30) in the wired network and the 5G network.
Relocation time reduction [h]	Flexibility	The time difference between the relocation time of an industrial machine (e.g., Rudolph F30) in the wired network and the 5G network.
Annual cost savings [€]	Sustainability	Cost saving for all industrial machines (e.g., Rudolph F30) used in the plant when operated over a 5G network.

Table 9: Validation KPIs for the industrial machine Rudolph F30 to its backend server communication

The estimation of the annual cost savings, relocation time reduction, as well as the commissioning time reduction KPIs, apply to all industrial machines (including Rodolph F30s) in the entire plant that could be expected to be endowed with 5G communication capabilities. The values used in the calculation of those KPIs are estimated by the plant based on their experience.

Commissioning time reduction

For large production plants like the semiconductor factory in Reutlingen, commissioning of a production machine is a frequently happening activity due to the need for replacing a faulty production machine or commissioning new ones due to increase or change in the production demand. Commissioning of an industrial machine in the factory requires accomplishment of several tasks pertaining to the physical installation. Even if a basic IT-infrastructure already exists in a production site, every additional industrial machine has to be integrated individually based on specific requirements of the industrial machine itself, as well as the requirements of the individual location of the machine. Therefore, the following process steps have to be taken into consideration for commissioning an industrial machine:

- Installation of one or more cables interfacing the industrial machine
- Patching of sockets and connecting cables to those sockets
- Connection of socket cables to switches inside server cabinets

The necessary IT Infrastructure (e.g., network sockets, switches, server cabinet) and the production machines are separated into two production layers. While the industrial machines as well as network sockets are located in the main production shopfloor, switches and server cabinets are located in the so-called Sub-Fab, which is a support shopfloor layer under the main production shopfloor. The distance from an industrial machine to network sockets as well as from a network socket to switch/server cabinet varies for almost every industrial machine. For example, if an industrial machine is located directly over a server cabinet, then their distance is very short, within 10m to 25m range. Since the necessary IT-infrastructure (e.g., network sockets, switches, server cabinet) is not always available at the specific location, it might be required to install new infrastructure, resulting in an additional delay in commissioning time. Considering the mentioned aspects, the total commissioning time of an industrial machine by cable, $T_{cable}^{commissioning}$, is estimated as an average value:

$$T_{cable}^{commissioning} = 8h$$

With the 5G network, most of the tasks (e.g., installation of cables and network sockets) required to commission an industrial machine with a cable are no longer required. Therefore, assuming that the 5G network has already been deployed and there is no coverage or capacity limitations at the location of industrial machines, the total commissioning time of an industrial machine in the 5G network, $T_{5G}^{commissioning}$, is mostly the time required to provision a SIM card, connect the 5G modem to the industrial machine and configure it. $T_{5G}^{commissioning}$ is estimated to be on average:

$$T_{5G}^{commissioning} = 1h$$

Given the total commissioning time of an industrial machine in a wired network $T_{cable}^{commissioning}$ and in a 5G network $T_{5G}^{commissioning}$, the commissioning time reduction, $T_{5G-cable}^{commissioning}$, of an industrial machine is computed as follows:

$$T_{cable-5G}^{commissioning} = T_{cable}^{commissioning} - T_{5G}^{commissioning} = 8h - 1h = 7h$$

Thus, an industrial machine can be commissioned in a 5G network on average 8 times faster than in a wired network.

Relocation time reduction

Depending on the production needs, some industrial machines are often relocated within the factory shopfloor. This is normally a time-consuming process, which like commissioning of an industrial machine requires several tasks to be accomplished. More specifically, if the industrial machine is connected by cable to the IT-infrastructure at the old position, besides all specific requirements, the following process steps have to performed to relocate an industrial machine to a new position:

- Removal of existing wired connection (cable – network socket – cable – switch)
- Installation of a new infrastructure (network sockets, server cabinets, switches, cables in IT-room)
- Commissioning of the industrial machine at new location as described above

Thus, in addition to the aspects considered in the estimation of “Commissioning time reduction” KPI, there are several aspects that need to be taken into account. Deinstallation of the IT-Infrastructure (e.g., cables, switches, network sockets, server cabinets) at the old position of the industrial machine depends on the specific reuse-strategy at that position. Therefore, the process time $T_{cable}^{decommissioning}$ to remove existing wire connection can be estimated as half the commissioning time on average:

$$T_{cable}^{decommissioning} = 0.5\, T_{cable}^{commissioning} = 4h$$

With respect to these aspects the relocation time $T_{cable}^{relocation}$ of an industrial machine by cable can be estimated as an average value as follows:

$$T_{cable}^{relocation} = T_{cable}^{decommissioning} + T_{cable}^{commissioning} = 4h + 8h = 12h$$

With the introduction of the 5G network it is expected to have a significant reduction in the relocation efforts thanks to the streamlined mobility of industrial machines enabled by the 5G network. Since it is assumed that the 5G network is already deployed at the plant, the relocation time of an industrial machine in the 5G network $T_{5G}^{relocation}$ would mainly be the time to reconfigure the 5G modem connected to the industrial machine, which in some of the cases might not even be required. Taking the mentioned aspects into account, the relocation time of an industrial machine in the 5G network is estimated to be:

$$T_{5G}^{relocation} = 0.5h$$

Given the relocation time in the wired network as well as in the 5G wireless network, the relocation time difference $T_{cable-5G}^{relocation}$ can be estimated as follows:

$$T_{cable-5G}^{relocation} = T_{cable}^{relocation} - T_{5G}^{relocation} = 12h - 0.5h = 11.5h$$

Thus, an industrial machine can be relocated in a 5G network on average 24 times faster than in a wired network.

Annual cost savings

Annual cost savings are mainly interesting for the second part of the use case: The integration of production equipment to central production controlling systems like databases, central decision systems, and others. In 5G-SMART, this aspect is represented by the integration of production equipment Rudolph F30 to its backend. In the annual cost savings calculations, it is in general assumed that a basic 5G-infrastructure is already available, and 50% of the existing 2000 industrial machines are already attached to 5G-modems and have 5G communication capabilities. 50% of 5G communication rate of the existing machines is justified by the fact that for some industrial machines the 5G communication would not be beneficial, and not all production sites in the plant would have 5G network installed due to, for example, the physical infeasibility of deploying antennas, IT security requirements, or potential Electromagnetic Compatibility (EMC) issues with some of the semiconductor products, which has been tested in the factory as one of the key activities in this testbed [5GS22-D42].

The annual cost savings are calculated for the use case of production equipment which has to be connected to central production site IT-systems. This connection can be realized via a wired network, which is normally the case in the plant, or via a 5G mobile network. There are two main sources for the annual cost savings. Commissioning of new 5G-enabled industrial machines per year is the first source of cost savings due to reduced setup time, effort, and cost compared to the wired setup, which required installation of cables, switches, racks, network sockets and more working hours. Relocation of the existing 5G-enabled industrial machines within the plant per year, which is a frequently occurring process due to production needs, is the second source of cost savings obtained from reduced setup time, effort, and cost like when commissioning new industrial machines. As opposed to commissioning a new industrial machine in a wired network, relocation of an existing industrial machines requires some effort first to remove the exiting wired connection and then install it in its new location. However, some of its components, such as cables could be re-used in the installation of this industrial machine in its new location. Therefore, it is assumed that commissioning of a new industrial machine and relocation of an existing industrial machine would require the same approximate setup cost in the wired network (e.g., cabling, switches, racks, network sockets and effort (working hours)) $C_{wired\_setup} = €2500$. It is also assumed that on average $N_{new\_machines} = 50$ 5G-enabled new industrial machines are commissioned every year in the plant, while on average $N_{existing\_machines} = 50$ existing 5G-enabled industrial machines are relocated in the plant. Thus, taking into account a 5G wireless modem cost $C_{5G\_modem} = €250$, the annual cost savings S_{annual} can be computed as follows:

$$S_{annual} = N_{new\_machines} * \left(C_{wired\_setup} - C_{5G\_modem}\right) + N_{existing\_modem} * C_{wired\_setup}$$
$$= 50 * (2500 - 250) + 50 * 2500 = €237.500$$

It is important to note that the annual cost savings does not take into account the capital expenditure (CAPEX), such as the deployment cost of the 5G network components (e.g., RDs), and operational expenditure (OPEX), such as power consumption, firmware uprades, security patches, of the 5G network.

4.4 Baseline scenario: C2C communication over wired Ethernet connection

The traffic characteristics of C2C communication for wired Ethernet connectivity between the two PLCs is discussed in the following. In the basic setup in Figure 26, we replaced the network with a single Ethernet cable. Since no significant delay can be expected over a single Ethernet link, the testing setup is reduced to a single tracing device to inspect the traffic (cf. Figure 28).

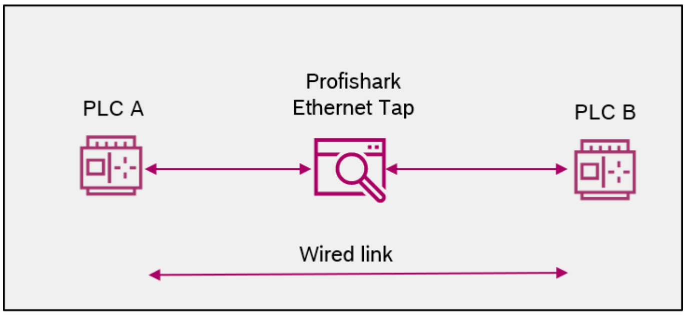

Figure 28: C2C communication setup.

The results are shown in the following four figures. The data rate in kbps of queries from PLC A to PLC B, as well as the rate of responses in the opposite direction, is shown in Figure 29. The unsteady behavior is caused by the re-connections performed each second by the PLCs, which results in a roughly 200ms interruption. Additionally, the average rate over a time window of 5s for the query and response rate is shown, respectively. The data rate of responses, which carries the relevant process data is with 28 kbps on average roughly twice the rate of queries. An inspection of the traces reveals that from this data rate, only around 27.4% or 67.6% accounts for the actual Modbus payload of the queries and responses, respectively. The remainder is caused by packet headers. A high frame rate and high overhead in comparison to the payload is typical for C2C communication. This usually poses several challenges to wireless networks that are usually tailored for bandwidth-intensive applications.

The round-trip time (RTT) is depicted in Figure 30, over time and as the complementary cumulative distribution function (CCDF). Since there is only a single tap used in the setup, we cannot derive any latency values for the network, which however would be negligible in this scenario. The decisive factor for the RTT is the time between the reception of a query and transmission of the corresponding response on the PLC, which we define as the PLC processing time in the following. The RTT in this baseline scenario is approximately equal to the PLC processing time. Depending on the processing task on the PLC, i.e. read or write register commands, the RTT value varies. In case of write queries, the PLC B sends out the response immediately, leading to a RTT of approximately 0.8ms. For read commands, the values range from 0.8ms up to 9ms. This delay between reception and response is not only due to active processing on the PLC, but also accounts for alignment delays. Sometimes the sensor readings are only sent back in the subsequent cycle, in which case the PLC remains in an idle state until the next cycle starts. The implication for the network is that not all transmissions happen in a deterministic way. While the PLC A requests data in a mostly periodic fashion, the responses are sent out depending on the processing and alignment delays of the PLC.

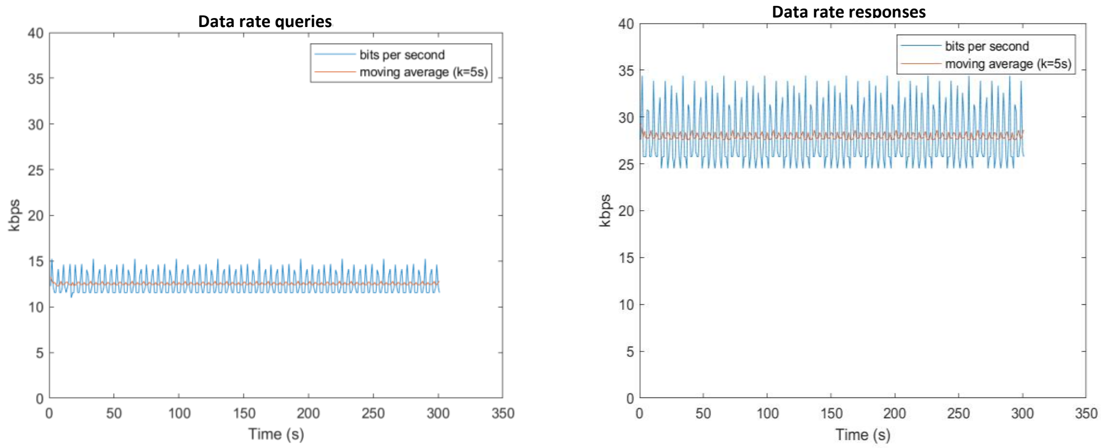

Figure 29: Data rate of queries (left) and responses (right).

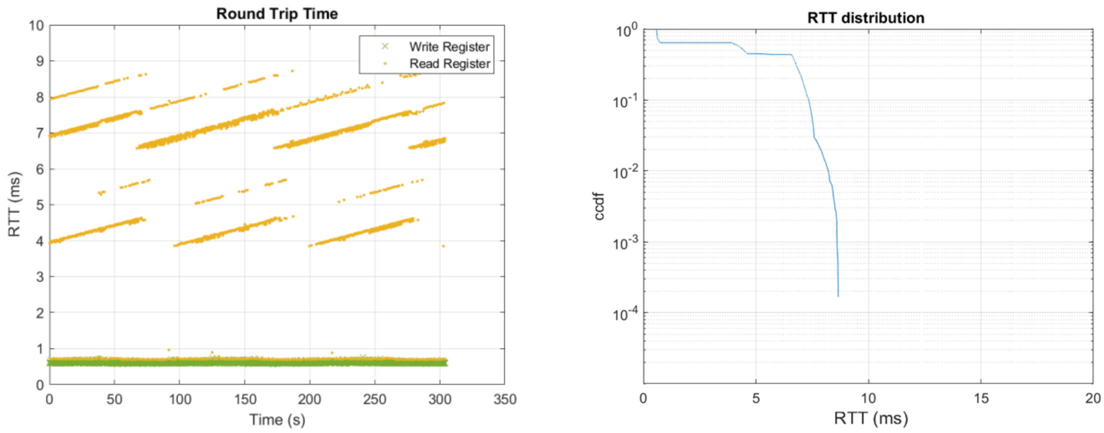

Figure 30: Round-trip time including PLC processing time.

4.5 Validation scenario 1: C2C communication over the 5G network without background traffic

4.5.1 Description of the scenario

In this validation scenario, the PLCs are communicating with each other over the 5G network, and there is no background traffic (i.e., any other traffic) in the network, as depicted in Figure 31. A pair of Ethernet taps (i.e., Profitap 1G+) are connected to the PLCs, one between the PLC A and the 5G UE, and the other between the core network and the PLC B, to capture the traffic between the PLCs both in the uplink (i.e., from the PLC A to the PLC B) and in the downlink (i.e., from the PLC B to the PLC A) in order to accurately measure the KPIs such as UL/DL delay in the 5G network. The main objective of this validation scenario is to determine how much latency is introduced by the 5G system comprising the network and the UE compared to the wired setup. Furthermore, we can determine the impact on the controller traffic in terms of the C2C communication efficiency which describes how many packets can be transmitted over 5G compared to the Ethernet solution.

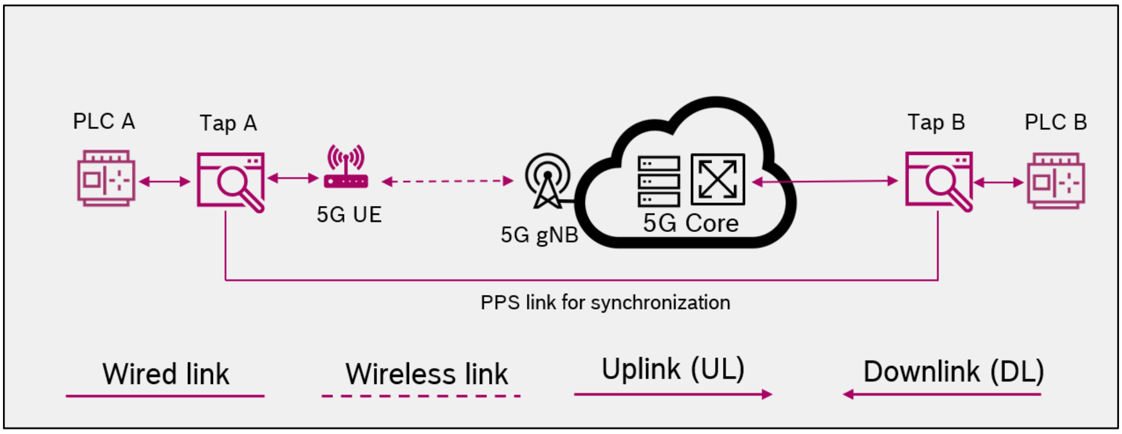

Figure 31: Setup of the validation scenario 1.

4.5.2 Validation results

In the general test setup in Figure 26 the network is now replaced by the 5G network. Hence, we can use the two taps to determine latency values for the end-to-end system. The data rate in UL and DL for the C2C communication are displayed in Figure 32. The average values match the results from the wired results in Figure 29. However, we observe higher fluctuations in the rate when using 5G compared to Ethernet.

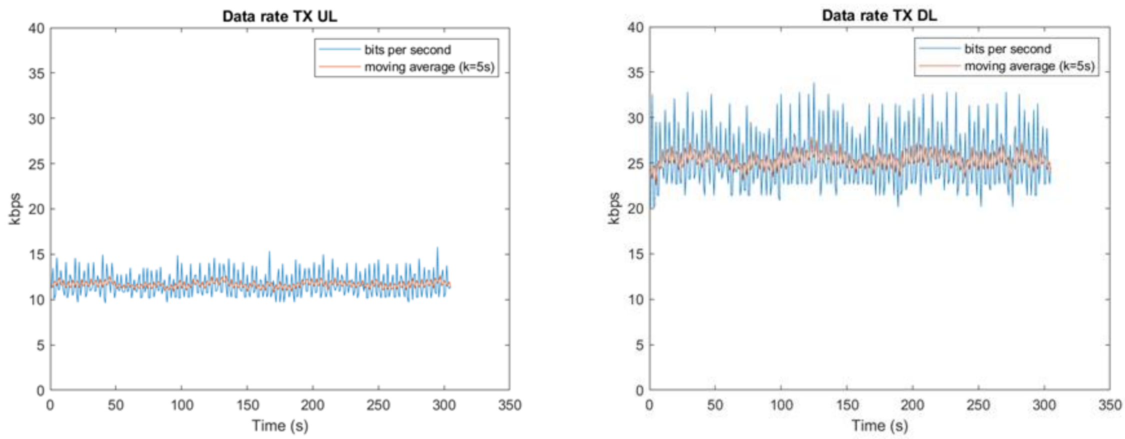

Figure 32: Data rate in uplink and downlink

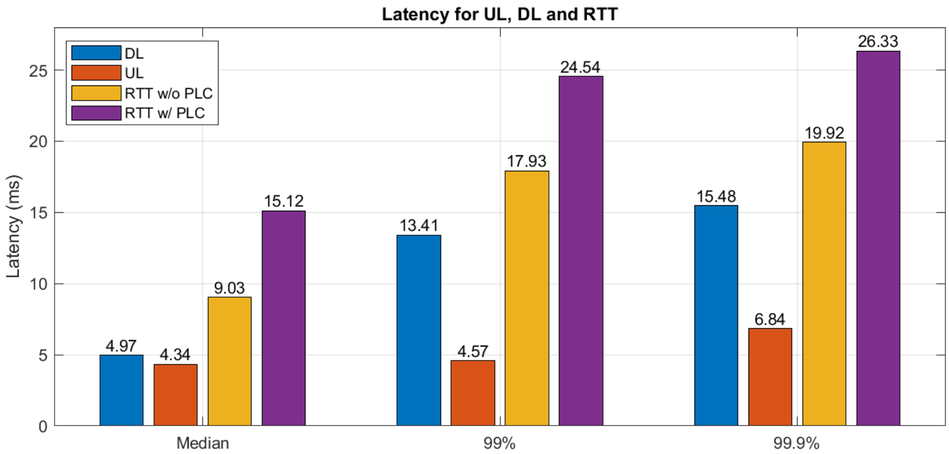

Figure 33: Downlink latency, uplink latency and round trip time with and without PLC processing time.

For the latency evaluation, we show in the remainder of this work bar graphs that indicate the median, the 99th and 99.9th percentile of the CCDF. The latency in DL and UL, as well as the RTT with and without the PLC processing time is displayed in Figure 33. We can observe that in UL, 99.9% of the packets are transmitted within 6.8ms, in DL within 15.5ms. These values represent the end-to-end 5G latency values measured close to the endpoints.

From the measurements, it appears that the uplink latency has a lower bound compared to the downlink. However, it is important to notice that we do not have a fair comparison, given the rate on UL and DL is not identical as shown in Figure 32. Nevertheless, it is worth noting that the application data rate is multiple orders of magnitude below the supported bandwidth of the 5G system. In the used setup, 5G was able to provide a data rate close to 100 Mbps and 1000 Mbps for the UL and DL, respectively. Hence, we do not expect significant queueing and scheduling delays inside the 5G system for either UL or DL.

The round-trip time is shown in the same Figure 33. Unlike in the wired test case (cf. Figure 30), the network contributes now significantly to the overall RTT value. Comparing the results for 99.9% of the transmissions, the RTT is below 27ms. Subtracting the PLC processing time, this value reduces to around 20ms.

The results on the transmission latency and RTT directly impact the cycle time and hence the C2C communication efficiency as defined in Section 4.3. The reason is that the Modbus/TCP protocol running on the PLC only initiates a new transmission after one query-response cycle has been concluded. While this approach works very well in cabled bus systems with very low latency, it is not well suited for wireless systems that experience longer latency. The RTT hence provides a lower bound on the minimum transmission rate as detailed in Figure 27: Dependency of round-trip time, transfer interval time, and cycle time. Since we are using a cycle time of t_{cycle} = 10ms, the PLC can send at most every 20ms a new read request. From the RTT results in Figure 33, however, we can derive an RTT 27ms for the 99.9% bar. For each packet exceeding the 20ms limit, an additional t_{cycle} is consumed by the communication, reducing the overall efficiency . In this scenario, we measured that around 10% of the packets will require two cycles to complete the transmission.

4.5.3 Discussion

As expected, changing the communication means between the controllers from the wired Ethernet network to the 5G wireless network comes with the cost of increased RTT, which in turn has a negative impact onto the C2C communication efficiency as discussed in the previous section. This cost is expected to become even higher when introducing other kind of background traffic in the network. Depending on the network configuration, such as traffic prioritization, one could expect to see impact of the background traffic onto the controller traffic in terms of increased RTT and, therefore, decreased C2C communication efficiency. This is what is tested in the validation scenario 2.

It is worthwhile to mention that in a wired setup, the PLCs are configured to use a 1ms cycle time by default. However, the C2C communication over the 5G network would fail with this configuration, since the TCP handshake between the controllers leads to timeouts due to the 5G network transmission delay. During our tests we have seen that 10 ms is the lowest feasible cycle time applicable to the 5G setup. It is an important finding that such PLCs with their default configuration cannot be operated directly over the 5G network but require modifications on the application level. This is especially important for applications which might have stricter requirements on low cycle times such as closed-loop controls. The same cycle time has been used in all validation scenarios including the baseline scenario (i.e., C2C communication over the wired Ethernet network) for the sake of fair comparison of the C2C communication efficiency KPI.

From this validation scenario it becomes apparent that the industrial protocol itself is not well suited to be operated directly over a wireless network for this type of application due to its dependency on the RTT. While there are certain applications that can only send out the next data packet when the previous feedback has been received, such as close-loop control applications, the monitoring task in this use case does not require this. In the wired domain this Modbus/TCP protocol is well established and operates effectively over an Ethernet link. On the wireless side however, there are potential efficiency gains for the application that can increase data rate and reduce packet overhead by tailoring the protocol to operate over a network with long RTT. This can includee.g. the aggregation of data to larger packets (if there are no real-time constraints) instead of requesting each sensor value with an individual request packet. Furthermore, the usage of TCP over wireless networks with tpically longer RTT is not an ideal, alternatives such as QUIC or UDP should be considered to improve performance.

4.6 Validation scenario 2: C2C communication with background traffic and without controller traffic prioritization

4.6.1 Description of the scenario

The high-level architecture of this validation scenario is displayed in Figure 34. There are two types of streams in this validation scenario. The first stream is between the PLCs, while the other stream is between the iPerf client and its server. It is important to mention that the traffic of the PLCs is not prioritized over the iPerf traffic. Thus, they have the same default priority in the 5G network. The iPerf stream is used to generate a different amount of background traffic in both the uplink and downlink directions in order to analyze its impact onto the communication between the PLCs, which is the main objective of this validation scenario. This is done by gradually increasing the background iPerf traffic in the uplink/downlink and calculating the efficiency of the C2C communication, as defined in Section 4.3.1. In the uplink, we increase UDP traffic in 10 Mbps steps using the iPerf (version 3.7) tool up to 90 Mbps, which is the maximum throughput we were able to achieve and hence the upper bound of the 5G network under the given test conditions. In the downlink, we increase UDP traffic in 100 Mbps steps up to 900 Mbps, which is close to the limit of the 1 Gbps Ethernet interface used as the interface to the 5G network. In total, 9 experiments in the uplink and 9 experiments in the downlink have been conducted each with 5 minutes duration for the PLC and the iPerf streams. In the experiments, both the 5G UE connected to the PLC A and the 5G UE connected to the iPerf have a line-of-sight visibility with the Radio Dot, being away from it around 3 meters and 30 meters, respectively.

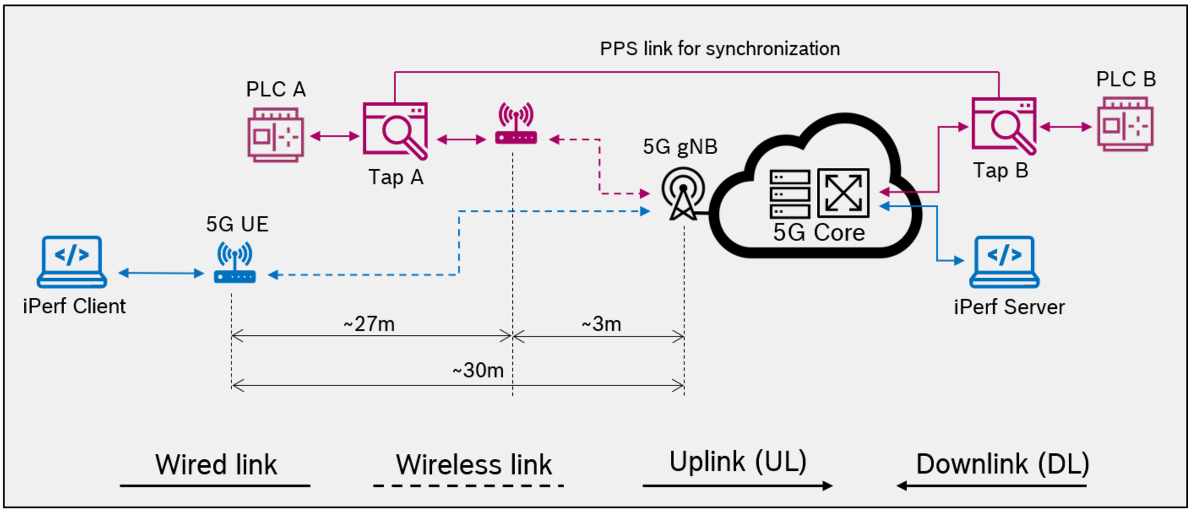

Figure 34: Setup of the validation scenario 2.

4.6.2 Validation results

The results section consists of two parts. In the first part, we investigate the network performance of the C2C traffic while loading background traffic in uplink direction. In the second part, we load the downlink direction with background traffic to analyze its impact onto C2C traffic. The results are compared to the validation scenario 1 without background traffic to analyze the impact of high-bandwidth cross-traffic in either transmission direction on the critical C2C packets. In general, for both cases, we observed no packet losses on application level throughout the entire testing session.

Uplink background traffic

In the first part, background traffic is sent only in uplink direction, i.e. from the iPerf client attached to the 5G UE to the iPerf server on network side as shown in Figure 34. The load of the background traffic ranges from 10 Mbps to 90 Mbps in 10 Mbps steps. For better visibility, the results for 0, 10, 50 and 90 Mbps are shown in the figures throughout the section. In Figure 35 and Figure 36, the latency of the C2C traffic in DL and UL, respectively, is depicted. Unsurprisingly, the DL latency is unaffected by the background traffic, since this scenario only considers UL background traffic. In UL, the C2C traffic is competing with the background traffic for the same network resources. As a result, higher UL latency is observed when introducing background traffic. For 0 Mbps, we achieve a delay of 6.8 ms for 99.9%

of the packets. With background traffic the latency increases to around 12 ms. The combined RTT is shown in Figure 37. Without background traffic we observe an RTT of under 27 ms for 99.9% of all C2C packets, with background traffic 31.5–33 ms. Subtracting the PLC processing time, we have 20ms and 25.5-26.6ms for the 99.9% target with and without background traffic, respectively. In general, we can observe that the presence of background traffic increases the latency, however, the actual bandwidth of the background traffic does not have a significant impact on the latency increase. The slight variation in latency between the 10, 50 and 90 Mbps test cases is not purely caused by the background traffic, but is also influenced by other 5G-internal effects. For comparison, the wired test case from the baseline scenario is shown with 8.2 ms delay for the 99.9% target.

The impact of RTT on the number of transmissions and the resulting C2C communication efficiency is depicted in Figure 38. We observe that an increase in the RTT reduces the communication efficiency. With longer RTT, the PLC can process fewer request and feedback cycles, resulting a lower total number of transmissions within the same time frame. Taking the wired case as the baseline, we observe a drop in number of transmissions by 10% when using the wireless 5G network without background traffic. Introducing additional background traffic decreases this value to 89%, 84% and 81% for 10 Mbps, 50 Mbps and 90 Mbps background traffic, respectively.

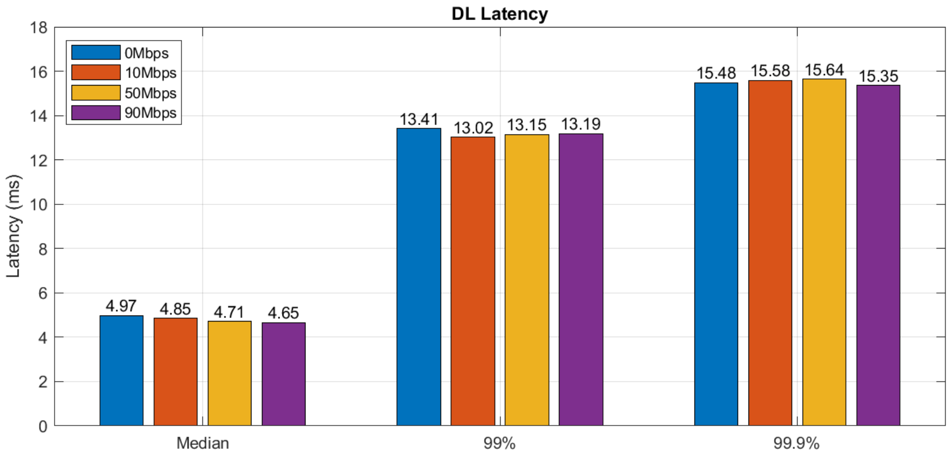

Figure 35: Downlink latency.

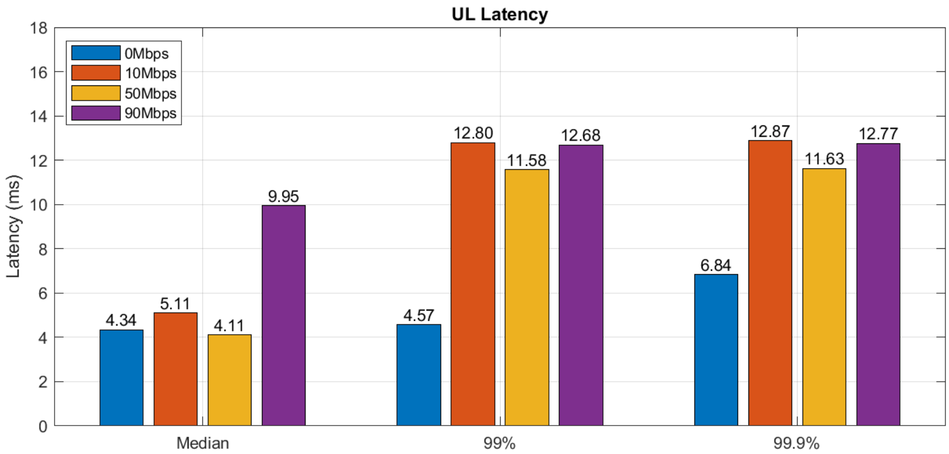

Figure 36: Uplink latency.

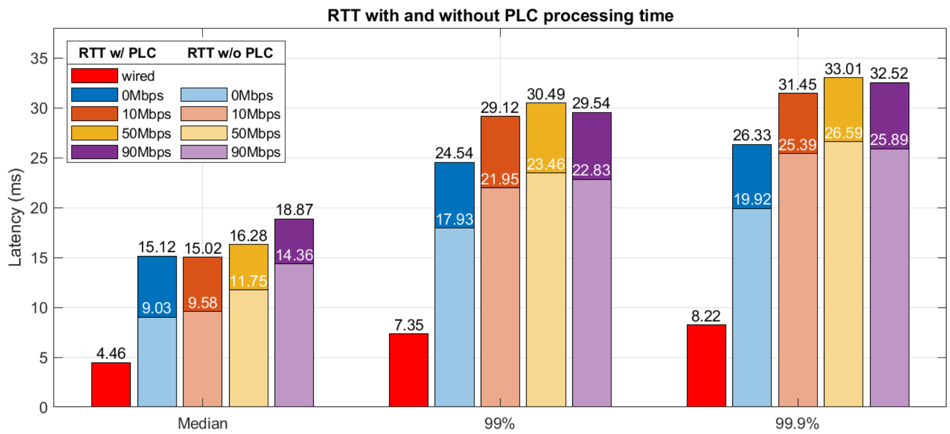

Figure 37: Round trip time with and without PLC processing time.

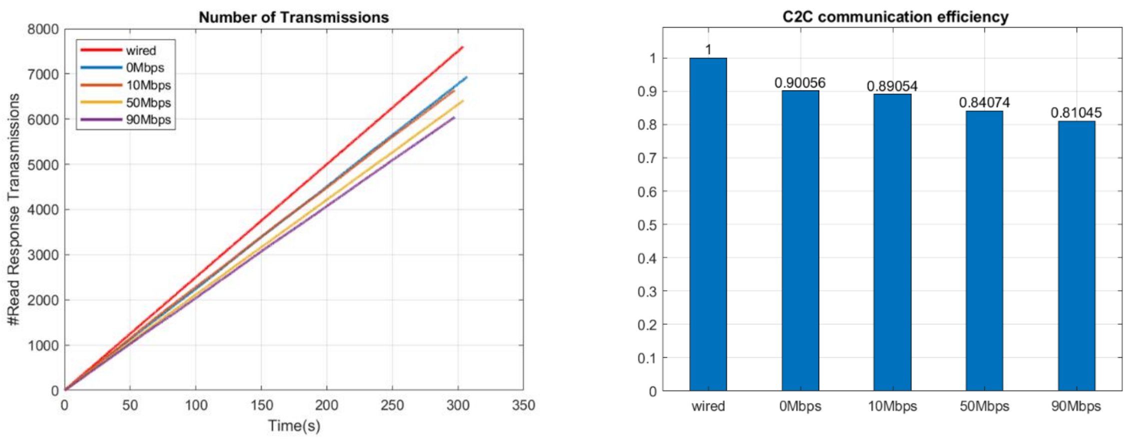

Figure 38: C2C communication efficiency with uplink background traffic.

Downlink background traffic

The 5G deployment supports in DL 10 times higher bandwidth compared to UL. Hence, we increase the background traffic step-wise in 100 Mbps increments from 0 Mbps to 900 Mbps. In general, the downlink background traffic has no strong impact on the C2C traffic performance. One explanation is that the load difference for the C2C traffic with 25 kbps and the background traffic with up to 900 Mbps is so high that queueing and scheduling delays inside the 5G system do not really impact the sparse C2C frames.

Detailed results on the latency are provided in Figure 39 for DL and Figure 40 for UL. We can observe a very slight increase in latency between low background traffic (0 Mbps, 100 Mbps) and higher background traffic (500 Mbps, 900 Mbps). The worst-case latency for 99.9% of the packets are in the range of 16ms to 17ms. In UL, no impact from the DL background traffic is expected. The slight differences in the measurements are caused by other factors, for example alignment delays to the TTI size of the 5G system, or alignment delays to the cycle time of the PLC. Similarly, the RTT results do not differ much between the test cases, showing a consistent worst-case RTT of 25.5-28ms with and 20-21ms without PLC processing time for the 99.9% target (cf. Figure 41).

Since the RTT is not heavily impacted by the background traffic, we also do not observe a strong reduction in the number of transmissions over the test duration. From Figure 42 we can see that the C2C communication efficiency compared to the wired test case decreases to 85.5%-88.5% when introducing background traffic in DL. Comparing this result to the UL background traffic test case in Figure 38, it becomes apparent that UL cross-traffic has a stronger impact on the efficiency.

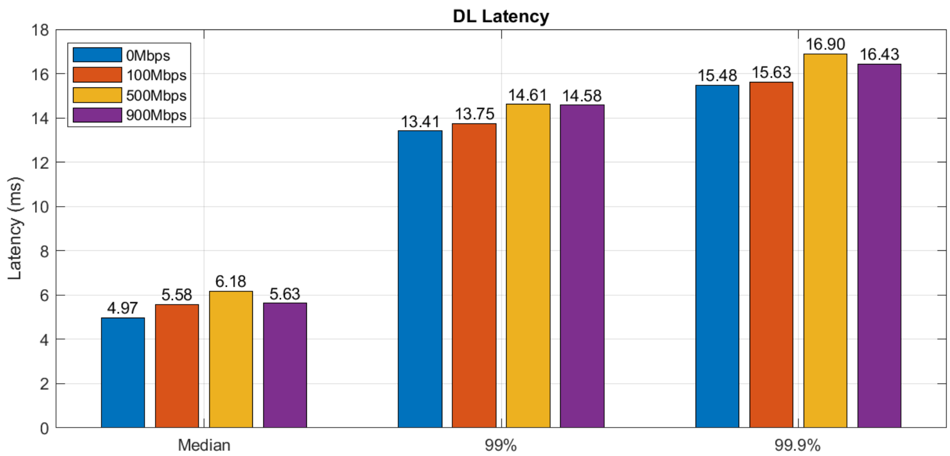

Figure 39: Downlink latency.

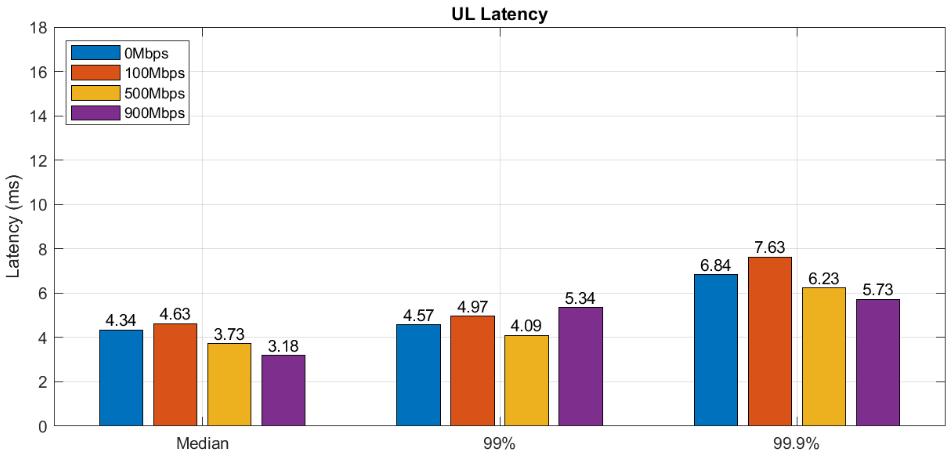

Figure 40: Uplink latency.

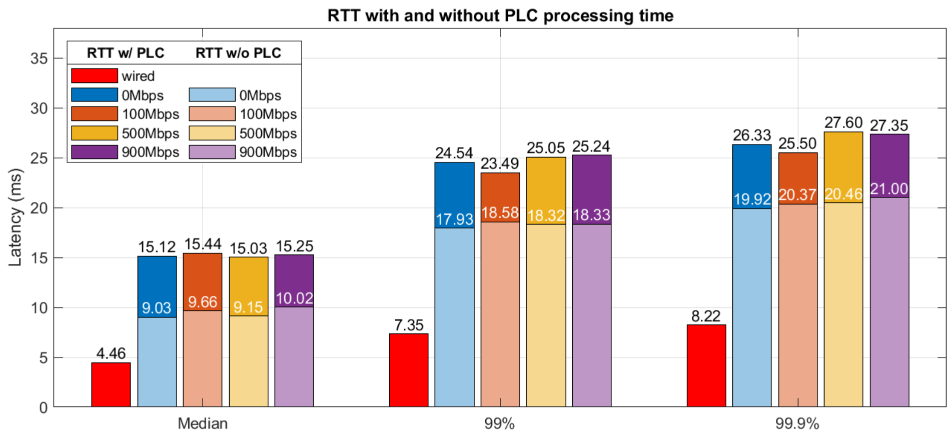

Figure 41: Round trip time with and without PLC processing time.

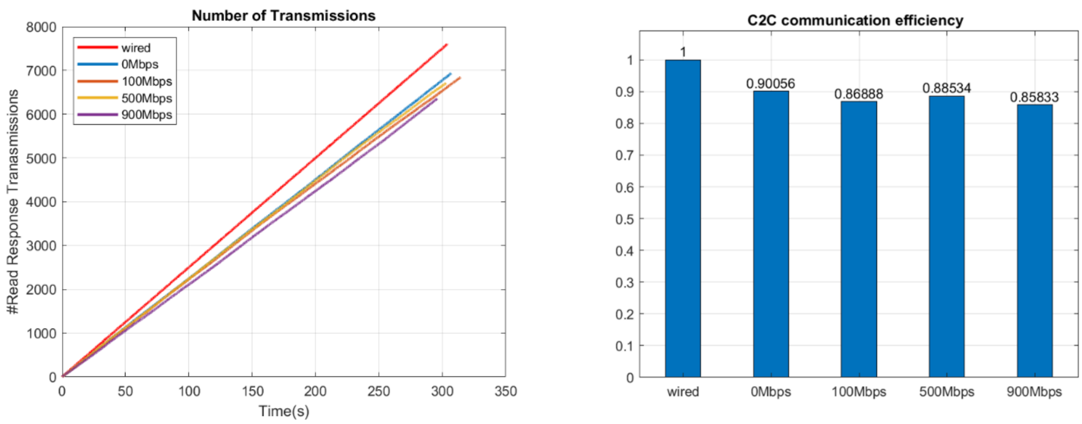

Figure 42: C2C communication efficiency with downlink background traffic.

4.6.3 Discussion

When the controller traffic and background iPerf traffic have the same default priority in the network, which is the case in this validation scenario, the background traffic has an impact onto the controller traffic. It has been shown that the level of impact of the background traffic onto the controller traffic, which is expressed in terms of increased RTT of transmitting the controller packets, depends both on the amount of the background traffic as well as on communication direction. The RTT of the controller packets, in turn, has a direct impact onto the C2C communication efficiency. From the validation scenario 1, we have seen that using 5G without background traffic, the efficiency reduced already by 10%. With background traffic, we have shown in Figure 38 and Figure 42, that the efficiency further decreases by at most 9% and 4.5% for UL and DL, respectively. Thus, it can be concluded that a certain level of controllers' performance degradation in terms of reduced C2C communication efficiency is to be expected when allocating the same priority to the controller traffic and the other kind of background traffic in the same network.

4.7 Validation scenario 3: C2C communication with background traffic and without controller traffic prioritization under different wireless channel conditions

4.7.1 Description of the scenario

Figure 43 shows the high-level architecture of this validation scenario, which resembles the previous validation scenario with the only difference that the 5G UE connected to the PLC A is co-located with the 5G UE connected to the iPerf client. Together with the serving gNB, the 5G UEs are located in the same corridor in the factory, having line-of-sight visibility with that gNB and being around 30 meters away from it, which means that the distance of the 5G UE connected to the PLC is 27 meters further away from the 5G radio dot compared to the evaluation in Section 4.6. Like in the validation scenario 2, the traffic of the PLCs is not prioritized over the iPerf traffic in this validation scenario and in total of 18 experiments have been conducted (i.e. 9 experiments per direction each with 5 minutes duration) with different amount of uplink and downlink background UDP traffic. The main objective of this validation scenario is to analyze the impact of the wireless channel condition by increasing the distance between the 5G UE and the 5G radio dot. In addition, the background iPerf traffic onto the traffic of the PLCs is investigated. This is achieved by gradually increasing the background iPerf traffic and calculating the efficiency of the C2C communication.

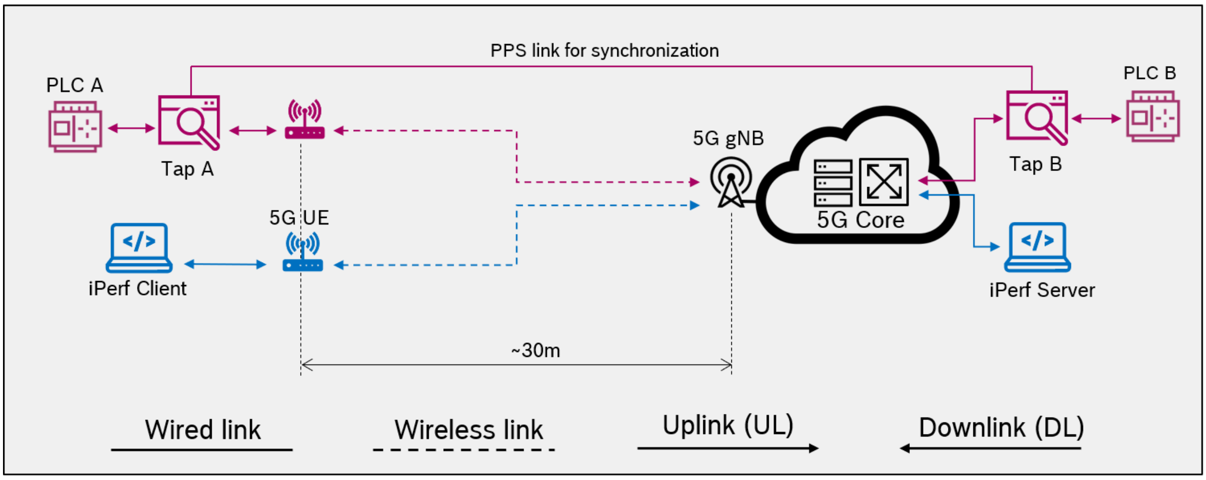

Figure 43: Setup of the validation scenario 3.

4.7.2 Validation results

The results on the validation scenario 3 are very similar to the validation scenario 2. In fact, we were not able to measure any performance degradation on the C2C traffic caused by the increased distance between the UE and gNB. This is due to similar wireless channel conditions (e.g., RSRP, RSRQ) of the 5G UEs connected to the PLC A in Figure 34 and Figure 43, as shown in Figure 44. We only show the test case for the UL background traffic test cases in the following.

	a)	b)
PLMN	26272	26272
Band	5G (n78)	5G (n78)
Channel	647424	647424
BandWidth(MHz)	100Mhz	100Mhz
PCI	79	79
NCI	N/A	N/A
TAC	N/A	N/A
4RX RSRP(dBm)	-57.90 / -59.83 / -55.13 / -53.23	-68.03 / -71.10 / -69.00 / -66.82
RSRQ(dB)	-10.47	-10.42
SSB-SINR(dB)	33.50	29.50
UL TX power(PUCCH)	N/A	-37
UL TX power(PUSCH)	-19	-10
CQI	15	12
SSB-Beam ID	0	0
DL MCS	25	23
DL Modulation	256QAM	256QAM
DL layers(DL rank)	4	4
DL RB number	16	273

Figure 44: Snapshots of the wireless channel status of the UEs connected to the PLC A in the validation scenarios 3 (a) and 4 (b).

The latency results in Figure 45 shows that the DL is unaffected by the UL background traffic. The UL latency in Figure 46 ranges from 6.8ms to 15.5ms for the 99th percentile when increasing background traffic from 0 Mbps to 90 Mbps. The higher variance observed here is caused by other effects such

5G-internal alignment delays and queueing delays that are not traceable with our end-to-end performance tests. Accordingly, also the RTT results, shown in Figure 47, comprising of the combination of UL and DL, do not reveal any significant differences to validation scenario 2. Finally, the C2C communication efficiency depicted in Figure 48 shows identical performance for the different UE locations for low background traffic. For higher background traffic some minimal variations can be observed which again are explained by 5G effects or alignment delays on the PLC that are not related to the positioning of the UEs.

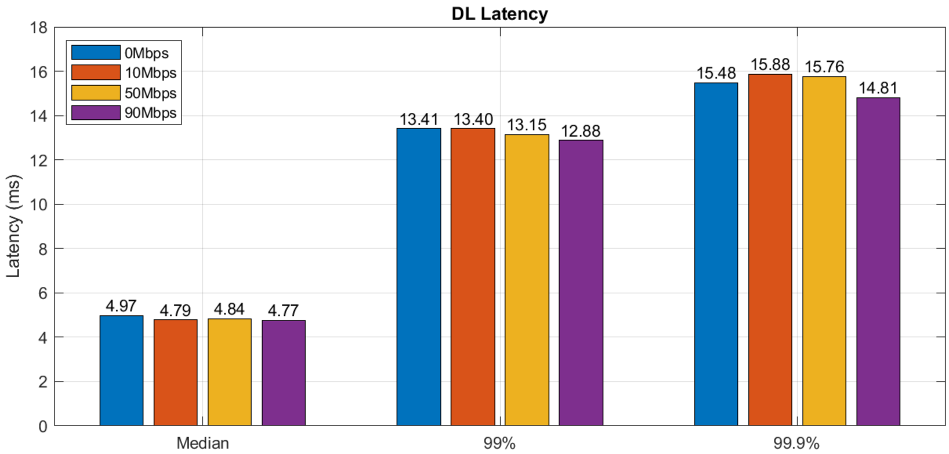

Figure 45: Downlink latency.

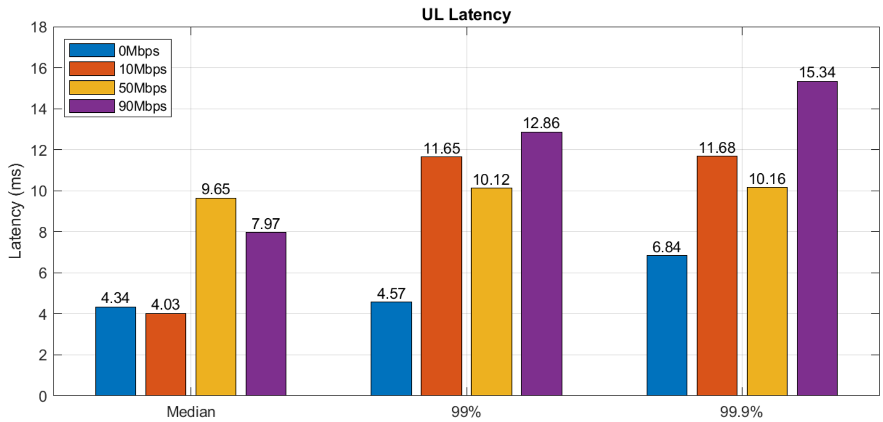

Figure 46: Uplink latency.

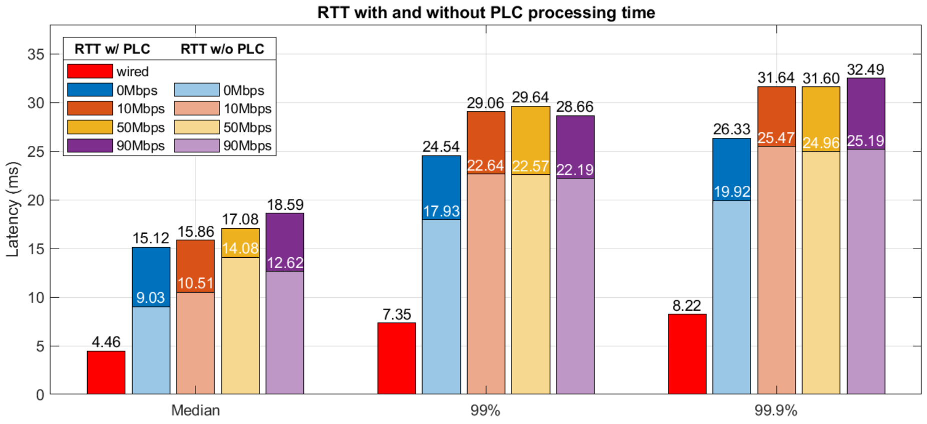

Figure 47: Round trip time with and without PLC processing time.

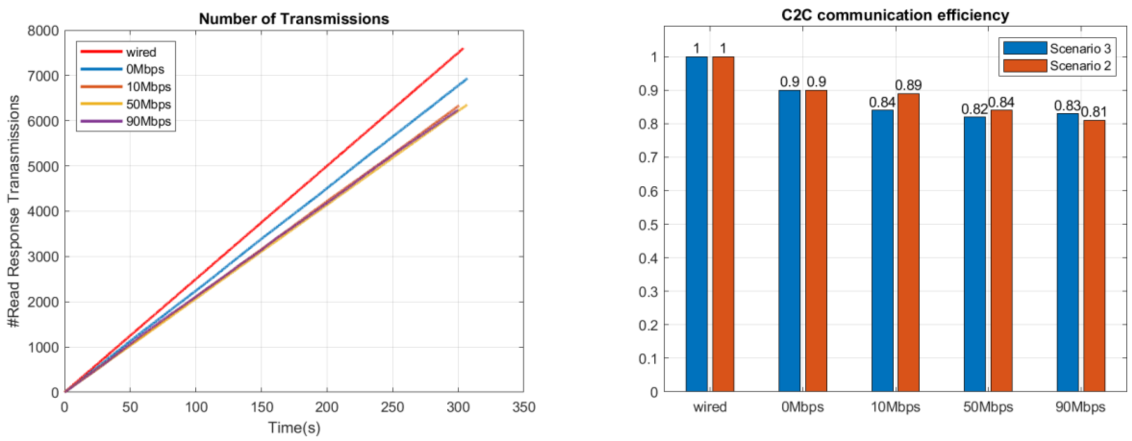

Figure 48: C2C communication efficiency with downlink background traffic for a 5G UE being close (Scenario 2) and far (Scenario 3) form the 5G radio dot.

4.7.3 Discussion

This validation scenario showed that the deployment of the 5G network is well optimized for the indoor environment. This concerns the performed radio planning, placement and orientation of multiple radio dots, as well as the transmit power. Although the results in this section only show performance measures for two sample locations, our experiments have shown that the 5G network achieves very good coverage across the factory hall. The conclusion from this test case is that at any given location, we can expect similar performance of the network without any significant performance degradation due to weak signal reception.

4.8 Validation scenario 4: Prioritized C2C communication with unprioritized background traffic

4.8.1 Description of the scenario

The physical setup of this validation scenario shown in Figure 49 and is identical with that of the validation scenario 2 in terms of the location of the used hardware devices. The difference of this scenario, however, lies in the fact that the traffic of the PLCs is prioritized over the iPerf traffic. The employed QoS configuration enables latency protection for a prioritized traffic flow. As a result, a consistently low transmission latency can be maintained for the prioritized traffic flow even in the presence of a high background traffic load.

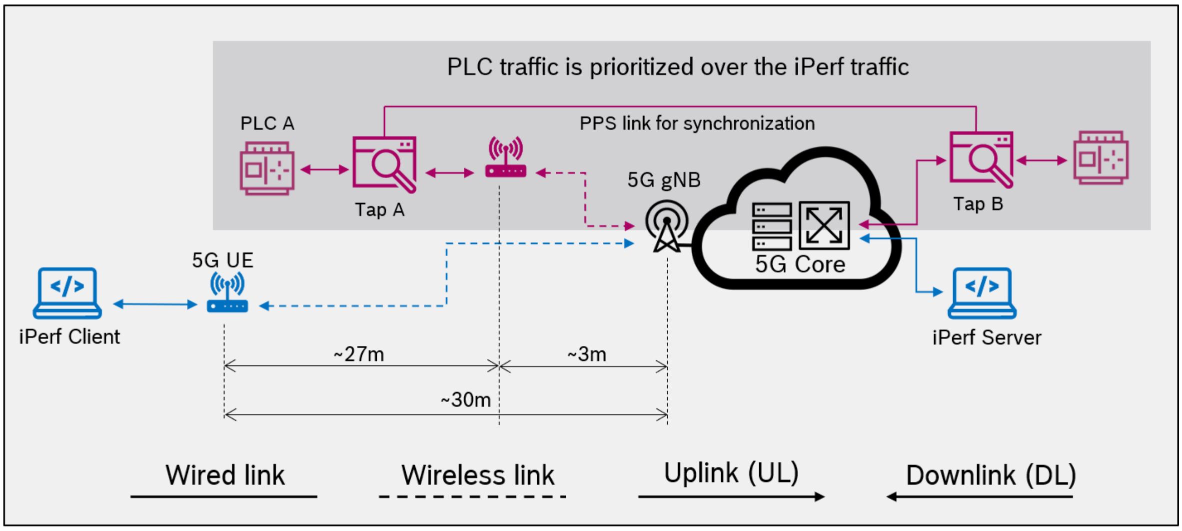

Figure 49: Setup of the validation scenario 4

4.8.2 Validation results

In the following, we compare the results from this scenario with validation scenario 2, where no traffic prioritization is in place. Again, we focus on the UL traffic in the presence of UL background traffic. In DL, we see in both scenarios 2 and 4 no effect of the background traffic on the C2C communication (c.f. Figure 39, Figure 50).

Unlike in scenario 2 (without prioritization), we can observe now that UL background traffic has no observable impact on the controller traffic. The controller traffic with prioritization maintains a worst-case delay of around 8.5ms, as shown in Figure 51, compared to the non-prioritized scenario with up to 13ms delay in the worst case. Due to limited impact of the background traffic on the prioritized controller traffic, also the RTT is around 5ms lower compared to the non-prioritized scenario.

Regardless of the traffic prioritization, the C2C communication efficiency is at least 10% lower in the 5G network compared to the wired setup. With traffic prioritization, however, the C2C communication efficiency only drops negligibly and maintains at least 88% regardless of how much background traffic is introduced (c.f. Figure 53). This is contrary to what we have seen in scenario 2, where efficiency decreases steadily for increasing amount of background traffic.

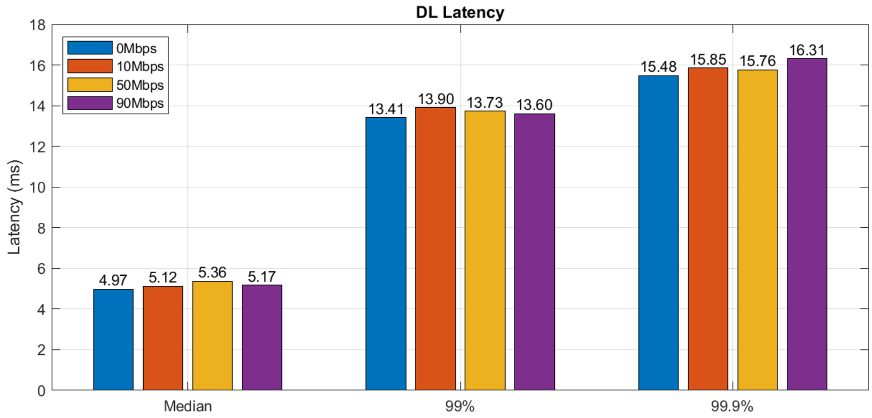

Figure 50: Downlink latency.

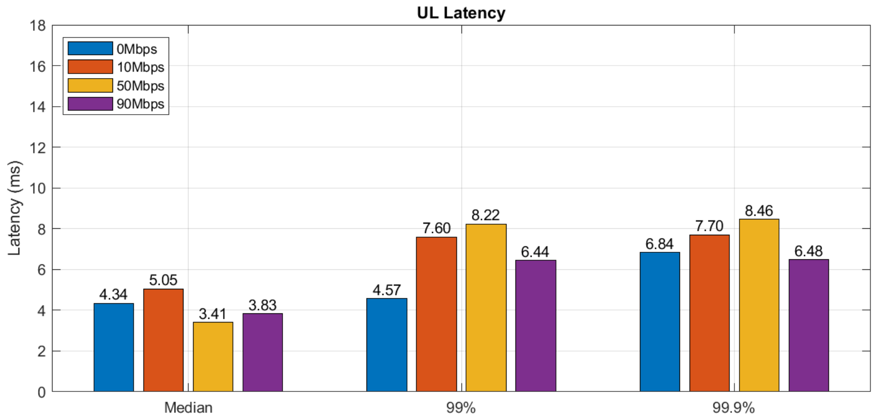

Figure 51: Uplink latency.

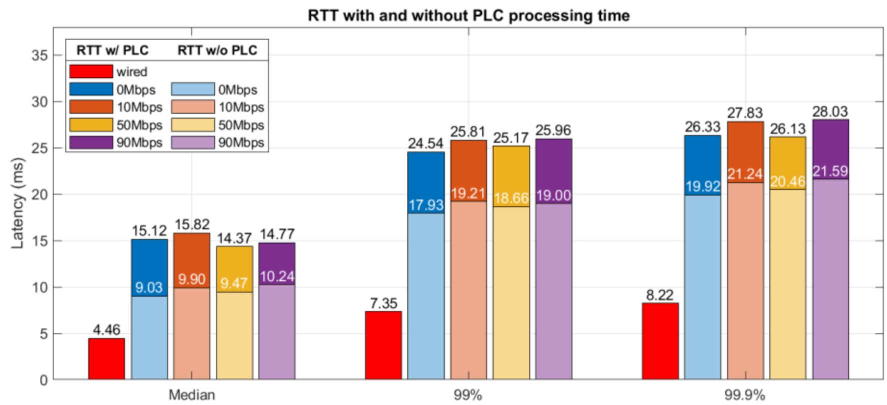

Figure 52: Round trip time with and without PLC processing time.

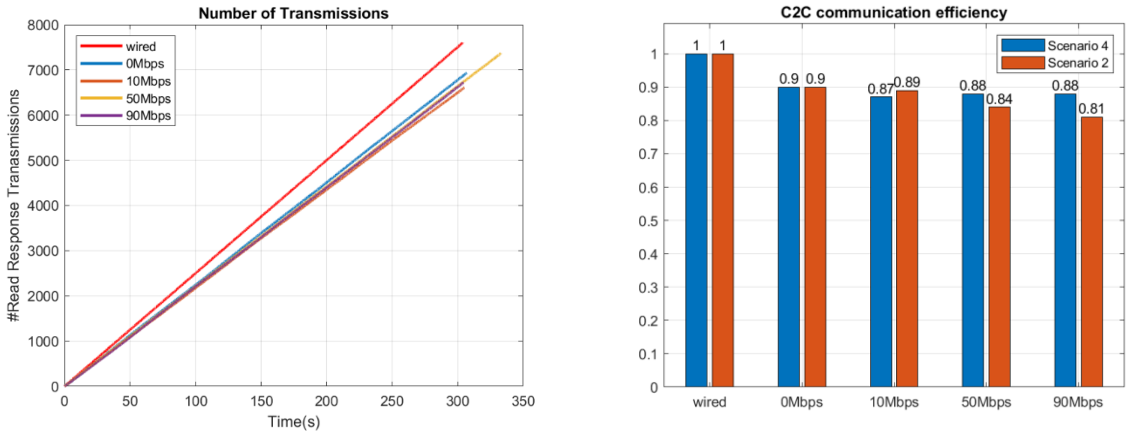

Figure 53: C2C communication efficiency with uplink background traffic.

4.8.3 Discussion

It has been demonstrated that by prioritizing the controller traffic over the background traffic, the impact of the latter on the former can be significantly reduced. However, it has to be mentioned that the employed traffic prioritization feature by QoS configurations is specifically designed latency protection of low-latency applications. Thus, the level of impact of the background traffic on the controller traffic is highly dependent on (i) if the controller traffic is prioritized over the background traffic and (ii) if yes, how the traffic prioritization is executed in the 5G network and especially inside the RAN to address the low-latency requirements.

4.9 Validation scenario 5: Machine-to-server communication with background traffic without machine’s traffic prioritization

4.9.1 Description of the scenario

Figure 54 shows the setup of this validation scenario involving machine-to-server communication. The original plan with this application was to test the machine-to-server communication over the 5G network using the industrial machine Rudolph F30 and its backend server with the latter being located in the Bosch network. This machine-to-server communication would require interconnection between the 5G network and the Bosch network since the communication goes beyond the 5G network. Unfortunately, it has not been possible to enable the communication between these networks due to security requirements at the plant as detailed in the following.

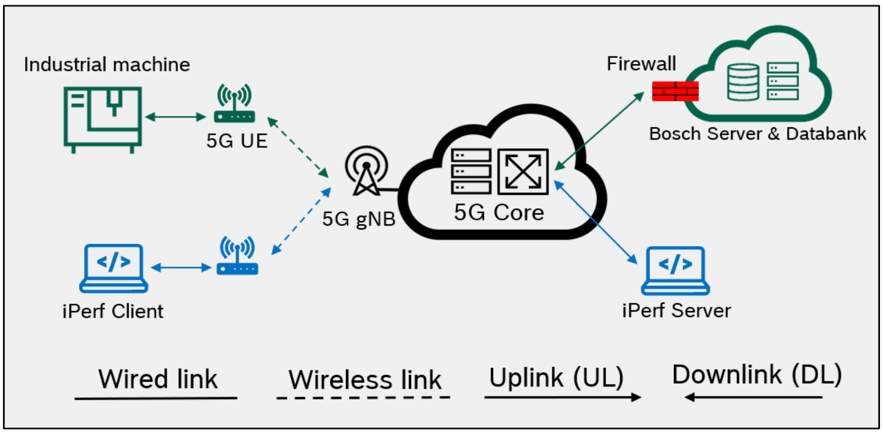

Figure 54: Original setup of the validation scenario 5, which has not been implemented due to the security requirements at the plant.

There are security standards and well-defined guidelines (blueprints) at Bosch for integrating and operating existing communication technologies used at the Bosch plants. Also, there is a security zoning concept implemented at Bosch plants, which allows for traffic segregation between different industrial devices in the same or different networks. The security zones have different security requirement and, therefore, require different security measures (e.g., ACL implementation, traffic filtering with a firewall, traffic segregation with VLANs, etc.) to implement these requirements. Depending on the purpose of the industrial/networking devices, they, for example, may only be allowed to communicate with a subset of industrial/networking devices within a specific network subnet at the plant.

Private 5G mobile network technology is considered to be a new technology to be used for industrial applications. Therefore, the integration of 5G into the industrial network requires a new security blueprint development and implementation, which is a time-consuming and challenging process. Since the blueprint shall be Bosch-wide applicable and not a project-specific, it has to undergo many review and approval processes. While over the course of the 5G-SMART project a security blueprint has been developed and approved at Bosch for integrating 5G within industrial networks, due to its challenging and time-consuming requirements and dramatically increased delivery time of IT-systems (e.g., firewall), it has not yet been possible to implement the blueprint at the semiconductor factory in Reutlingen. It has therefore been decided to change the setup of this validation scenario.

Figure 55 shows the modified setup of this validation scenario with emulated machine-to-server communication. In this setup, the two green laptops represent the industrial machine and its backend server, respectively, with the one being connected over a 5G router to the UE side of the network, and the other being directly connected to the network on the 5G core network side. In this validation scenario, we stream the data collected from the industrial machine Rudolph F30 to the server over the 5G network for 5 minutes, mimicking the real machine-to-server communication. In parallel, we generate various amount of background traffic (e.g., from 10 Mbps to 90 Mbps with 20 Mbps steps with different test runs) with iPerf3 tool running on the blue laptops and repeat the test. In this validation scenario, both the iPerf3 stream and industrial machine stream have the same default priority and the objective of this scenarios is to evaluate the impact of the background traffic onto the emulated machine-to-server communication.

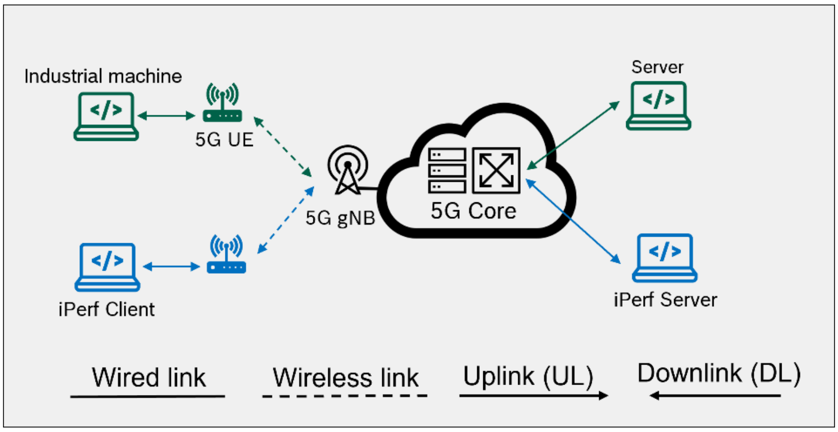

Figure 55: Modified setup of the validation scenario 5, which has been implemented.

4.9.2 Validation results

Generating F30 machine traffic

To generate the data stream, we first captured the real traffic stream between the industrial machine and the server backend during regular operation. The machine processes one wafer at a time and collects several image files that are evaluated on the machine itself. It searches for impurities and manufacturing flaws such as scratches on the surface of the wafer. Depending on the properties of each individual wafer, the process time per wafer and the amount of data generated varies. Once the inspection of a wafer has concluded, the machine pushes all collected image files, as well as further process data, in a single batch to the server backend using TCP, which is then stored there for quality control.

To recreate this behavior with the emulated setup, we generated the same traffic pattern by using the original wafer image files and sending them from one laptop to the other in uplink direction.

Measurement results

An excerpt on the data rate measurement for the emulated machine traffic is depicted in Figure 566. The processing of each wafer is distinguishable by the spike in data rate when the batch of images is transmitted at the end of the inspection process. The measurements have been conducted in a first step over a wired Ethernet connection between the machine and server. The results thereof represent the current state in the factory. When operating the machine over 5G, we notice a slight increase in the width of each transmission burst, resulting in an overall longer time to send the image batch to the server backend.

This becomes more obvious in the following Figure 577, which depicts the transmission time for the Ethernet and 5G setups, respectively. For each wafer, corresponding to the data peaks in the preceding figure, the varying transmission times are depicted. This difference between individual wafers is due to the various data sizes. The more imperfections the machine detects, the more images will be collected and transmitted. In this test run, the number of images per wafer ranges from 21 to 88 with a total batch size ranging from 10 to 33 Mbyte. The transmission time in the wireless setup is roughly 4 times longer compared to the wired scenario. One explanation for this observation is the bandwidth in the UL of the 5G system that limits the rate at which the images can be sent to the server backend. As in the C2C communication test case, also in the machine-to-server communication validation scenarios no packet drop has been observed.

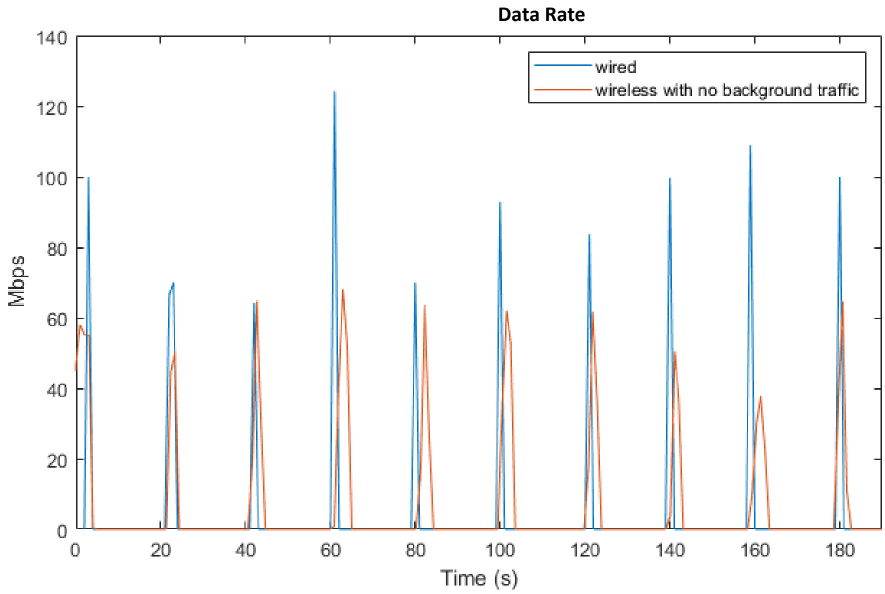

Figure 566: Data rate of F30 machine to server backend when connected over Ethernet (wired) and over 5G (wireless).

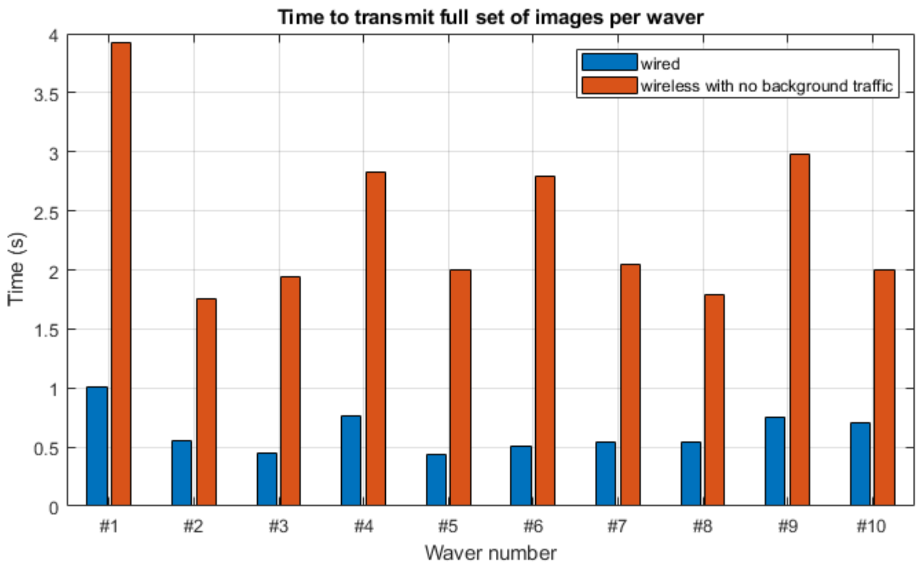

Figure 577: Time to transmit the full set of quality control data (i.e. images) per wafer when connected over Ethernet (wired) and over 5G (wireless).

In the following, additional background traffic in UL is introduced to test the impact on the machine-to-server communication. In this step, no traffic prioritization is in place. Given the variation of data size per wafer, the average transmission time per image is considered, which allows for a better comparison of the results. Figure 588 shows the average transmission time per image, represented as a boxplot over all wafers within the 5-minute trace. The lower and upper end of the box represents the 25th and 75th percentile, respectively, while the whiskers the minimum and maximum values, respectively. The red bar shows the median. The first two boxplots represent the wired and wireless test case without background traffic, which are also partially shown in Figure 577. With increasing background traffic (10, 50, and 90Mbps), the transmission duration increases, for the worst-case scenario, almost doubling the required transmission time per wafer compared to the 0 Mbps case. This is mostly due to the fact that both the machine traffic and the background iPerf3 traffic have the same default priority in the RAN, resulting in queueing delays inside 5G.

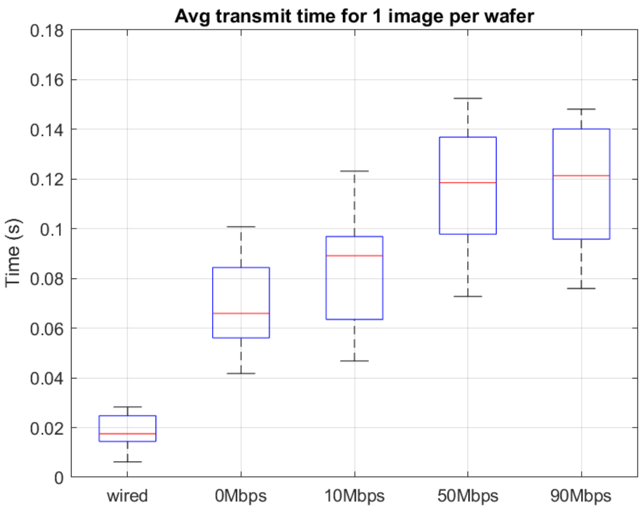

Figure 588: Average transmission time per image. Industrial machine’s traffic is not prioritized over the background traffic in UL.

4.9.3 Discussion

The key takeaway from this validation scenario is that transmission times in the machine-to-server communication increase when operated over a wireless network. It further increases when introducing background traffic. There are several potential explanations for this observation. The 5G system has an upper bound of the uplink data rate (around 90Mbps) under the given wireless channel condition. This can be mitigated, for example, by allocating more frequency bandwidth for the uplink transmission. With the presence of background traffic, the transmission time of the quality control data (i.e. images) per wafer increases even more due to the fact that both the background iPerf3 stream and the industrial machine stream share the same priority in the 5G network. In case of very high background traffic, the joint data rate exceeds the throughput limit of the 5G network, which impact the transmission time. Packet losses were observed in these tests, but since TCP is used as the streaming transport layer protocol, lost packets are recovered by TCP and all image files are retrieved error-free at the receiver eventually. The transmission time increase due to packet losses and retransmissions can be mitigated, for example, by prioritizing the machine’s traffic over the background iPerf3 traffic, which is tested in the validation scenario 6. However, it is important to mention that the increase of the average transmission time per image or batch of images per wafer has no impact on the wafer inspection process. This is because after a wafer inspection, the collected data (i.e., images) is stored in the local memory of the industrial machine Rudolph F30 and the inspection process of the next wafer starts, while the collected data is sent to the backend server in the background.

4.10 Validation scenario 6: Prioritized machine-to-server communication with background traffic

4.10.1 Description of the scenario

This validation scenario shown in Figure 59 resembles the previous validation scenario, with the only difference being that the machine-to-server traffic is prioritized over the background iPerf3 traffic. The used QoS configuration enables traffic prioritization to ensure that the throughput of one high-priority data flow is protected against other traffic flows with lower priority. As a result, the throughput of the prioritized traffic is not affected by a high background traffic load.

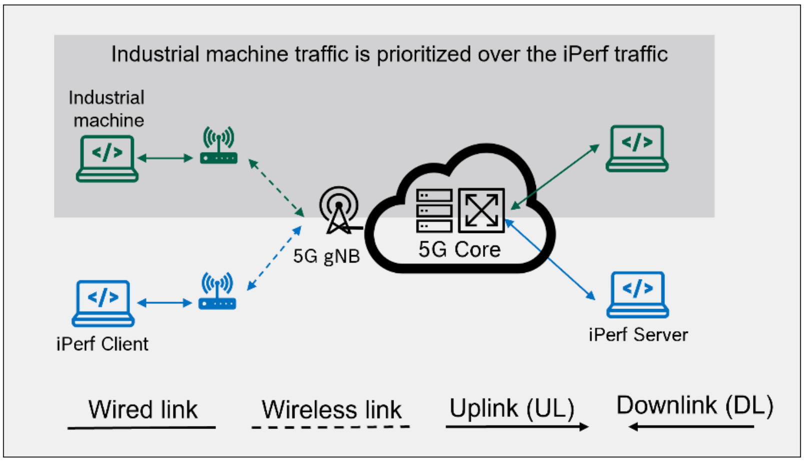

Figure 59: Setup of the validation scenario 6.

4.10.2 Validation results

The same test procedure as in Section 4.9 has been conducted in this case. By introducing traffic prioritization, we can observe that the background iPerf3 traffic has almost no effect on the transmission time of the Rudolph F30 data stream. Only in the extreme case of 90 Mbps background traffic, the transmission time increases slightly for each image as shown in Figure 590. Figure 601 displays the overall transmission time per wafer, comparing the scenarios with no background traffic, with 90 Mbps background traffic without prioritization of the machine traffic, and 90 Mbps background traffic with prioritization of the machine traffic. Again, we can observe that background traffic has almost no impact on the data stream, increasing the transmission duration by at most half a second.

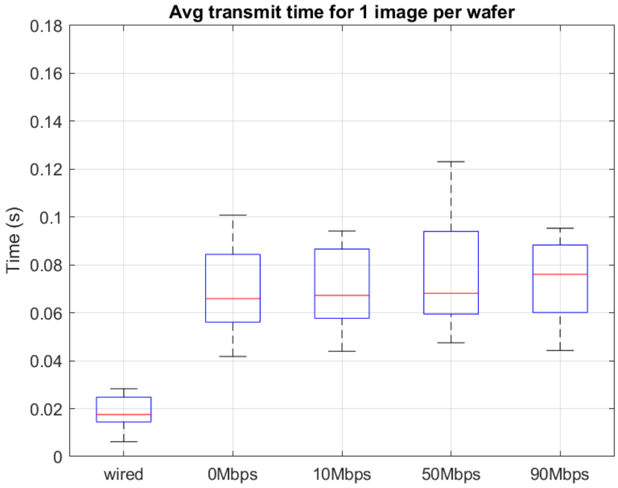

Figure 590: Average transmission time per image. Industrial machine's traffic is prioritized over the background traffic in UL.

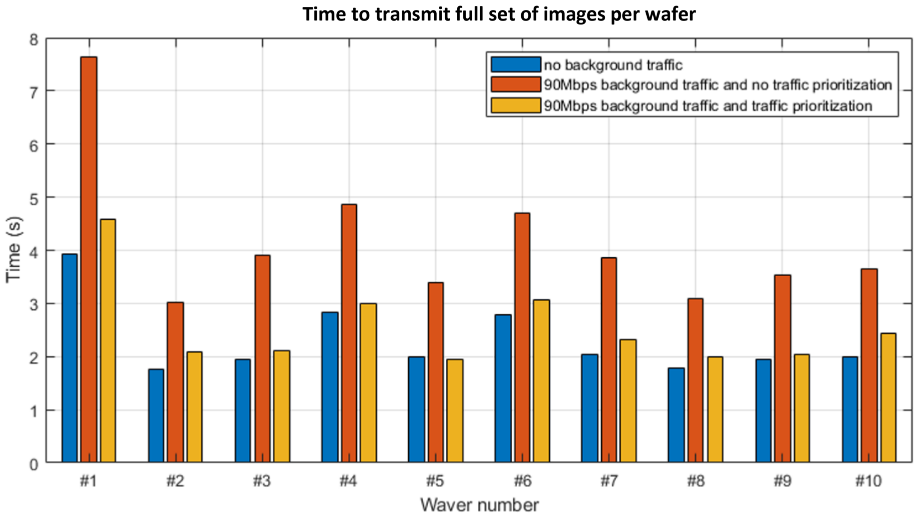

Figure 601: Time to transmit the full set of quality control data (i.e. images) per wafer with and without traffic prioritization.

4.10.3 Discussion

The important takeaway from this use case is that this type of application can be supported by 5G regardless if traffic prioritization is applied or not. The application protocol can adjust to the available bandwidth such that all images are always received error-free. With the employed QoS differentiation and corresponding prioritization, we can however significantly curtail the data transmission time in the presence of cross-traffic, which nonetheless has no impact onto the wafer inspection process.

4.11 Conclusion

In the TSN/I-LAN over 5G use case, we have validated the capability of the 5G network for supporting the traffic of industrial LAN applications, which include the C2C communication and machine-to-server communication. Eight KPIs have been defined to evaluate *flexibility*, *sustainability*, *quality*, and *productivity*, which have been set by the factory management as industry goals for this use case. Based on the results, we can conclude that the 5G network provides more *flexibility* compared to the wired network when it comes to customizing and modifying the production cells by significantly reducing the commissioning and relocation time of industrial machines like Rudolph F30. We can also conclude that replacing the wired communication technology with the 5G wireless technology in this use case has a positive impact on its *sustainability*, resulting in a significant amount of annual cost savings due to, for example, curtailed amount of cabling, maintenance effort, and network sockets. As for the industry goal *quality*, which is expressed in terms of the number of packet drops in the C2C and machine-to-server communication, it can be concluded that both wired and 5G wireless network deliver the same level of communication quality, both resulting in no packet loss. Finally, for what concerns the *productivity* industry goal, which is expressed in terms of the C2C communication efficiency, it can be concluded that the productivity slightly drops when using the 5G network compared to the wired network.

The applications considered in the TSN/I-LAN over 5G use case have been tested in various validation scenarios. Specifically, the C2C communication has been tested in four validation scenarios. In the validation scenario 1, the C2C communication over the 5G network has been tested without any background traffic in the network, demonstrating that the RTT of the controller's packet is many times more compared to the traditional C2C communication over the wired Ethernet network, which led to a reduced C2C communication efficiency. In the validation scenario 2, the C2C communication over the 5G network has been tested with presence of various amount of background iPerf traffic, where both the controller traffic and the iPerf traffic had the same default traffic priority and shared the same RAN. It has been concluded that the background traffic increases the communication delay both in the uplink (more than in the downlink) and downlink directions, resulting in a lower C2C communication efficiency compared to that in the validation scenario 1. In the validation scenario 3, the controller connected to the 5G UE has been moved further away from the closest RD and then the same test of the validation scenario 2 has been performed with the goal of evaluating the impact of the changed wireless channel condition onto the C2C communication. However, it has been found out that there was no significant change in the wireless channel conditions when moving the 5G UE around 30 meters away from the RD versus having the 5G UE around 4 meters away from the same RD. Therefore, the observed performance of the C2C communication was similar to the one in the validation scenario 2, leading to a conclusion that with the proper radio planning in the factory, a similar C2C communication performance can be achieved, regardless of location of the 5G UEs in the factory. Finally in the validation scenario 4, the C2C communication has been prioritized over the background traffic with the main conclusion that the used traffic prioritization mechanism significantly curtails the impact of the background traffic onto the controller traffic in terms of the packet transmission delay in the 5G network.

Two validation scenarios have been tested in the machine-to-server communication as a part of the TSN/I-LAN over 5G use case. In the first scenario (validation scenario 5), we have machine-to-server communication together with background traffic both having the same default priority in the RAN. It has been demonstrated that the data (i.e., wafer images) transmission time is more over the 5G network without any background traffic compared to the wired network and significantly more increases with presence of background traffic when both machine and backgroud traffic share the same priority in the RAN. It has also been demonstrated that by prioritizing the machine's traffic over the background traffic, as tested in the second scenario (validation scenario 6), the impact of the background traffic onto the machine's traffic in terms of the data transmission time can be mitigated. Nonetheless, it has been concluded that the observed data transmission time regardless of machine's traffic prioritization has no impact onto wafer inspection process, which is the main application in this machine-to-service communication.

5 Summary of all activities and results from the Reutlingen trial site

Three major objectives have been defined for the 5G trial at the Bosch semiconductor factory in Reutlingen. The first objective of WP4 was to evaluate the deployability of 5G system in the semiconductor factory, in terms of radio propagation, coverage, and electromagnetic compatibilities (EMC). To achieve this objective, we have performed wireless channel measurements both in a mid-band frequency centered around 3.71 GHz in the production area of the plant in order to get an insight into the channel characteristics in this rather specific radio propagation environment full of narrow corridors, metallic object, production machines and mobile robots. Overall, it has been concluded that in terms of coverage in the factory, the 3.71 GHz channel measurements show promising results. The details of the wireless channel measurement scenarios as well as the results can be found in 5G-SMART Deliverable D4.2 [5GS22-D42]. These measurements provide useful insights into this quite challenging radio propagation environment and will be taken into consideration when planning a 5G network deployment in the factory for productive use. Since the channel measurements have been performed in area of wafer processing above, the obtained results are excellent base for further activities in semiconductor production e.g. for clean room space extension.

Before deploying a 5G mobile network in the plant for productive use, it is of paramount importance to make sure that the 5G signal does not impair the semiconductor production process. As a part of the first objective, EMC tests have been designed and performed to analyze the impact of the 5G signal onto the running production of various types of semiconductor products at the Bosch factory in Reutlingen. The EMC tests have been conducted in all the three test areas, i.e., wafer test, final test and sensor backend area, on a subset of semiconductor products, which have carefully been selected to have a representative result at the end. A test network setup has been developed and used to conduct the EMC tests on the shop floor with the carrier frequency was fixed at 3.7 GHz. The measurements on the shop floor have been conducted in uplink (UL), where the 5G device was at close proximity (20cm) to the device under test (DUT) sending data to base station, as well as in downlink (DL), where 5G base station installed in the ceiling was approximately 2 meters away from the DUT, both in 20 MHz and 100 MHz channel bandwidth. Moreover, different electric field strengths (i.e., 12.2 V/m and 38.7 V/m in the UL, 5.5 V/m and 10.9 V/m in the DL) have been used in the EMC tests.

Based on the EMC test results, it has been concluded that with the actual tested samples, deploying a 5G network in the final test and the sensor backend areas can be considered without any concerns, while the wafer test area cannot be considered unless actions are taken to ensure EMC. It is important to mention that negative effects of 5G signals on untested devices, or in future test processes, cannot be fully excluded. For the details on the EMC tests and the results, we refer the reader to 5G-SMART Deliverable D4.2 [5GS22-D42]. The outcome of these EMC tests represents a valuable input for Bosch when adopting the 5G network technology for productive use since the EMC tests have also been performed in the area of wafer test, which is one of the relevant areas for ensuring a high-quality level.

The second objective of WP4 was to design and install a 5G system testbed at the trial site. Taking into account the learnings from the wireless channel measurements and EMC tests, the 5G system testbed has been designed and deployed at the Bosch semiconductor factory in Reutlingen to validate 5G-based factory automation. This activity led to significant learnings with respect to the design and integration of a 5G network into the existing Bosch-network in fulfilling requirements of IT-security. The 5G network deployed at the trial site is a dedicated, on-premise, and standalone (SA) 5G network based on 3GPP Release 15 specification. More details about the network deployment as well as the trial site constraints can be found in the 5G-SMART's Deliverable D4.1 [5GS20-D41].

The third objective of WP4 was to develop and trial Cloud-based mobile robotics and TSN/Industrial-LAN over 5G use cases. The cloud-controlled mobile robots showed positive impact on the efficiency of the robot operation. Efficiency can be significantly increased by optimizing AGV routes in real factories when cloud-based control is realized over the 5G network. We managed to quantify the gain in terms of mission execution time in real environment with obstacles and multiple paths, which can be directly exploited when the use case is applied in the real production flow. Learnings and insights from the development and validation of the AGV use cases include the importance of certifications on the shopfloor, as well as the fact that the deployed 5G network could easily support all of the tested validation scenarios without any deterioration on the performance of the robots.

The research activity related to the robot HW and SW architectures developed in WP4 has been taken up in further research activities outside the action already. For instance, Ericsson has continued the activity internally, as well as in form of an approved university collaboration to extend the implementation with advanced use cases such as camera-based real-time and highly accurate self-positioning and adding social layer information to the common map from camera to categorize obstacles as humans or objects. The Ericsson internal activity will focus on evaluating the realized trajectories of mobile robots or robot arms and investigating the potential gains of adaptive control mechanisms such as dynamic control cycle time or control domain configurations while the original quality of experience is maintained, which can further improve the efficiency of the 5G network when it comes to connecting large fleet of (mobile) robots.

The main purpose of the TSN/I-LAN over 5G use case was to validate and evaluate C2C and machine-to-server communication over the 5G network. Diverse validation scenarios have been designed and tested for both of the applications in this use case. More specifically, the performance of these application has been evaluated under various load conditions at the RAN in prioritized and non-prioritized traffic scenarios. Moreover, several KPIs have been defined and used to evaluate industry goals such as flexibility, productivity, etc. defined by the plant management in different validation scenarios. It has been demonstrated that without traffic prioritization for the C2C as well as machine-to-server communication, the background traffic may have a significant impact onto the controller traffic and the machine traffic in terms of increased delay, thereby affecting their performance.

Several issues have been encountered during the integration and testing phases of use case. While most of them have been easily resolved and the major part of the tests have been conducted as originally planned, it has not been tested the machine-to-server communication as per the original plan. Since this communication goes beyond the 5G network (i.e., the server of the industrial machine is located in the Bosch network), testing the machine-to-server communication was not feasible due to the IT security requirements at Bosch. As an alternative solution, however, the machine-to-server communication was emulated with the server being located in the 5G network.

All in all, with the activities in the Reutlingen trial site Bosch has gained more hands-on experience with testing and evaluating the applicability of 5G network industrial applications, which allows the Bosch plants to better understand the advantages and limitations of the 5G mobile network technology for industrial applications, which in turn helps with the evaluation of the overall 5G performance for productive use.

References

[Eth] "Bosch Rexroth S20-ETH-BK, S20-ETH-BK"

[Indra] "Bosch Rexroth IndraControl, IndraControl XM22"

[HEBI] "HEBI Robotics, https://www.hebirobotics.com/"

[RDS] The Ericsson radio dot system https://www.ericsson.com/en/portfolio/networks/ericsson-radio-system/radio/indoor/radio-dot-system

[5GS21-D12] 5G-SMART's Deliverable D1.2, "Analysis of business value creation enabled by 5G for manufacturing industries", May 2021.

[5GS20-D11] 5G-SMART Deliverable D1.1, "Forward looking smart manufacturing use cases, requirements and KPIs", June 2020.

[5GS20-D41] 5G-SMART Deliverable D4.1, "Report on design and installation of the 5G trial system in Reutlingen", November 2020.

[5GS22-D42] 5G-SMART Deliverable D4.2, "Report on 5G radio deployability in the factory", May 2022.

[5GS20-D43] 5G-SMART Deliverable D4.3, "Report on the development of 5G use cases", June 2021.

[5GS20-D51] 5G-SMART Deliverable D5.1, "First Report on new technology features to be supported by 5G standardization", May 2020.

[5GS21-D53] 5G-SMART Deliverable D5.3, "First Report on new technology features to be supported by 5G standardization", November 2021.

[ROS] "Robot Operating System", https://www.ros.org/".

[Gazebo] "Gazebo simulation tool, http://gazebosim.org/".

[MiR] "MiR 100, https://www.mobile-industrial-robots.com/en/solutions/robots/mir100/".

[MiRRest] "MiR 100 REST API documentation, https://www.mobile-industrial-robots.com/media/11210/mir_mir100_rest_api_21002.pdf".

[ROS2] "Robot Operating System version 2, https://docs.ros.org/en/foxy/index.html".

[AGVDemo] "Cloud based mobile robotics demo at Ericsson R&D Innovation Days, https://youtu.be/xBWDwmFGPm8, Budapest, Hungary, 10-11 November 2020".

[AMCL] "Adaptive Monte Carlo localization method, http://wiki.ros.org/amcl".

[5GS20-CT] "5G-SMART – 5G Common terminology", June 2020.

[5GACIA-IE] https://5g-acia.org/wp-content/uploads/2021/04/5G-ACIA_Integration-of-Industrial-Ethernet-Networks-with-5G-Networks-.pdf

[Profitap] "Profishark 1G+, https://www.profitap.com/profishark-1g-plus/".

Appendix

List of abbreviations

AGV	Automated Guided Vehicle
API	Application Programming Interface
BBU	Baseband Unit
C2C	Controller-to-Controller
CaaS	Container as a Service
CAD	Computer-Aided Design
CCDF	Complementary Cumulative Distribution Function
CE	Conformité Européenne (French for "European Conformity")
DDS	Data Distribution Service
DL	Downlink
HW	Hardware
I-LAN	Industrial Local Area Network
IRU	Indoor Radio Unit
LiDAR	Light Detection and Ranging
NR	New Radio
PC	Personal Computer
PLC	Programmable Logic Controller
PPS	Pulse Per Second
RD	Radio Dot
RDS	Radio Dot System
ReST	Representational State Transfer
RTT	Round-Trip Time
ROS	Robot Operating System
SLAM	Simultaneous Localization and Mapping
SW	Software
TSN	Time Sensitive Networking
UE	User Equipment
UI	User Interface
UL	Uplink
VM	Virtual Machine
YM	Yield Manager

Table 10: List of abbreviations