



# Article Performance of 5G Trials for Industrial Automation

Junaid Ansari <sup>1,\*</sup>, Christian Andersson <sup>1</sup>, Peter de Bruin <sup>1</sup>, János Farkas <sup>1</sup>, Leefke Grosjean <sup>1</sup>, Joachim Sachs <sup>1</sup><sup>(D)</sup>, Johan Torsner <sup>1</sup>, Balázs Varga <sup>1</sup>, Davit Harutyunyan <sup>2</sup>, Niels König <sup>3</sup><sup>(D)</sup> and Robert H. Schmitt <sup>3,4</sup>

- <sup>1</sup> Ericsson, Torshamnsgatan 21, Kista, 164 83 Stockholm, Sweden; christian.x.andersson@ericsson.com (C.A.); peter.de.bruin@ericsson.com (P.d.B.); janos.farkas@ericsson.com (J.F.); leefke.grosjean@ericsson.com (L.G.); joachim.sachs@ericsson.com (J.S.); johan.torsner@ericsson.com (J.T.); balazs.a.varga@ericsson.com (B.V.)
- <sup>2</sup> Robert Bosch GmbH, 71272 Renningen, Germany; davit.harutyunyan@de.bosch.com
- <sup>3</sup> Fraunhofer Institute for Production Technology IPT, 52074 Aachen, Germany; niels.koenig@ipt.fraunhofer.de (N.K.); r.schmitt@wzl.rwth-aachen.de (R.H.S.)
- <sup>4</sup> Laboratory for Machine Tools and Production Engineering (WZL), RWTH Aachen University, 52074 Aachen, Germany
- \* Correspondence: junaid.ansari@ericsson.com

Abstract: Wireless- and 5G-enabled industrial automation is expected to include a plethora of different applications with a wide variety of requirements. In this article, evaluations are undertaken for the deployment of 5G in realistic industrial production environments with realistic deployment settings. Both deployments using commercial 5G systems and a 5G prototype system including precommercial and standard compliant URLLC functionality have been investigated. Systematic latency and reliability measurements were performed, over the air and in live networks, for different packet sizes, different devices, and networks with different capabilities (at different sites) to characterize the expected performance. The results indicate that today's 5G latency performance significantly depends on packet size, transmission direction (uplink or downlink), and network configuration as well as on the end device's design and capabilities. Our over-the-air measurements also empirically show that 5G technology and future networks have the capability of providing one-way latency of around 1 ms in both uplink and downlink for the various packet sizes tested. It is concluded that the requirements for very low latencies can be achieved with high reliability guarantees, as required in some of the most stringent industrial IoT applications.

Keywords: 5G; industrial automation; latency; empirical results

### 1. Introduction

The development of 5G promises to provide novel opportunities for industrial IoT and smart manufacturing by bringing wireless mobile connectivity with high reliability and low latency to the shopfloor. Today's industrial applications are generally characterized by periodic or cyclic communication, where messages need to be communicated with strict latency bounds. Failure to deliver messages within the given delay bound may lead to service downtime. It is therefore crucial to investigate typical industrial use cases while looking at key performance indicators, including message size, communication service availability, and latency.

In industrial setups, message sizes typically vary from few 10s of bytes to more than 1000 bytes. The required communication service availability is often in the range of 99% to up to 99.9999%. Typically, an industrial application can tolerate a survival time of several transmission cycles, where service can be maintained even if messages are occasionally not delivered within the given latency bounds [1,2]. The availability of a communication service y can be translated into network reliability based on the survival time. For instance, with the approximation described in [1], a network reliability of 99.9% can provide a communication service availability of 99.999%.



Citation: Ansari, J.; Andersson, C.; de Bruin, P.; Farkas, J.; Grosjean, L.; Sachs, J.; Torsner, J.; Varga, B.; Harutyunyan, D.; König, N.; et al. Performance of 5G Trials for Industrial Automation. *Electronics* **2022**, *11*, 412. https://doi.org/ 10.3390/electronics11030412

Academic Editor: Raed A. Abd-Alhameed

Received: 10 December 2021 Accepted: 20 January 2022 Published: 29 January 2022

**Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Copyright:** © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/).

So far, only a few 5G over-the-air (OTA) latency results have been reported for industrial deployments. Some of the recent 5G OTA latency measurement results that have been reported [3–5] were performed in non-industrial environments and consider traffic profile types that are closer to eMBB (enhanced Mobile Broadband) or V2X (Vehicle-to-everything). While some published work [3] provides significant insights into the real-world performance of 5G networks, aspects such as tests in real production halls with traffic profiles derived from real use-cases, RAN configuration effects, and the empirical benefits of using URLLC features have not been covered. A comparison of Wi-Fi 6 with 5G focuses on simplified testing and does not consider industry-application traffic profiles or deployment in a realistic production environment [6]. While the paper concludes in favor of 5G when it comes to coverage and resilience, scalability aspects or 5G URLLC features were been empirically evaluated. Furthermore, using higher percentiles than the 95th (which is more suitable for eMBB-type applications) would be of interest, especially when considering the industry reliability requirements for mission-critical industry applications [6]. Often, round trip times (RTT) are measured for the latency or delay characterization of a 5G network. However, industrial traffic is typically characterized as one-way periodic traffic. Since uplink (UL) and downlink (DL) delays are different and can be optimized or configured differently, RTT measurements are not always representative of industrial scenarios.

This paper contributes to the understanding of the practical aspects of performance in 5G deployments for industries by presenting results from systematic and comprehensive OTA latency measurements. The measurements were conducted in real industry deployments with traffic profiles derived from actual industry use cases, and latency performance results to the 99.9th percentile are presented, separating UL and DL performance. The tests compare different network features, including RAN configurations, use of over-the-top redundancy schemes, as well as the use of URLLC capabilities. Additionally, device-dependency aspects on the end-to-end latency performance are highlighted.

The rest of the article is structured as follows: Section 2 describes the use case requirements and gives an overview of the use case implementations at the different trial sites. In Section 3, a brief overview of the 5G deployments at the different 5G-SMART trial sites is given. Section 4 describes the measurement setups and the evaluation of 5G URLLC features and the combination of 5G with TSN (Time Sensitive Networking) features such as FRER (Frame Replication and Elimination for Redundancy). The test results are summarized and explained in Section 5, and a summary of the article is given in Section 6.

#### 2. Industrial Use Cases and Requirements

A range of use cases were identified and described in, e.g., 3GPP [1], 5G-ACIA [7], and the EU project 5G-SMART [8]. These use cases represent the typical characteristics of today's industry applications in terms of cycle times and packet sizes as well as the requirements for latency, reliability, and survivability.

The evaluation undertaken in this article was carried out for the 5G system deployments at the 5G-SMART project trial sites. For each of the trial sites, a variety of use cases was considered for evaluating and validating 5G for smart manufacturing. These use cases are described below, where their benefits as well as their key requirements are outlined.

### 2.1. Kista Trial Site

At the Kista trial site, 5G-enabled industrial robotics are being validated in three distinct use cases, out of which the following two are most relevant for the performance evaluation carried out in this article: 5G-enhanced collaborative robots and human–robot interaction. The setup is located in an Ericsson smart factory. The setup consists of two stationary robots and a mobile robot, which in the first use case, collaboratively solve a task where material needs to be transported between to robot workstations. The robots are integrated over 5G, allowing major parts of the robot control functionality to be removed from the robot itself and to be placed into the edge cloud. Machine vision via a video system and video analysis in the edge cloud support robot collaboration.

Interconnecting robots over 5G infrastructure and the edge cloud has several benefits, e.g., improved production organization flexibility on a factory floor, support for accelerated changes in the production processes, or a reduced robot footprint by miniaturizing the hardware that needs to stay onboard the robot units.

Another use case aims to explore 5G-enhanced human–robot interaction. One aspect that is considered is whether safe real-time human–robot interaction should be allowed, e.g., a technician operates on the factory floor while mobile robots safely move around the technician. Another aspect considers enhancing real-time human–robot interactions, e.g., a technician demonstrates an arm motion to a robot by performing contactless teaching, so called lead-through programming. In both cases, robot motion planning software is deployed at the edge cloud and is fed by a video camera system, providing the robot with vision capabilities.

Offloading robot software and advanced industrial robotics applications from its dedicated hardware to a common edge cloud platform puts stringent requirements on the reliability and latency of the 5G communication network. The requirements on bounded communication latency and communication service availability of the use cases at the Kista trial site are listed in Table 1. Further details of the use cases can be found in [8,9].

Use Cases	Message Size (Bytes)	Cycle Time and Latency Bound (ms)	Comm. Service Availability (# of Nines)
Robotics motion planning (Kista)	500	5–50	99.99%
Controller-to-controller (Reutlingen)	500	4–10	99.9–99.999%
Cloud-based mobile robotics (Reutlingen)	500	10–100	99.999%
Workpiece monitoring (Aachen)	1024	10	99.999%

Table 1. Industrial use cases and their requirements.

#### 2.2. Aachen Trial Site

The Aachen trial site is part of the 5G-Industry Campus Europe. It is located on the shopfloor of Fraunhofer IPT, which is equipped with multiple different machine tools. The infrastructure reflects a production landscape that is representative of many manufacturing companies, ranging from small and medium enterprises to large enterprises. Trial activities are focused on using 5G non-public networks for monitoring critical manufacturing processes, contributing to the paradigm of networked, adaptive production. One of the use cases is the monitoring of the tool condition in cutting processes such as milling using a 5G wireless acoustic emission (AE) sensor system. These kinds of systems are widely applied today to detect critical incidents such as tool breakage, collisions, and even tool wear, helping companies to better use their tools throughout the tool's lifetime. A wireless AE system can be applied closer to the tool without the need to feed cables through the complete machine, making it ideally suited for easy retrofitting. The AE sensor system uses a probe that is attached close to the machining area, which is monitored by the sensor system at a 1 MHz sampling rate. The wireless sensor system also contains an FPGA-based platform for signal processing. The measurement data are transmitted over 5G at 8 Mbit/s. Low latency with high reliability provided by 5G is crucial since short reaction times upon critical incidents are needed. The requirements for the bounded communication latency and communication service availability for the workpiece-monitoring use case are listed in Table 1. More details of the use case can be found in [8,10].

### 2.3. Reutlingen Trial Site

There are two different 5G use-cases trialed at the Bosch semiconductor factory in Reutlingen, see Figure 1, as part of the 5G-SMART project. The cloud-based mobile robotics

use case focuses on investigating the feasibility, flexibility, and performance of wirelessly controlled Automated Guided Vehicles (AGVs) that are connected to the 5G network. Decoupling the closed-loop control of the AGV from its embedded system and placing it into a factory cloud enables the physical platform and control intelligence to evolve separately, which contributes to the simpler hardware architecture of the robot device and longer operation time with a single battery charge. Moreover, it creates several operational benefits such as enabling enhanced route selection in real-time through a common map. This use case employs two types of AGVs: a commercial AGV and a prototype AGV realizing partial and full cloud-based control, respectively. In this context, the timely exchange of control messages between the AGVs and their control stacks over 5G becomes of great importance for the desired operation of the AGVs.



**Figure 1.** High-level architecture of the use cases realized in Reutlingen, where the arrows indicate the traffic flow between the PLCs in the TSN/Industrial LAN over 5G use case and the AGV and its control application in the cloud-based mobile robotics use case. Source: Bosch.

The TSN/Industrial LAN over 5G use case, on the other hand, centers around investigating and validating the applicability of 5G for transporting the traffic for Time-Sensitive Networking (TSN)/industrial LAN (I-LAN) applications. Limited flexibility for setting up new production lines or for restructuring an existing production line as well as complex and costly maintenance are the major drawbacks of wired I-LAN realizations. This use case aims at evaluating the feasibility of partially replacing fixed interconnections between TSN/I-LAN nodes with 5G mobile communications. This will reduce wear and tear of the cables and connectors for the mobile machines/controllers, resulting in reduced maintenance costs. Furthermore, replacing the cables for communication between controllers and machines with 5G also means greater flexibility for the implementation and adaptation of the industrial manufacturing infrastructure, which can improve manufacturing productivity. The requirements of the use cases are shown in Table 1. Further details describing these use cases and their implementation status can be found in [8,11].

# 3. Trial Systems in the 5G-SMART Project

The 5G standard specifies non-public networks (NPN) that can be deployed at enterprise locations to support industrial IoT services for industry users [12–15]. In the 5G-SMART project, we carried out on-premises trial deployments of 5G NPNs in factory production halls at different industrial sites. In this article, we present the over-the-air (OTA) evaluation results from these industrial deployments. Our performance analysis focuses on the industrial use-case scenarios described in Section 2; the concrete implementation of the different use cases is ongoing work. Compared to a recently available measurement study on a campus network [16], our work focuses on both the mid-band (MB) and high-band (HB) performance of 5G networks with ultra-reliable low-latency communication (URLLC) features in real industrial deployment scenarios across different trial sites. Moreover, we focus on high percentile values in our latency analysis, which is of key importance for mission-critical industrial IoT applications as opposed to average performance metrics.

In total, there are six different 5G network deployments within the 5G-SMART project. These are briefly described below and are summarized in Table 2. The architectures and frequency bands for the different sites were selected to enable the evaluation of different deployment options.

Trial Site	Trial System	5G Architecture	5G Frequency Band	4G Frequency Band	5G Bandwidth [MHz]
Aachen	MB NSA	NSA	3.6 GHz (n78)	2300 MHz (B40)	100
Aachen	MB SA	SA	3.6 GHz (n78)	NA	100
Aachen	URLLC	SA	28 GHz (n261)	NA	100
Kista	HB NSA	NSA	28 GHz (n257)	1800 MHz (B3)	200
Västerås	MB NSA	NSA	3.6 GHz (n78)	1800 MHz (B3)	80
Reutlingen	MB SA	SA	3.6 GHz (n78)	NA	100

Table 2. 5G-SMART trial system summary.

Aachen: As part of the 5G Industry Campus Europe [16] in Aachen, Germany, a 5G non-standalone (NSA) system has been deployed on a shopfloor with an area of 3000 m<sup>2</sup> at the Fraunhofer Institute for Production Technology. In parallel to the 5G NSA system, a 5G Standalone (SA) system has also been built to investigate the integration of the use cases with both the 5G NSA and SA systems. The 5G RAN uses 100 MHz of bandwidth in the locally licensed mid-band Time-Division Duplex (TDD) spectrum at the 3.7–3.8 GHz spectrum (5G band n78). This trial system is based on commercial 5G network components and can be used with commercially available 5G terminal devices. In addition, a 5G prototype system operating in the high-band TDD spectrum at 28 GHz (5G band n261) has been installed in the shopfloor and includes pre-commercial and standard compliant URLLC functionality with a standalone core network. It targets some of the more demanding and future requirements of industrial IoT use cases, hereafter referred to as the URLLC testbed, and is depicted in Figure 2. The prototype system is not based on the commercial hardware and software but contains all of the relevant parts of the 5G protocol stack. Further details describing the 5G trial systems can be found in [10].

Kista: At the 5G-SMART trial site in Kista, Sweden, a 5G NSA system operating at 28 GHz (5G band n257) using 200 MHz of bandwidth is deployed within an Ericsson factory. In this trial system, the 5G-connected industrial robot use cases are investigated and demonstrated. The industrial robots are provided by ABB, and the use cases have been developed at an ABB facility in Västerås, where a 5G NSA system operating on a test license at 3.7 GHz (5G band n78) has also been deployed. More details describing the 5G trial systems can be found in [9].

Reutlingen: Another 5G-SMART trial site is in Reutlingen, Germany, where a 5G SA system has been deployed in the clean room of a Bosch semiconductor production facility. The 5G system operates in the locally licensed spectrum at 3.7–3.8 GHz (5G band n78) using a bandwidth of 100 MHz. Further details on the 5G deployment can be found in [11].

We have carried out signal strength measurements in different factory halls, and all the deployment setups ensure excellent coverage in the factory halls even behind machine chambers and cabinets.

We have conducted extensive testing on the 5G trial networks in Aachen to validate the performance of the 5G networks. Tests on the trial systems in Kista and Reutlingen are ongoing and will intensify as the industrial equipment for the use cases has been installed.



**Figure 2.** Selected parts of the URLLC testbed. The middle rack includes the RAN and core network with a phased-array antenna module on top. On the far right, the UE prototype is visible also with an antenna mounted on top.

# 4. Evaluation Setup

# 4.1. Measurement Setup

In order to carry out latency and reliability measurements systematically in realistic application scenarios, we have developed an FPGA-based tool that allows the generation of configurable industrial automation protocol traffic. Having the possibility to configure the periodicity and data sizes of the individual messages of the selected protocol for traffic streams using the FPGA-based tool, various application scenarios using Layer-2 (e.g., PROFINET) and Layer-3 (e.g., UDP over IP) have been evaluated. The FPGA-based tool not only allows the emulation of user-configurable application traffic in a highly precise manner but also allows latency measurements to be conducted in the FPGA hardware with sub-microsecond-level accuracy. The measurement setup is illustrated in Figure 3.



**Figure 3.** Measurement setup for on-premises 5G network on the shopfloor, where an FPGA-based tool is used for industrial traffic emulation and latency measurements. Source: Ericsson.

As industrial automation applications do not necessarily have traffic with specific round-trip behavior from the 5G user equipment (UE) to the edge server and back, simple round-trip timing (RTT) measurements (such as ICMP ping measurements) have not been considered. Instead, separate uplink (UL) and downlink (DL) one-way measurements of the user data traffic have been used in our measurement campaign in order to observe the end-to-end (E2E) latency behavior of a 5G network for edge applications. In Figure 3, a bidirectional arrow is used to indicate the traffic flow from the 5G device to the base station radio for UL, while the opposite indicates the DL direction. The setup allows the UL and DL directions to be tested separately.

As the support for native Layer-2 communication over 5G is not yet typical in commercial 5G devices and networks, we applied Layer-2 data tunneling over IP-based tunneling on Generic Routing Encapsulation (GRE) for the tests involving PROFINET traffic. The direct transmission of Ethernet frames over 5G has only been carried out in the URLLC testbed.

In one set of experiments, we empirically evaluated the benefits of using FRER, which is part of the TSN standard suite. FRER was designed for fixed Ethernet networks to provide, e.g., resilience against equipment failures. In the FRER scheme, the application traffic with specific VLAN tags is replicated over two separate transmission paths followed by potential elimination at the receiving end. While the replication of the application traffic on the redundant paths consumes more resources, FRER provides resilience against transmission failures on one of the paths. FRER can also be applied to wireless networks such as 5G networks. It can be used to establish redundant network paths via different independent network nodes [13] and can provide protection against, e.g., node failures. When the same 5G network nodes are used, FRER can still protect against, e.g., UE failure or performance variations. In addition, FRER can improve the achievable reliability for meeting certain latency bounds whenever wireless links have larger latency variations since the end application can utilize the earliest data received via the separate FRER paths. A simplified setup that can be used for FRER tests over the 5G system is illustrated in Figure 4, where the application data (with specific IEEE 802.1Q VLAN tags to separate from background traffic) are transmitted over two user-plane functions using two UEs. Please note that for simplicity, Figure 4 does not show the Layer-2 tunnel setup or the traffic flow directions.



**Figure 4.** A simplified FRER setup where the application traffic at the ingress point is replicated and transmitted over two independent 5G user plane paths using two different UEs and is eliminated at the egress point.

We carried out our OTA measurements during 5G trial deployments on factory shopfloors where the UEs are placed at various location, including the inside of a milling machine chamber, mounted on robots and AGVs, behind cabinets, etc., and including both line-of-sight and non-line-of-sight scenarios between the 5G UE and the 5G base station. Compared to office and lab setups, factory shopfloors have strong multipath effects with lots of metal, moving machine parts, multiple automation cells and assembly lines, AGVs, and workers moving around, etc. We obtained a large sample size of at least 100,000 (and typically 1 million) measurements for each network, deployment, and traffic setting in order to generate statistically significant results. This amounts to an overall measurement time that spans several days. In this article, we present the results from a single 5G terminal device. We conducted the tests when the radio network was loaded with additional traffic

that had been created by up to three background UEs. The load had little influence on the performance results. Systematic performance measurements investigating the load situations of multiple devices and traffic types have been left for future work.

# 4.2. 5G Configuration Options

With the ongoing global deployment of commercial 5G networks, the focus of the deployed devices and infrastructure is currently on the public consumer markets for mobile broadband applications and fixed-wireless access. Those 5G systems also are currently the basis for deployments addressing vertical use cases and are based on products that are primarily targeted toward mobile broadband (MBB) traffic with limited customization for industrial use cases. In our measurements, we used different commercially available devices, all of which are based on the modem chipsets that have been primarily designed for the MBB market. The performance results with the default network settings indicate that such devices suit several industrial applications, as described in the subsequent section. This configuration is referred to as 5G-MBB, and the default configuration used for the evaluations is shown in Section 5. Obviously, the achievable performance is dependent on traffic characteristics, such as the message size as well as the service requirements, e.g., the required reliability level for data communication.

We have experimented with different RAN configuration parameters in the trial networks related to scheduling and link robustness and observed that parameter tuning could allow significantly reduced latency to be achieved. In particular, latency reduction for uplink transmission is highly desirable, especially when considering that most of the industrial use cases have uplink-focused traffic profiles. RAN configurations, resulting in reduced uplink latency, are referred to as 5G-MBB, which have an adapted configuration in the later description of the results. We also observed that the achievable performance varies substantially when using different 5G devices, even when they are based on the same 5G modem chipset.

Redundant data transmission with FRER creates multiple independent traffic flows for the same data application. In cases where the flows traverse independent paths, this can be used to protect against equipment failure, which is beyond the scope of this article. When FRER is applied over a 5G network, another benefit can be identified: FRER may be used to create multiple flows, even over a single path, providing improved reliability performance. When different wireless devices are locations with similar radio propagation and that are transmitting the same data over the same network with the same network load and that are located nearby, then the average performance perceived by those devices is mostly similar. However, in occasional cases, one of the devices may face wireless transmission errors caused by, e.g., a blockage; such transmission errors are then recovered by retransmissions at the cost of an increase in transmission latency. Moreover, the scheduling algorithm may lead to slightly different latencies on redundant paths. In such cases, the redundant transmission of data flows over separate data paths and 5G devices can provide performance benefits, as the better-performing data path determines the end-to-end performance. Related measurements are presented in Section 5.

Finally, we empirically study the achievable performance of the pre-commercial standard compliant URLLC testbed. This testbed has been designed and optimized for URLLC and includes features such as scheduling enhancements, robust control and data channels, traffic prioritization schemes, and latency-optimized TDD patterns, some of which have been explicitly specified to support time-critical communication use cases, such as those in industrial automation. The measurements determined for the URLLC testbed are referred to as 5G-URLLC in the figures in Section 5.

# 5. 5G Performance for Industrial Control

In our latency measurements of industrial traffic over a 5G network, we assumed fixed message sizes with periodic transmission intervals. Periodic messages with an inter-arrival time of 8, 16, and 32 ms have been evaluated as they would be in industrial control processes

using protocols such as PROFINET, CIP, etc. We obtained at least 100,000 samples for each measurement in order for our measurement results to have sufficient statistical significance. The standard deviation for each percentile value shown for a given packet size in the UL or DL direction with different inter-packet periodicity intervals generally varied from 0.3 ms to 1.2 ms in the results presented in this section. Different message sizes were considered: 100 bytes (as shown in Figures 5 and 6) and 1024 bytes (as shown in Figures 7 and 8). A first observation is that for the typical small message sizes for industrial control [1], the MBB-focused 5G system as deployed today can achieve low average latencies: 3.7–6.9 ms in uplink and 5.2–6.2 ms in downlink. This is visible from the median latency in Figures 5–8 When the message size is increased by a factor of approximately 10, the latency increase is rather small.



Uplink Latency [ms] - 100 bytes

**Figure 5.** Uplink latency for 100-byte messages and different reliability percentiles (median, 99th percentile, 99.9th percentile).



Downlink Latency [ms] - 100 bytes

**Figure 6.** Downlink latency for 100-byte messages and different reliability percentiles (median, 99th percentile, 99.9th percentile).











**Figure 8.** Downlink latency for 1024-bytes messages and different reliability percentiles (median, 99th percentile, 99.9th percentile).

What becomes apparent is that the latency increases significantly when considering the 99th percentile or even the 99.9th percentile values. In other words, a 5G system focused on MBB services has very good average performance, but this performance cannot be necessarily guaranteed with high reliability.

For uplink, the latency that can be provided for 99% of the transmissions increases from the median value by 1.05–4.1 ms when small messages that are 100 bytes in size are transmitted. For large messages of 1024 bytes in size, the corresponding latency increase is between 3–10.6 ms. The equivalent latency value increase in the downlink is 4.7–4.8 ms for 100 bytes messages and 5.5–5.8 ms for large 1024 bytes messages. The latency guarantees at higher levels of 99.9% lead to a latency increase from the median uplink value of 1.5–9.4 ms for small messages and 6.5–12.8 ms for large messages. In the downlink, the corresponding increases in latency are 5.8–8.8 ms for small messages and 6.1–9.3 ms for large messages.

Based on those results, it can be concluded that MBB-based 5G systems are primarily designed and configured to provide good average latency performance while providing stringent guarantees on the latency bounds has not been the focus.

We investigated two different 5G devices in the measurement setup, which are shown in the green and blue graphs in Figures 5–8 and are marked as device A and device B. It is noteworthy that the latency performance can differ substantially between devices, even if both 5G devices use the same 5G modem chipset (i.e., Qualcomm x55). For example, in the uplink, Device B has a 2.9 ms lower median latency and a 4.6 ms higher latency for the 99.9th percentile. In general, the differences in the latencies between the two devices are in the range of up to 6 ms.

The 5G network can be configured in various ways. In particular, during uplink, the link can be configured with higher link robustness, while medium access times can be reduced by pre-scheduling. Figures 5 and 7 show the impact of the RAN configuration on the latency performance. It can be observed that the gains for the selection of the RAN parameter values is particularly visible for large message sizes (see Figure 7), where the achievable median latency is reduced from 12.5 ms to 4.5 ms. For small message sizes, the reduction (see Figure 5) in the latency is only up to 1.5 ms, and for device A, the 99.9th percentile of the latency is even increased by 3.1 ms.

One way to increase the reliability of 5G connectivity is to introduce redundant communication paths, which is enabled by, for instance, FRER, as shown in Figure 4. The benefit of redundant connectivity is that if a single communication path is experiencing link degradations and potentially long transmission times, the other link can still provide message transmission with lower latency. The FRER performance is shown in Figures 9 and 10 for uplink and downlink, respectively. The figures show that, not surprisingly for the median latency, FRER does not provide any benefits, as its primary target is improved reliability. As the average latency of the 5G system performs in-line with the configuration, it is not expected that the second transmission path can provide any substantial benefit. However, at the 99th or 99.9th percentile values, FRER provides extra robustness. As a result, the guaranteed latency that can be provided to the industrial system can be reduced by up to several milliseconds.



Uplink Latency [ms]

Figure 9. Uplink latency with and without redundant transmissions using FRER.



Downlink Latency [ms]

Figure 10. Downlink latency with and without redundant transmissions using FRER.

1024 bytes - 5G-MBB

For 5G, many features have been standardized, and these features allow ultra-low latency and high reliability to be achieved in stringent industrial IoT use cases as needed. We have conducted experiments with the URLLC testbed system described in Section 3, in which several of the 3GPP Rel. 15 and Rel. 16 standardized URLLC features were implemented in the network and the device. The results in in Figures 5–8. show that with the proper design and usage of standardized 5G URLLC functionality, much lower latencies are able to be achieved and can also be guaranteed with high reliabilities in comparison to MBB-focused 5G systems. For small and large messages, the median latencies are below 0.8ms in uplink and between 0.8–0.9 ms in downlink. Even for guaranteed latencies at 99.9th percentile, the latencies remain below 1.09–1.12 ms. A FRER configuration with redundant traffic flows over the URLLC testbed was not considered, as it is not expected to reduce latency below the already low latency bounds. While the presented results on the commercial 5G eMBB system in this article focuses on the 5G-SMART use cases (cf. [11]), we carried out the latency performance evaluation in both the UL and DL directions on the URLLC testbed for a wide range of packet sizes ranging from 100 bytes to 1500 bytes. Our empirical results indicate that the URLLC features guarantee a latency bound of approximately 1 ms for both the UL and DL directions for these packet sizes up to 99.9th percentile value.

1024 bytes - 5G-MBB + FRER

# 6. Summary

We have presented 5G over-the-air performance measurements conducted in realistic industrial production environments with realistic deployment settings. Our empirical results provide insightful understanding of 5G system performance for some manufacturing use cases. We have applied traffic models to represent realistic industrial traffic as described in 5G-ACIA, 5G-SMART, and 3GPP. A proper realization of industrial use cases over 5G and its validation is ongoing work within the 5G-SMART project [9–11].

We have shown that besides the use of higher subcarrier spacing and more suitable TDD patterns, 5G systems with scheduling enhancements, prioritization, robust control and data channels, faster control feedback mechanisms, etc., can achieve the very low latencies that are required for some of demanding industrial use cases. However, for early 5G networks and devices that are designed for MBB services, we observed an increase in the latency for higher reliability guarantees. The achievable latency is further dependent on the 5G device as well as on the network configuration. We also observed that redundant transmissions, for instance, via FRER, can reduce the E2E latency when high reliability is

required. We evaluated 5G features that have been specified for ultra-reliable and lowlatency communication. We observed that with these features, very low latencies could be achieved with high reliability guarantees, which are required in some stringent industrial IoT applications.

**Author Contributions:** Conceptualization and validation, J.T., J.S., P.d.B., L.G., J.F., B.V. and J.A.; methodology, software, formal analysis, and research, J.A. and C.A.; writing, J.A., C.A., P.d.B., J.S., L.G., B.V., J.F., N.K., D.H. and R.H.S. All authors have read and agreed to the published version of the manuscript.

**Funding:** This work was performed in the framework of the H2020 project 5G-SMART co-funded by the EU under grant number 857008. 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.

**Acknowledgments:** The authors would like to acknowledge the contributions of their colleagues from ABB, Bosch, Ericsson, and Fraunhofer IPT.

Conflicts of Interest: The authors declare no conflict of interest.

#### References

- 3GPP. Service Requirements for Cyber-Physical Control Applications in Vertical Domains. Technical Specification TS 22.104 V18.1.0. June 2021. Available online: https://www.etsi.org/deliver/etsi\_ts/122100\_122199/122104/16.05.00\_60/ts\_122104v160 500p.pdf (accessed on 19 October 2021).
- 5G-ACIA. Key 5G Use Cases and Requirements. White Paper. Available online: https://5g-acia.org/wp-content/uploads/2021 /04/Key\_5G\_Use\_Cases\_and\_Requirements\_DOWNLOAD.pdf (accessed on 19 October 2021).
- Rischke, J.; Sossalla, P.; Itting, S.; Fitzek, F.H.P.; Reisslein, M. 5G Campus Networks: A First Measurement Study. *IEEE Access* 2021, 9, 121786–121803. [CrossRef]
- 4. Xu, D.; Zhou, A.; Zhang, X.; Wang, G.; Liu, X.; An, C.; Shi, Y.; Liu, L.; Ma, H. Understanding Operational 5G: A First Measurement Study on Its Coverage, Performance and Energy Consumption. In Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication, Virtual Event, 10–14 August 2020.
- Ficzere, D.; Soós, G.; Varga, P.; Szalay, Z. Real-Life V2X Measurement Results for 5G NSA Performance on a High-Speed Motorway. In Proceedings of the IEEE International Symposium on Integrated Network Management, Bordeaux, France, 17–21 May 2021.
- Hoppari, M.; Uitto, M.; Mäkelä, J.; Harjula, I.; Rantala, S. Performance of the 5th Generation Indoor Wireless Technologies-Empirical Study. *Future Internet* 2021, 13, 180. [CrossRef]
- 5G-ACIA. 5G for Connected Industries and Automation Second Edition. White Paper. Available online: https://5g-acia. org/wp-content/uploads/2021/04/WP\_5G\_for\_Connected\_Industries\_and\_Automation\_Download\_19.03.19.pdf (accessed on 19 October 2021).
- 5G-SMART. Forward Looking Smart Manufacturing Use Cases, Requirements and KPI's. Deliverable D1.1. Available online: https://5gsmart.eu/wp-content/uploads/5G-SMART-D1.1.pdf (accessed on 19 October 2021).
- 5G-SMART. Design of 5G-Based Testbed for Industrial Robotics. Deliverable D2.1. Available online: https://5gsmart.eu/wpcontent/uploads/5G-SMART-D2.1.pdf (accessed on 19 October 2021).
- 5G-SMART. Report on Implementation of Options for Monitoring of Workpiece and Machines. Deliverable D3.3. Available online: https://5gsmart.eu/wp-content/uploads/5G-SMART-D3.3-v1.0.pdf (accessed on 19 October 2021).
- 11. 5G-SMART. Report on Development of the 5G Use Cases. Deliverable D4.3. Available online: https://5gsmart.eu/wp-content/uploads/5G-SMART-D4.3-v1.0.pdf (accessed on 19 October 2021).
- 12. Sachs, J.; Landernäs, K. Review of 5G capabilities for smart manufacturing. In Proceedings of the International Conference on Wireless Communication Systems (ISWCS), Berlin, Germany, 31 August–2 September 2021.
- 13. 5G-SMART. First Report on 5G Network Architecture Options and Assessments. Deliverable D5.2. Available online: https://5gsmart.eu/wp-content/uploads/5G-SMART-D5.2-v1.0.pdf (accessed on 19 October 2021).
- 14. 5G-ACIA. 5G Non-Public Networks for Industrial Scenarios. White Paper. Available online: https://5g-acia.org/wp-content/uploads/2021/04/WP\_5G\_NPN\_2019\_01.pdf (accessed on 19 October 2021).
- 15. 3GPP. System Architecture for the 5G System (5GS). Technical Specification TS 23.501 V17.1.1. Available online: https://www. 3gpp.org/ftp/Specs/archive/23\_series/23.501/23501-h20.zip (accessed on 19 October 2021).
- 16. 5G Industry Campus Europe. Available online: https://5g-industry-campus.com/ (accessed on 19 October 2021).