

Assessing the Controllability of Modular UAV Swarms over 5G Networks via ROS 2

Kevin Tebbe¹, Robert Pommeranz¹, Gaurav Kapoor¹, Ralf Heynicke¹, Gerd Scholl¹

¹Chair of Electrical Measurement Engineering, Helmut-Schmidt-University, University of German Armed Forces, Hamburg

Abstract

Multi-UAV operations are constrained by legacy point-to-point radio links that scale poorly in range, bandwidth management, and quality of service. In this feasibility study, the controllability and reliability of a highly modular Robot Operating System 2 based UAV swarm interconnected via private 5G cellular networks are evaluated. After integrating multiple aerial vehicles into a 5G architecture, the communication performance is compared against standard WLAN by profiling end-to-end latency and throughput under representative mission and sensor traffic. The findings provide practical design guidelines for deploying UAV swarms over cellular infrastructure, particularly in Beyond-Line-of-Sight scenarios, where scalability and coverage outweigh raw latency performance.

Keywords: C-UAS, 5G, ROS 2, Swarm, PX4-Autopilot

Introduction

Unmanned Aerial Vehicles (UAVs) are increasingly deployed in complex, distributed scenarios such as disaster response, infrastructure monitoring, and cooperative multi-agent missions. Traditional UAV communication relies predominantly on point-to-point radio links, such as ExpressLRS, FrSky R9 or proprietary Links from DJI or Walksnail. While these systems provide low-latency links for individual platforms, they do not scale efficiently with increasing swarm size, nor do they support flexible multi-node communication patterns required for collaborative autonomy. As swarm sizes grow, the communication paradigm must evolve toward point-to-multipoint (P2X) or fully networked architectures. Cellular communication systems, particularly Fifth Generation (5G) Standalone (SA) networks, offer a promising alternative due to their inherent support for wide-area coverage, mobility management, and quality-of-service (QoS) differentiation. However, unlike local wireless technologies such as Wireless Local Area Networks (WLAN), 5G networks introduce additional architectural complexity, including radio scheduling, core network traversal, and transport-layer buffering. As a result, the communication network becomes an active component of the control loop rather than a transparent channel. This shift raises a fundamental question: to what extent do the latency and reliability characteristics of cellular networks constrain the controllability of UAV

swarms? To address this question, this paper presents an experimental evaluation of a modular UAV swarm architecture based on PX4-Autopilot and Robot Operating System 2 (ROS2). Unlike prior work focusing primarily on network-level performance metrics, this study explicitly links measured communication characteristics to control-relevant constraints such as achievable control loop rates, watchdog design, and synchronization accuracy. In doing so, the paper frames the communication network as a dynamic system component that directly influences swarm behavior and operational safety. The results provide a quantitative comparison of communication technologies and practical design guidelines for deploying UAV swarms over wide-area cellular infrastructure, particularly in Beyond-Line-of-Sight (BLoS) scenarios.

Related Work

Recent research on multi-UAV systems has increasingly focused on modular software and hardware architectures that support scalable swarm operation. In particular, middleware-based approaches built on PX4-Autopilot and ROS2 have gained attention because they enable flexible integration of heterogeneous vehicles, companion computers, and distributed mission logic. A modular and scalable ROS2/PX4-based architecture for heterogeneous UAV swarms was presented in [1], providing the system-level foundation for the platform used in this paper. In parallel, the role

of 5G communication for cyber-physical and industrial real-time systems has been investigated in a growing body of literature. Prior work has shown that private 5G SA networks can provide predictable connectivity and offer advantages such as managed spectrum access, mobility support, and QoS differentiation. The private 5G SA campus network used in this work, including its baseline performance characteristics, is described in [2]. However, existing studies often focus either on communication infrastructure metrics without explicitly linking network behavior to control-relevant constraints in real multi-UAV deployments. In particular, there is still limited experimental work on how end-to-end latency, jitter, and packet loss affect the practical controllability of ROS2-based UAV swarms operating over cellular 5G networks. This paper addresses the gap by experimentally evaluating a modular PX4/ROS2 UAV swarm by interpreting the measured application-level latency and reliability in terms of operational swarm control constraints.

System Architecture and Experimental Setup

The UAVs in this experiment are configured as modular research platforms. They consist of a Pixhawk 6X Flight Control Unit (FCU) running PX4-Autopilot version 1.15, an Nvidia Jetson Orin NX as a companion computer for high-level tasks, and a 5G modem (Quectel RMU500EK or Waveshare Dongle with RM520N-GL). Figure 1 illustrates the hardware used in this experimental setup, whereas Figure 2 depicts the logical network topology, including the protocols and communication pathways between the UAVs and the ground station. The logical connections within the 5G network are omitted from the diagram for clarity. The experimental setup involves interconnecting the Ground Control Station (GCS), designated as the benchmark node, and the UAVs via three different physical layers for comparison: LAN, WLAN, and a private 5G campus network at Helmut-Schmidt-University (HSU). Figure 3 illustrates the communication pathways and protocols between the vehicles and the GCS. This modular architecture enables a clear separation between flight control, onboard processing, and communication layers, allowing the isolated evaluation of network-induced effects on swarm operation.

Measurement Scenario

The benchmark system implements a distributed ROS 2 measurement campaign for one-way message latency measurements in multi-UAV scenarios. The system is structured

as a client-server topology, where a GCS serves as the time synchronization master and the sender nodes (UAVs) that are time-synchronized to the GCS via Network Time Protocol (NTP).



Fig. 1 UAV experimental setup showing Nvidia Jetson, Pixhawk 6X, and Quectel RMU500EK 5G modem mounted on a Holybro X650 Frame.




Fig.2 Simplified representation of the benchmark measurement setup for KPI measurement using LAN, WLAN and 5G cellular network.

In this work, latency in this publication is defined as the time difference between message generation at the UAV and reception at the GCS:

$$L_{\text{end-to-end}} = t_{\text{recv}} - t_{\text{send}}, \quad (1)$$

whereas t_{recv} is the timestamp of message reception at the GCS (ROS 2 system clock) and t_{send} is the timestamp of message publication at the UAV (ROS 2 system clock). Clock synchronization across the UAV and GCS is maintained via NTP. Although Precision Time Protocol (PTP) is generally more accurate than NTP, the achieved NTP synchronization accuracy of 1 ms – 2 ms is sufficient for the

latency measurements in this publication, which are in the order of tens to hundreds of ms [3], [4].




Fig. 3 Benchmark setup for ROS 2 based latency measurements in multi-UAV scenarios. The figure illustrates the generic data flow from a UAV to GCS.

The UAV messages from the PX4 firmware are not using sequence numbers, which makes it difficult to determine the number of lost packets. Therefore, packet loss is determined by analyzing the dedicated latency probe stream on a dedicated latency probe topic using Zenoh (a lightweight communication middleware, detailed in Sec. Network Benchmarking), which contains a sequence number in each JSON payload. By comparing consecutive sequence numbers in the probe, missing probes can be identified and packet loss calculated from the gap count.

Message-Callback Architecture:

The measurement framework follows a dual-path design to capture both end-to-end message delay and control-plane dynamics of the Zenoh edge-communication protocol [5]. The execution flow of each node is determined by command-line initialization arguments passed to a unified measurement script. When launched in GCS mode, the orchestration layer starts a periodic Zenoh dump process and latency subscribers in parallel. When launched in UAV mode, only the probe publisher is active. **Active Probe Path:** A ROS 2 probe publisher transmits JSON payloads containing sequence number and sender timestamp. The subscriber callback extracts the payload, computes the one-way latency, and appends samples to database. For sample i , latency is computed as defined in Eq. 1. with acceptance constraint.

$$0 \leq L_i^{\text{ms}} \leq 6 \times 10^5 [\text{ms}]. \quad (3)$$

Let L denote the sorted latency sample set and let $a \in (0,1)$ represent the target probability level for the desired quantile (e.g., $a = 0.95$ for the 95th percentile). The reported quantiles are then defined as:

$$Q_q = a_{[q(N-1)]}, \quad q \in \{0.50, 0.95, 0.99\}. \quad (4)$$

Snapshot Key Performance Indicator (KPI) Path:

While the primary objective of this study is the evaluation of the end-to-end latency to assess swarm controllability, network throughput is monitored as a secondary KPI. Throughput is recorded not to evaluate maximum capacity, but to verify that the network bandwidth is not exhausted, ensuring that the observed latencies are not artifacts of network saturation. To contextualize these primary physical-layer metrics, the system simultaneously monitors the internal state of the middleware. At fixed interval Δt , the system queries the Zenoh administrative REST API to store a JSON snapshot of the middleware's internal metrics. Runtime topics and endpoint metadata are parsed from each snapshot to compute per-interval KPIs. The evaluated metrics encompass total records, topic cardinality at runtime, the number of unique routers and participants, writers/reader counts, QoS class counts, and topic churn relative to previous interval. All measurements are performed under identical software and configuration settings to ensure comparability across different communication links. The use of an application-layer latency probe ensures that the measured values reflect end-to-end communication performance as experienced by the control application.

Throughput Estimation:

Two throughput estimates are logged per interval. Admin-snapshot throughput:

$$B_{\text{admin}} = \frac{\text{bytes}_{\text{dump}}}{\max(\Delta t, 10^{-6})}. \quad (7)$$

All metrics are appended as timestamped CSV rows, yielding a synchronized dataset for latency, topology, QoS, and bandwidth analysis.

UAV Network Performance Characterization

Robust and reliable communication link is essential for controlling UAVs, whether individually or in a swarm. The performance requirements vary depending on the level of autonomy [6]:

- **Manual Control:** The operator sends direct commands (stick input). This requires low latency to ensure responsive handling.
- **Automatic Control:** The UAV follows a pre-programmed flight plan or mission. Real-time operator input is minimal, but telemetry monitoring remains critical.
- **Autonomous Control:** The vehicle makes decisions based on onboard sensors and algorithms. Communication is primarily used for high-level coordination and state updates, but low latency is still important for timely decision making and responsiveness.

Latency requirements depend on the environment and the flight type (Line-of-Sight (LoS) vs. Beyond-Line-of-Sight (BLoS)). For manual control at speeds up to 60 km/h, a maximum latency of 40 ms is recommended by 3GPP [7] as the upper limit for safe operation. Studies indicate that while latencies below 170 ms are manageable, delays between 300 ms-500 ms can significantly degrade operator performance and necessitate active obstacle avoidance systems [8]. For video streaming, latency requirements range from 20 ms-200 ms depending on the application and video quality [7].

Network Benchmarking

This section details the measurement campaign used to evaluate the controllability of modular UAV swarms over various physical networks (Ethernet, WLAN, and 5G). This triad was chosen to systematically isolate the latency bottleneck of the system: Ethernet establishes the hardware and middleware baseline devoid of wireless interference, WLAN serves as a high performance local wireless reference using Layer-2 switching, and private 5G reflects the scalable wide-area cellular communication required for operational BLoS deployments. It should be noted that this comparison highlights architectural differences between local and wide-area communication systems rather than purely physical-layer performance. The measurements were conducted using a Private 5G SA Campus Network at HSU. The network operates in the 3.7 GHz band (n78) with a 100 MHz bandwidth, utilizing Ericsson 5G Core and Radio Access Network (RAN) components. For this study, an enhanced Mobile Broadband (eMBB) slice was used. The testbed follows a client-server topology where the GCS acts as the time synchronization master. Zenoh Router Role: In this benchmark setup, the GCS is configured as the central Zenoh router. It is important to note that for a production-ready, real-world deployment configuring Zenoh for

P2P communication would likely be the preferred approach to further minimize end-to-end latencies. However, for the purpose of this comparative study, a routed topology was chosen to ensure a consistent and reproducible baseline across all network types. While a P2P configuration could reduce absolute latency values, the routed setup was intentionally chosen to ensure reproducibility and comparability across all communication technologies.

- **Time Sync:** Synchronization between UAVs and the GCS is maintained via NTP (chrony [9]), achieving and accuracy of approximately 0.1 ms–2 ms, which is sufficient for latency measurement in the range of tens to hundreds of ms.
- **Data Flow and Middleware:** ROS2 communication is bridged to Zenoh using the *zenoh-bridge-ros2dds*, enabling efficient data transport over UDP-based communication links. It serves as an optimized alternative to the standard Data Distribution Service (DDS) in ROS 2, specially targeting performance improvements over constrained or wireless networks [5].
- **KPI Collection:** A dedicated latency probe publishes JavaScript Object Notation (JSON) payloads (sequence numbers and timestamps) at 20 Hz. One-way latency is calculated as shown in Eq. (2) based on timestamp differences between sender and receiver.

The 5G network is configured to provide different slices with different 5G QoS Indicator (5QI) values, which allows the evaluation of the impact of different QoS settings on the network performance and swarm control metrics. In this study, only eMBB slice was available, which is not optimized for ultra-low latency communication. The KPIs used for the evaluation in this publication were taken from [2] and shown in Tab. 1:

Tab. 1: Private 5G Network Performance (HSU) adapted from [2]

KPIs	Average
Round-Trip Time^a	20.4 ms ^b
Throughput (DL)^c	912 Mbps
Throughput (UL)^c	110 Mbps

^a IRTT [10]

^b: with 2 · 1.2 ms latency added from Ethernet

^ciPerf3

Measurement methodology and system architecture:

The testbed follows a two-role topology. On the aerial side, each node consists of a Pixhawk 6X and an NVIDIA Jetson Orin NX. On the ground

side, a dedicated ground station performs KPI collection. Communication is tunneled from ROS 2 through *zenoh-bridge-ros2dds* over Zenoh (UDP) across Ethernet, WLAN, and private 5G links. Before each run, clock synchronization is established using NTP (chrony). In GCS mode, the script acquires Zenoh middleware KPIs via the official Zenoh REST API and stores all outputs in a common measurement directory.

End-to-end latency:

Each UAV publisher emits *std_msgs/String* probe messages at rate f (default 20 Hz) on

$$\text{topic} = \text{TOPIC_BASE}/\text{uav_}\langle id \rangle. \quad (9)$$

The publisher period is defined as

$$T_{\text{pub}} = \frac{1}{\max(hz, 0.1)}, \quad (10)$$

where hz denotes the requested update frequency. Each payload contains a sequence number and sender timestamp t_s . To convert timestamp tuples (sec, nsec) to floating point seconds. A dedicated ROS 2 latency probe publishes probe messages on topic */zenoh_latency_probe*. Each probe payload is a JSON object containing a sequence number and sender timestamp:

$$\text{payload} = \{\text{seq}":n,\text{send_ns}":t_{\text{send}}\}. \quad (12)$$

Samples with $\text{latency}_{\text{ms}} < 0$ or $\text{latency}_{\text{ms}} > 600,000$ are discarded as anomalies (clock violations or extreme outliers). Every N_{report} samples (default $N_{\text{report}} = 100$), percentile KPIs are computed from sorted latency samples a :

$$p_q = a_{\lfloor q \cdot (n-1) \rfloor}, \quad q \in \{0.50, 0.95, 0.99\}, \quad (14)$$

corresponding to median (p_{50}), tail latency (p_{95}), and extreme tail latency (p_{99}), and logged to console for real-time monitoring.

Network topology snapshot KPIs:

Every T seconds (default $T = 5$ s), queries the Zenoh admin space through the official Zenoh REST API and parses the returned JSON records.

Network throughput:

Two complementary throughput estimators are logged per interval. Application-layer proxy from admin dump size:

$$\text{admin_bps} = \frac{B_{\text{dump}}}{\max(\Delta t, 10^{-6})}. \quad (17)$$

Wire-level proxy from UDP capture:

$$\text{udp_bps} = \frac{B_{\text{udp}}}{\max(T_{\text{sample}}, 10^{-6})}. \quad (18)$$

Both throughput metrics, alongside the raw latency timestamps and sequence numbers, are logged locally as comma-separated values

to form the experimental dataset. The offline post-analysis subsequently parses these raw logs to compute the statistical network KPIs. All measurements were conducted under comparable system configurations and identical software setups to ensure a fair comparison between communication technologies. The three latency series are summarized in Fig. 3 to 5. Ethernet is included as a cable-bound reference only and is not considered a deployable link for UAV swarm operation.




Fig. 4 Normalized Ethernet Latency histogram




Fig. 5 Normalized WLAN Latency histogram




Fig. 6 Normalized 5G Latency histogram

Its measured latency level is lowest (*median* ≈ 0.74 ms, $p_{95} \approx 1.34$ ms), which defines the benchmark floor for the setup. For the relevant wireless comparison, WiFi and 5G show clearly

separated latency ranges. WLAN is centered at lower one-way delay (*median* ≈ 1.31 ms, $p_{95} \approx 2.31$ ms), while 5G exhibits higher central tendency and broader spread (*median* ≈ 15.46 ms, $p_{95} \approx 24.58$ ms). The reported means and standard deviations are shown directly in each histogram. At this stage, this section remains descriptive and does not derive control performance implications. The benchmark and wireless runs had durations of approximately 60 min yielding 10^5 one-way latency samples, respectively. Sequence-gap-based packet-loss estimates over all UAV streams were 0.000% (Ethernet), 0.000% (WLAN), and 0.035% (5G). Notably, this minimal packet loss on the 5G link occurred despite the static setup and optimal Received Signal Strength Indicator (RSSI) / Reference Signal Received Quality (RSRQ) conditions. The minimal packet loss observed exclusively on the 5G link can be attributed to the inherent characteristics of the cellular radio environment combined with the UDP-based transport layer. Unlike the controlled Ethernet baseline and the short-range WLAN environment, 5G signals are subject to transient channel fading. Because the evaluated network slice was configured for eMBB rather than URLLC, the MAC-layer Hybrid Automatic Repeat Request (HARQ) mechanisms possess strict retransmission limits. Once these lower-layer limits are exceeded during brief signal degradation, the packet is discarded. Since Zenoh relies (in this setup) on UDP without transport-layer acknowledgements in this configuration, such discarded frames translate directly into application-level packet loss. Consequently, in dynamic flight scenarios involving handovers and BLoS conditions, this packet loss rate is expected to increase, further emphasizing the need for robust onboard fail-safe mechanisms. All measurements used synchronized one-way probe logging and identical GCS-side KPI extraction; thus, packet-loss values are reported as application-level estimates derived from missing sequence numbers.

Analysis

The latency histograms in Fig. 5 and Fig. 6 reveal a clear separation between the two wireless communication technologies. The 5G latency distribution is centered around a higher median and has a broader spread compared to WLAN. The median latency for 5G is approximately 15.46 ms, while for WLAN it is around 1.31 ms. This corresponds to an approximately 12-fold increase in mean latency for 5G compared to WLAN. This significant difference in latency can be attributed to the architectural characteristics of the respective

communication systems. The architectural divergence is visually summarized in Fig. 6, which contrasts the direct localized data pathways of a WLAN Access Point (AP) against the multi-stage transversal required within the 5G RAN and CN. In contrast to LAN or WLAN, 5G networks behave more like a Wide Area Network (WAN) with higher latency and jitter due to the additional processing and network traversal involved [11]. Furthermore, Time Division Duplex (TDD) patterns, especially using eMBB slices, increase the latency. This pattern (typically 4:1) means that the network is only available for uplink transmission of 20% of the time, which can significantly increase the latency for uplink data. In contrast WLAN uses a completely different Medium Access Control (MAC) and layer-2 switching mechanism, which can provide lower latency under low-load scenarios. Even though the measured latency for 5G is significantly higher than WLAN, it is still within the defined latency requirement for controlling a UAV with direct stick input at a speed of 16.7 ms (60 km/h), which was defined as 40 ms in [5]. This highlights that network-induced delay must be explicitly considered as part of the control loop design rather than treated as a negligible transport effect. Operating an autonomous UAV swarm communicating over wireless networks is limited by the worst-case latency due to the need for low latency and high reliability for control and coordination messages. The higher latency and jitter must be taken into account when designing the control algorithms and operations a UAV swarm is capable of. For example, a swarm that relies on real-time communication for collision avoidance or formation control must be constrained in terms of maximum speed and maneuverability to ensure safe operation. This becomes particularly critical for tightly coupled swarm behaviors such as formation control or synchronized task execution. Algorithms requiring inter-UAV temporal alignment should implement a jitter buffer or use GNSS-time-stamped command activation (execute at $t_s + \Delta$) rather than arrival-triggered execution [12]. Evaluating dedicated low-latency 5G slices (Ultra-Reliable Low-Latency Communication (URLLC)) rather than the eMBB slice used in these campaigns represents a direct next step for reducing the latency and tail variability [13]. Special NR slices with low-latency configurations could potentially reduce the latency, but in this measurement campaign, the only supported eMBB slice was used, which is not optimized for low latency.




Fig. 7 Simplified comparison of data streams over cellular 5G and WLAN

Tail Latency and Watchdog Design:

Beyond median values, the tail of the latency distribution is critical for safety-critical operation. In the 5G trials, while the median was 15 ms, late arrivals reached over 400 ms. Watchdog mechanisms must reliably distinguish between such delays and actual packet loss. Consequently, system design must be driven by worst-case latency rather than average performance. It is recommended to derive fail-safe thresholds from the 99th percentile ($p_{99} \approx 30\text{ms}$) with an additional safety margin. Crucially, critical failsafe behaviors (e.g., Return-to-Home) must reside onboard the UAV to avoid dependence on delayed GCS commands.

Control Loop Rate Constraints:

A median one-way delay of 15.46 ms implies a minimum round-trip time of 31 ms. This limits the effective GCS-side control loop to approximately 32 Hz. For predictable operation considering the p_{95} values, a schedulable rate of 20 Hz is more realistic. High-frequency safety-critical reactions must therefore be handled locally on the UAV.

Jitter-Aware Command Timing:

The 5G standard deviation of 5.24 ms introduces timing jitter that is non-negligible relative to the median latency. For time sensitive swarm maneuvers — synchronized arming sequences, coordinated trajectory transitions, or timed payload releases — this jitter introduces an inter-UAV temporal uncertainty of $\pm 5\text{ ms}$ per node.

Middleware QoS and Message Scheduling:

The measured 5G application-layer throughput is stable at 0.140 MB/s with a narrow distribution, indicating that bandwidth saturation is not the primary constraint in the observed scenario. Instead, middleware-layer scheduling becomes relevant. ROS 2 via Zenoh bridges provides configurable QoS profiles: reliable (with retransmit acknowledgment) and best-

effort. For safety-critical topics — emergency stop, mode transitions, collision-avoidance triggers — a reliable QoS profile with a retransmit timeout aligned to p_{99} is advisable to avoid silent message loss. Non-critical telemetry, such as periodic GPS state logging, should use best-effort QoS profiles to prevent head-of-line blocking effects that would inflate the latency of critical messages. This separation prevents head-of-line blocking and preserves latency for critical control messages. Furthermore, the Zenoh session and transport layer parameters (scouting intervals, transport unicast/multicast configuration, and ROS 2 domain ID isolation) should be tuned per deployment link: configurations optimized for low-latency WLAN may not perform equivalently over a cellular 5G link with different PDCP and RLC retransmit behavior. Evaluating dedicated low latency 5G slices (URLLC) rather than the eMBB slice used in this campaign represents a direct next step for reducing the observed median latency and tail variability.

Latency Optimization and Deployment Considerations:

The analysis of 5G performance (median latency $\approx 15.46\text{ ms}$) vs. WLAN ($\approx 1.31\text{ ms}$) shows a clear 12-fold difference. As discussed in the methodology, the use of a central router at the GCS adds a specific processing overhead. While moving to a Zenoh P2P configuration could reduce the absolute latency values, the fundamental bottlenecks of the 5G eMBB slice — specifically the TDD pattern would remain the dominant factor for the 5G lag. Therefore, the relative performance gap between WLAN (Layer-2 switching) and 5G (WAN-like traversal) is expected to be robust against middleware-level topology changes.

Conclusion on Network Selection:

Although this campaign reports lower latency for WLAN under local test conditions, private 5G remains necessary for operational UAV swarms once deployment extends beyond confined, interference-controlled environments. Compared to WLAN, private 5G provides managed spectrum access, mobility support with handover, network slicing, and enforceable QoS policies over larger service areas, which are essential for predictable command-and-control performance under multi-user load. Accordingly, WLAN should be treated as a high-performance short-range link, while private 5G represents the scalable wide-area control infrastructure. In addition, when swarm operation can leverage public cellular infrastructure, the mission-level communication range is theoretically unbounded (i.e., constrained by network coverage rather than hotspot footprint), which is fundamentally not

achievable with standalone WLAN architectures. Therefore, the use of private 5G is justified not because it always minimizes latency, but because it enables controlled, secure, and geographically scalable swarm operations with a practical pathway toward beyond-visual-line-of-sight deployment.

Conclusion

Although WLAN provided lower latency in local tests ($p_{95} \approx 2.31$ ms), Private 5G ($p_{95} \approx 24.58$ ms) is the superior choice for operational swarms due to its managed spectrum, mobility support (handover), and scalability over large areas. While WLAN serves as a high-performance short-range link, 5G enables geographically unbounded operations. The use of private 5G is justified not by minimizing latency in all cases, but by providing a secure, scalable, and professional infrastructure BLoS deployments.

Acknowledgement

This work is a part of the “Digital Sensor-2-Cloud Campus Platform” (DS2CCP) project, which is funded by the Federal Ministry of Defense under the dtcc.bw program. Project website:

<https://dtccbw.de/home/forschung/hsu/projekt-ds2ccp/projekt-ds2ccp>

References

- [1] R. Pommeranz, K. Tebbe, R. Heynicke, and G. Scholl, “A Modular and Scalable System Architecture for Heterogeneous UAV Swarms Using ROS 2 and PX4-Autopilot,” Oct. 31, 2025. doi: 10.48550/arXiv.2510.27327.
- [2] H. Beuster, K. Tebbe, T. R. Doebbert, and G. Scholl, “Measurements of the Safety Function Response Time on a Private 5G and IO-Link Wireless Testbed,” in *2024 IEEE 29th International Conference on Emerging Technologies and Factory Automation (ETFA)*, Sep. 2024, pp. 1–4. doi: 10.1109/ETFA61755.2024.10710762.
- [3] D. L. Mills, “Internet time synchronization: the network time protocol,” *IEEE Trans. Commun.*, vol. 39, no. 10, pp. 1482–1493, Oct. 1991, doi: 10.1109/26.103043.
- [4] *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*. doi: 10.1109/IEEESTD.2008.4579760.
- [5] *eclipse-zenoh/zenoh*. (Apr. 15, 2026). Rust. Eclipse zenoh. Accessed: Apr. 15, 2026. [Online]. Available: <https://github.com/eclipse-zenoh/zenoh>
- [6] Luftfahrt Bundesamt (LBA), “Operations Manual / OM Betriebshandbuch für den Betrieb von unbemannten Flugsystemen (UAS).” Jan. 28, 2026.
- [7] ETSI, “Unmanned Aerial System (UAS) support in 3GPP (3GPP TS 22.125 version 17.6.0 Release 17).” Apr. 2022. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/122100_122199/122125/17.06.00_60/ts_122125v170600p.pdf
- [8] T. Zhang, “Toward Automated Vehicle Teleoperation: Vision, Opportunities, and Challenges,” *IEEE Internet Things J.*, vol. 7, no. 12, pp. 11347–11354, Dec. 2020, doi: 10.1109/JIOT.2020.3028766.
- [9] “chrony – Introduction.” Accessed: Apr. 15, 2026. [Online]. Available: <https://chrony-project.org/>
- [10] P. Heist, *IRTT*. (May 15, 2021). [Online]. Available: <https://github.com/heistp/irtt>
- [11] M. Seliem, D. Pesch, U. Roedig, and C. J. Sreenan, “Comparative Analysis of 5G and Wi-Fi Integration with TSN for Industrial Applications,” in *2025 IEEE Wireless Communications and Networking Conference (WCNC)*, Milan, Italy: IEEE, Mar. 2025, pp. 1–6. doi: 10.1109/WCNC61545.2025.10978747.
- [12] R. Imai and R. Kubo, “Introducing jitter buffers in networked control systems with communication disturbance observer under time-varying communication delays,” in *IECON 2015 - 41st Annual Conference of the IEEE Industrial Electronics Society*, Yokohama: IEEE, Nov. 2015, pp. 002956–002961. doi: 10.1109/IECON.2015.7392552.
- [13] L. Guijarro, J.-R. Vidal, and V. Pla, “Competitive advantage of URLLC vs. eMBB for supporting timeliness-relevant services,” Apr. 26, 2024, *arXiv: arXiv:2402.03922*. doi: 10.48550/arXiv.2402.03922.