

Aerial Monitoring with Drones: Experimental Insights on Communication Challenges

(Invited Paper)

Micha Sende
Lakeside Labs GmbH
Klagenfurt, Austria
0000-0003-0589-4819

Christian Raffelsberger
Lakeside Labs GmbH
Klagenfurt, Austria

Christian Bettstetter
University of Klagenfurt
and *Lakeside Labs GmbH*
Klagenfurt, Austria

Abstract—A system comprising multiple unmanned aerial and ground vehicles is implemented and evaluated for autonomous monitoring of freight containers in a container terminal. Reliable control and operation require stable, high-throughput, and low-latency communication. We share experimental results and lessons learned, demonstrating how drones benefit from 5G and Wi-Fi mesh connectivity and act as relays to enhance the connectivity and latency of ground vehicles. Current technologies provide most required functionality but have limited scalability.

Index Terms—UAV networks, aerial connectivity, latency, relaying, hybrid robot systems

I. INTRODUCTION

Aerial inspection and monitoring of large, complex, and dynamic environments is increasingly performed by semi-autonomous cyberphysical systems consisting of multiple drones. Application domains span transportation, agriculture, defense, industrial, and logistics. One of the key building blocks is the wireless communication required for tasks like command and control, sensing, coordination, and offloading of computation tasks. This is why significant research and development on communication and networking of drones has been made in recent years [1], [2]. Two core requirements are short roundtrip latency for command and control as well as high uplink rates for delivery of payload data like videos. The standardization body 3GPP has addressed the integration of drones into cellular networks and incorporated a range of drone-related features into its specifications [3]. Besides 5G, also Wi-Fi is employed to enable short-range and license-free communications and to form ad hoc mesh networks between drones.

Our own research on networked multi-drone systems has, for many years, concentrated on system-level and experimental work [4]–[7], often conducted in collaboration with practitioners to identify relevant requirements and approaches for real-world use cases. This paper continues this stream of hands-on research. Specifically, we report on communication-related insights gained from experiments with an aerial monitoring system that we developed, implemented, and evaluated in a container terminal. The system employs multiple aerial drones and a ground rover communicating via both 5G cellular and

Wi-Fi mesh systems with a local ground control station (GCS) and a cloud server. The system was used for two use cases, where drones (i) use the network to coordinate autonomously to identify and localize freight containers and (ii) support a rover operating on the ground to improve network coverage.

By implementing, deploying, and testing complex systems in real-world conditions, we report our findings regarding the communication requirements, including:

- measurements of distributed service discovery and solutions to avoid bottlenecks,
- measurements of visual monitoring payload and discussion of its scalability, and
- measurements demonstrating how an aerial relay provides connectivity for a ground vehicle.

The paper is structured as follows: Section II describes the system components and their interaction, Section III presents measurement results. Section IV discusses the findings. Finally, Section V recaps.

II. SYSTEM DESCRIPTION

The system we developed and assessed consists of several drones for aerial inspection and a single rover operating on the ground. These mobile vehicles are supported by computing infrastructure, both on-site for command and control and remote for data processing. The system architecture is shown in Figure 1.

Both Wi-Fi and 5G are employed for wireless connectivity between the system entities. A Wi-Fi mesh network enables the exchange of control commands and telemetry data, facilitating low-latency local coordination without external infrastructure. At the same time, 5G is used to transmit payload data to remote cloud services for processing and analysis.

As drones, we use twinFOLD GEO hexacopters by the Austrian company Twins. They are controlled by an ArduPilot flight control unit (FCU) and a Raspberry Pi 4b onboard companion computer running Ubuntu with the robot operating system (ROS). This allows drones to operate autonomously while coordinating each other and allowing remote command and control. The drones use the Quectel RM500Q 5G modem to connect to the Internet. The Wi-Fi mesh is based on




Fig. 1: System architecture with drones, rovers, cloud services, and control station connected with mesh Wi-Fi and 5G

the Wallys DR4029 wireless access points providing Wi-Fi 5 (802.11ac) in the 5 GHz band. They run the OpenWrt firmware that implements the Meshmerize multipath routing protocol [8].

The GCS controls the mission by sending commands to all or individual drones. To do so, it receives and visualizes the MAVLink telemetry and ROS topics forwarded by the onboard companion over an IP-based network. It can be placed on-site or at a remote site connected via the Internet.

The ROS-based flight stack on board the companion computer sends high-level navigation commands over a local MAVLink connection to the FCU, which then performs the low-level control to maneuver the drone. Specifically, for aerial monitoring, we implement a distributed approximation of the multiple traveling salesmen problem. The drones coordinate using ROS topics (publish-subscribe) and services and actions (request-response).

III. SYSTEM EVALUATION

The system is evaluated in a container terminal, where containers are transferred between road vehicles and train wagons. Containers are typically temporarily stored in grouped stacks before being transported further.

A. Use Cases

The communication capabilities are evaluated in two use cases: aerial monitoring and aerial relaying. These use cases address different combinations of vehicles: aerial vehicles only, mixed teams of aerial and ground vehicles, and ground vehicles only.

The **aerial monitoring** use case is illustrated in Figure 2. It consists of three drones that coordinate in a distributed way to scan a set of container stacks. They are supported by an on-site GCS that controls the mission and monitors the status of the drones and by a cloud server that processes images and identifies and localizes containers. Once the GCS initiates the mission, each drone selects a container stack and performs a circular flight around it while capturing pictures in a periodic way. The drones upload the pictures compressed with JPEG over the 5G network to a cloud server using the secure




Fig. 2: Aerial view of a container stack as seen by a drone with detected container IDs in green




Fig. 3: Rover and drone operating between the container stacks

file transfer protocol. They constantly share their scanning progress to the GCS and other drones. Once a stack has been fully scanned, the drone proceeds to the next stack (or returns to its takeoff location once all stacks have been scanned).

The **aerial relaying** use case is illustrated in Figure 3. Its goal is to demonstrate the ability of drones to provide networking services to ground users. The focus is to offer reliable low-latency communication for real-time services such as remote operation. Specifically, using Wi-Fi mesh, we stream a video from an autonomously operating ground rover to a GCS. A drone follows the rover's path in the air to provide wireless relaying functionality, thus improving the quality of the video if required. In good conditions, the rover transmits the video stream directly to the GCS. However, whenever the radio signal experiences strong shadowing, the stream is transmitted in a multihop way via the drone.

We perform several experiments for both use cases and present selected results and insights.

B. Aerial Monitoring

During the **initialization phase**, each drone advertises the services it provides and queries the services it requires to establish a local network for direct, low-latency coordination. This generates a significant traffic volume in a short period. This traffic burst places high demands on the network and can even break the communication when causing queue overflows.




Fig. 4: Discovery overhead

To address this problem, we analyze the communication middleware of ROS, which is based on data distribution services (DDS) for distributed discovery, serialization, and transportation. We set up a testbed consisting of two drones and one GCS connected via a Gbit-Ethernet switch. We start the GCS software to control the mission, and on each drone we start the complete ROS flight stack consisting of 32 nodes, 62 topics, 210 services, and 5 actions. The resulting number of packets over time is shown in Figure 4.

There are different vendors providing middleware. Our comparison of the standard option, eProxima Fast DDS, with the alternative middleware, Eclipse Cyclone DDS, shows only minor differences. However, when we replace the inter-drone communication with the Eclipse Zenoh protocol, we experience significant improvements in a way that the number of packets decreasing by more than an order of magnitude and the discovery time reduces to fraction of its previous value. This is achieved by the design of Zenoh’s discovery mechanism and more efficient messages in comparison to DDS: Instead of advertising all publishers and subscribers, only the desired subscriptions are advertised and grouped by drones (as opposed to the individual several dozen ROS nodes per drone) for compression [9].

During the **monitoring mission**, the drones regularly take pictures and offload them to cloud services for post-processing using the on-board 5G modem. The associated traffic volume depends on the flight speed and the spacing between capture locations. Our experiments are with 2m and 5m spacing, resulting in 38 to 111 capture locations for one container stack, depending on the size of the stack (see Fig. 5). Flights with a maximum speed of 5 m/s lead to a capturing period of 3s to 4s with an average size of approximately 10 MB. The resulting payload data rate is shown in Figure 6. The average value is between 20 and 30 Mbit/s, depending on the waypoint spacing.

To generalize these experimental results and study the


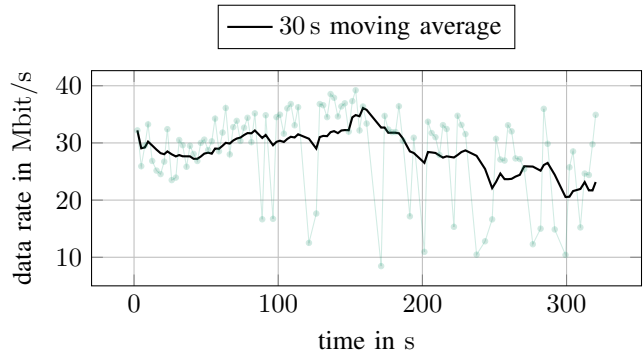



Fig. 5: Monitoring mission with three drones circling container stacks and taking pictures on marked photo waypoints (© Mapbox © OpenStreetMap)



a: 2 m waypoint spacing



b: 5 m waypoint spacing

Fig. 6: Measured data rate per drone for visual monitoring payload for different flights (colored) and average (black)

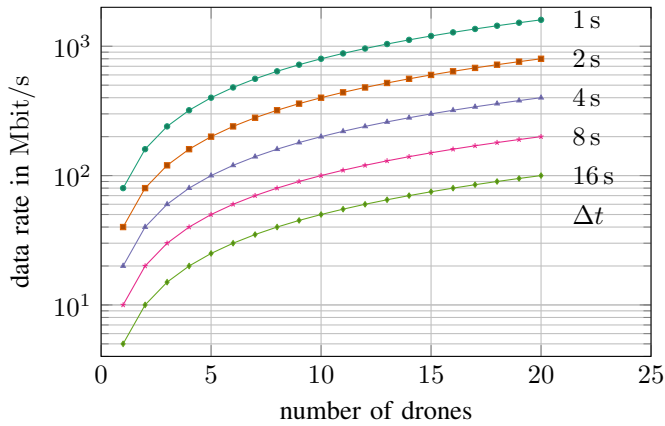


Fig. 7: Calculated aggregated payload data rate for visual monitoring over the number of drones for varying capture interval Δt

scalability of the system, we express the payload data rate as $r = \bar{S}/\Delta t$ with average image size \bar{S} and capture interval Δt . The aggregated payload data rate produced by N drones is then $L = \sum_{i=1}^N r_i = N \bar{S}/\Delta t$, which is shown in Figure 7 for $\bar{S} = 10$ MB. In our experiments, the rate remained below 100 Mbit/s, but increasing the number of drones or capture rate quickly raises the demand to several hundred Mbit/s. In conclusion, care must be taken when scaling the aerial monitoring system, not to exceed the network's capacity.

C. Aerial Relaying

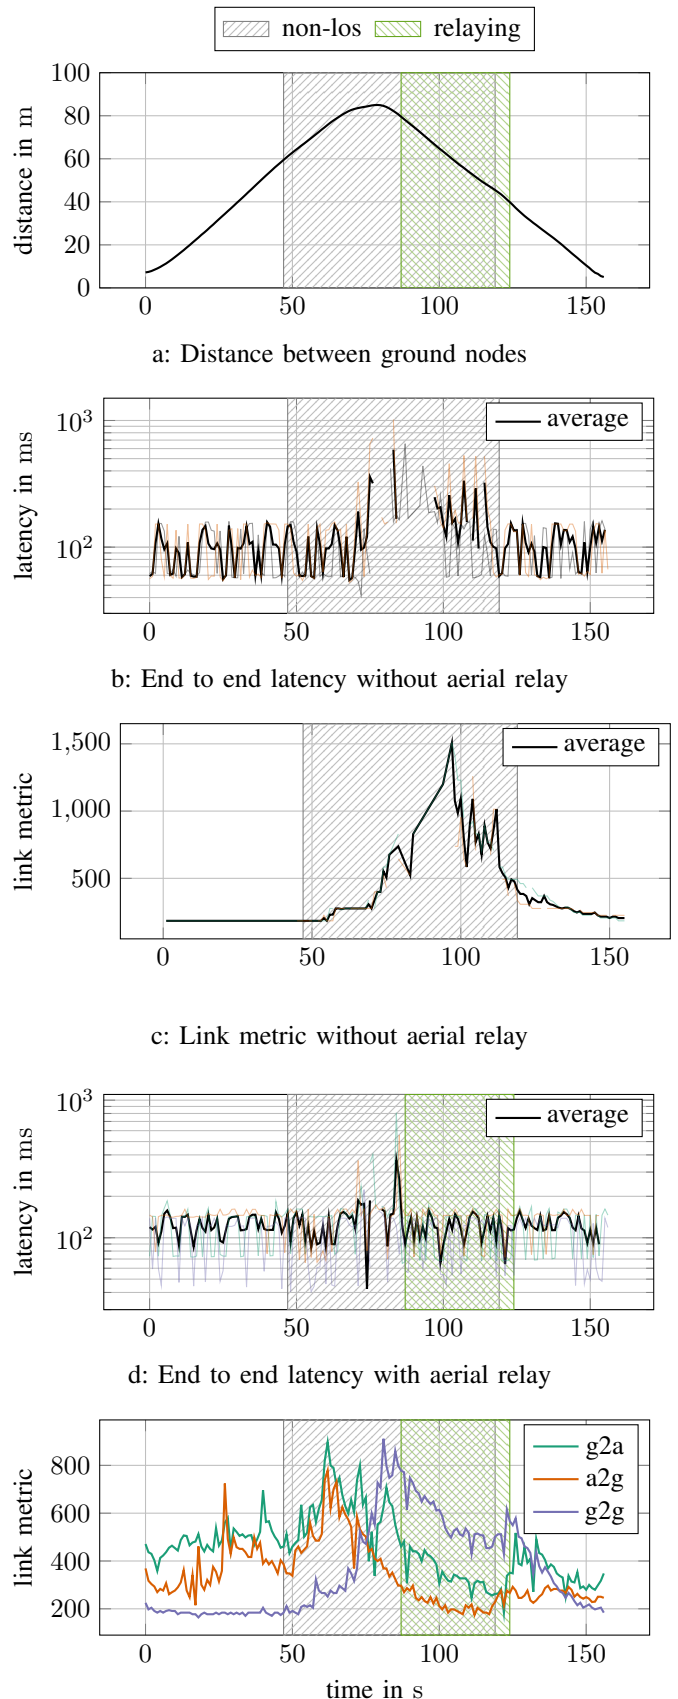
Rovers working on the ground experience severe shadow fading due to the massive container stacks. The drones can improve their connectivity by acting as aerial relays. We measure the end-to-end latency of a video stream transmitted from a rover to a local GCS with mesh Wi-Fi. Figure 8 shows the results both with and without aerial relaying.

Fig. 8a shows the distance between rover and GCS. The gray background marks the period during which there is no direct line-of-sight between them. Fig. 8b shows the end-to-end latency without an aerial relay. The average latency is about 100 ms in line-of-sight conditions. It increases to multiples of this value when the line-of-sight is obstructed. At some points in time, even link outages occur. This behavior is supported by the sharp increase in the link metric shown in Fig. 8c which is computed from the signal strength.

Fig. 8d shows the end-to-end latency if an aerial relay drone is employed during the period marked with a green background. The latency first increases (in the gray period) but they stabilizes once the relay becomes active (during the green period). The link metric shown in Fig. 8e supports this behavior.

IV. LESSONS LEARNED

Let us summarize some of the lessons learned in the course of our experimental work. These are clearly not fundamentally new scientific insights, but are intended to be a experience



e: Average link metrics g2a (ground to air), a2g (air to ground), and g2g (ground to ground)

Fig. 8: Latency of communication link with and without aerial relay for different experiments (colored) and average (black).

report that we wish to share with other researchers performing hands-on field work with multi-drone systems.

First, stable high-rate communication is absolutely crucial for autonomous operation. Control algorithms often provide a multitude of services that rely heavily on communication. This is especially critical during the initial service discovery phase, where rapid and reliable data exchange is needed to establish connections and functionalities. In our system, which is rather small, we measured packet rates up to 3.8 million packets per second during the first three seconds of a mission, which makes the demand for suitable communication infrastructure evident.

Second, the integration of visual monitoring with multiple drones significantly amplifies the volume of payload data generated. Each drone produces approximately 30 Mbit/s, and data volume increases linearly with the number of drones and capture rate. Current wireless technologies can manage this volume up to about ten drones and capture intervals in the order of several seconds. However, challenges arise if the number of drones exceeds ten or if the capture interval is below two seconds.

Third, communication breakdowns of ground vehicles in complex environments with massive objects that are moved around can be effectively mitigated by employing aerial relays, yielding reasonable latency values. The use of ad hoc routing protocols further enhances the connectivity by switching between direct links and multi-hop paths. This adaptability is crucial for maintaining reliable communication in dynamic environments with changing shadowing constellations.

V. CONCLUSIONS AND OUTLOOK

Our experiments with a multi-drone system capable of autonomous aerial inspection of freight containers in a container terminal demonstrated some communication challenges and approaches. The drones identify and localize containers and act as aerial relays to improve connectivity for ground vehicles. Measurements underscore the importance of careful system dimensioning to prevent exceeding the network capacity. Future research could focus on systematically defining the requirements for next-generation networks.

ACKNOWLEDGMENTS

Suggestions for language improvements were obtained with the aid of tools like DeepL and ChatGPT.

This work was supported in part by the CELTIC-NEXT project, 6G for Connected Sky (6G-SKY), with funding received from the Austrian Federal Ministry for Climate Action, Environment, Energy, Mobility Innovation and Technology (BMK) via FFG under grant agreement no. 891478 and by the 3D-NET project, with funding received from the Austrian BMK via FFG under grant agreement no. 923328.

REFERENCES

- [1] M. Mozaffari, W. Saad, M. Bennis, Y.-H. Nam, and M. Debbah, "A tutorial on UAVs for wireless networks: Applications, challenges, and open problems," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2334–2360, 2019.
- [2] Y. Zeng, Q. Wu, and R. Zhang, "Accessing from the sky: A tutorial on UAV communications for 5G and beyond," *Proceedings of the IEEE*, vol. 107, no. 12, pp. 2327–2375, 2019.
- [3] J. Stanczak, "NR support for UAVs," *3GPP Highlights*, May 2023.
- [4] E. Yanmaz, R. Kuschig, and C. Bettstetter, "Achieving air-ground communications in 802.11 networks with three-dimensional aerial mobility," in *Proc. IEEE INFOCOM*, (Turin, Italy), Apr. 2013.
- [5] J. Scherer, S. Yahyanejad, S. Hayat, E. Yanmaz, T. Andre, A. Khan, V. Vukadinovic, C. Bettstetter, H. Hellwagner, and B. Rinner, "An autonomous multi-UAV system for search and rescue," in *Proc. Workshop on Micro Aerial Vehicle Networks, Systems, and Applications for Civilian Use (DroNet)*, (Florence, Italy), May 2015.
- [6] S. Hayat, C. Bettstetter, A. Fakhreddine, R. Muzaffar, and D. Emini, "An experimental evaluation of LTE-A throughput for drones," in *Proc. ACM Workshop on Micro Aerial Vehicle Networks, Systems, and Applications (DroNet)*, (Seoul, Korea), pp. 3–8, June 2019.
- [7] R. Muzaffar, C. Raffelsberger, A. Fakhreddine, J. López Luque, D. Emini, and C. Bettstetter, "First experiments with a 5G-connected drone," in *Proc. ACM Workshop on Micro Aerial Vehicle Networks, Systems, and Applications (DroNet)*, (Toronto, Canada), June 2020.
- [8] S. Pandi, F. Gabriel, O. Zhdanenko, S. Wunderlich, and F. H. P. Fitzek, "Meshmerize: An interactive demo of resilient mesh networks in drones," in *IEEE Annual Consumer Communications & Networking Conference (CCNC)*, 2019.
- [9] A. Corsaro, L. Cominardi, O. Hecart, G. Baldoni, J. E. P. Avital, J. Loudet, C. Guimares, M. Ilyin, and D. Bannov, "Zenoh: Unifying communication, storage and computation from the cloud to the microcontroller," in *Euromicro Conference on Digital System Design (DSD)*, pp. 422–428, 2023.