

Alma Mater Studiorum Università di Bologna  
Archivio istituzionale della ricerca

Mobility-Aware Orchestration for UAV-Enabled IoT Networks in Emergency Scenarios

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

*Published Version:*

Heideker, A., Interdonato, G., Pizzi, S., Molinaro, A., Trotta, A. (2025). Mobility-Aware Orchestration for UAV-Enabled IoT Networks in Emergency Scenarios. New York : Institute of Electrical and Electronics Engineers Inc. [10.1109/mass66014.2025.00126].

*Availability:*

This version is available at: <https://hdl.handle.net/11585/1037091> since: 2026-01-14

*Published:*

DOI: <http://doi.org/10.1109/mass66014.2025.00126>

*Terms of use:*

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).  
When citing, please refer to the published version.

(Article begins on next page)

# Mobility-Aware Orchestration for UAV-Enabled IoT Networks in Emergency Scenarios

Alexandre Heideker\*, Giovanni Interdonato†, Sara Pizzi‡, Antonella Molinaro‡, Angelo Trotta\*

\*Department of Computer Science and Engineering, University of Bologna, Italy

† Department of Electrical and Information Engineering, University of Cassino, Italy

‡ Department of Information, Infrastructure and Sustainable Energy Engineering,  
University Mediterranea of Reggio Calabria, Italy

{alexandre.heideker, angelo.trotta5}@unibo.it, giovanni.interdonato@unicas.it, {sara.pizzi, antonella.molinaro}@unirc.it

**Abstract**—Unmanned Aerial Vehicles (UAVs) integrated with Internet of Things (IoT) systems represent a powerful solution for reestablishing connectivity in emergency scenarios where terrestrial infrastructure is damaged or overloaded. However, ensuring service continuity under UAV mobility and resource constraints poses significant orchestration challenges. This paper presents a mobility-aware orchestration framework tailored for UAV-enabled IoT networks in such critical contexts. Leveraging Kubernetes for service orchestration and Prometheus for telemetry monitoring, we design and implement a modular architecture that dynamically reallocates containerized services across mobile UAV nodes in real time. Our framework is validated through physical testbed experiments on Raspberry Pi-equipped UAVs, where container migration is triggered by geographic constraints and monitored across varying pod workloads and link conditions. Results demonstrate robust adaptability, sub-second recovery under favorable link quality, and a scalable orchestration strategy for mission-critical operations. This work advances the practical deployment of orchestrated UAV swarms, bridging the gap between theoretical frameworks and real-world mobility-aware computing.

**Index Terms**—UAV networks, emergency communications, container migration, mobility-aware computing, edge computing, real-world deployment

## I. INTRODUCTION

The integration of Unmanned Aerial Vehicles (UAVs) and the Internet of Things (IoT) has emerged as a critical enabler for resilient, flexible, and responsive network infrastructures. This becomes particularly vital in emergency contexts such as natural disasters, critical infrastructure breakdowns, or post-disaster recovery phases, where traditional terrestrial communication networks are often damaged, unavailable, or overloaded [1]. Ensuring communication continuity and supporting timely decision-making under these constraints is fundamental to enabling effective emergency response and coordination.

UAVs offer rapid deployment, flexible mobility, and the ability to operate in disconnected or infrastructure-less environments, making them highly suited for collecting sensor data, supporting localized processing, and relaying communication in emergency scenarios [2]. These capabilities enable UAVs to complement terrestrial and satellite systems by bridging coverage gaps and supporting delay-tolerant or latency-sensitive services for rescue teams, first responders, and IoT-based monitoring applications. However, these benefits come with significant challenges.

Emergency scenarios are characterized by high variability in both spatial and temporal dimensions, requiring real-time adaptation of UAV roles, paths, and service configurations. UAVs must operate under energy and bandwidth constraints, manage dynamic workloads, and deliver critical services with minimal human intervention. This necessitates a flexible, robust orchestration framework capable of monitoring UAV and network states and dynamically deploying computing and caching functionalities [3].

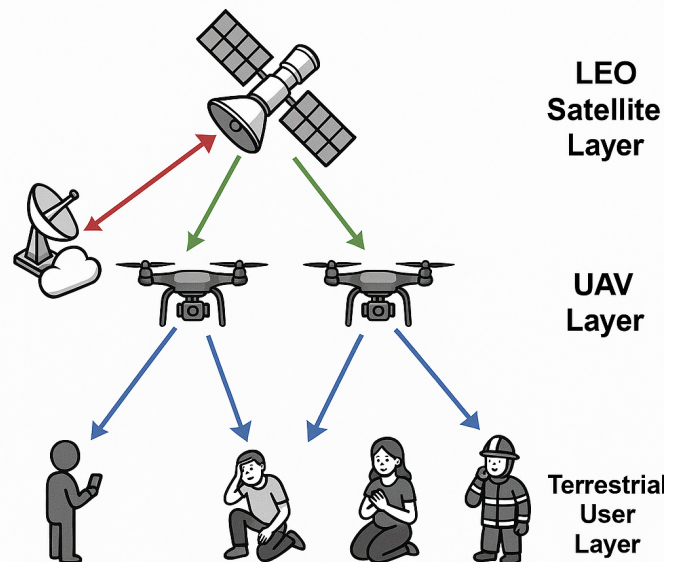


Fig. 1: Multi-layer architecture proposed for emergency scenarios. UAVs act as communication and computation relays between isolated users, IoT devices, and satellite backhaul.

In this work, we present a novel mobility-aware orchestration framework for UAV-enabled IoT networks operating in emergency conditions. The proposed architecture combines Kubernetes for orchestration and Prometheus for continuous telemetry collection. UAVs are equipped with onboard companion computers that host containerized services and coordinate with the orchestration logic to support offloading and adaptive communication.

The main contributions of this work are as follows:

- We propose and implement a mobility-aware orchestration framework tailored to the constraints and dynamics of UAV-enabled IoT deployments in emergency scenarios.
- We realize a fully integrated prototype using Kubernetes and Prometheus, and demonstrate its capabilities on real UAV hardware.
- We evaluate the system through both simulation and experimental validation, showing clear benefits in latency reduction, resilience, and task allocation efficiency.

Our proposed solution is situated within a broader three-layer communication architecture (see Fig. 1), where satellite links provide backhaul connectivity, UAVs operate as aerial relays and computing agents, and ground-level users or sensors generate data or require services [4]. UAVs dynamically coordinate to support the rescue team by bridging fragmented terrestrial links and hosting temporary edge services.

The remainder of the paper is structured as follows. Section II reviews related work and positions our contribution within the current research landscape. Section III presents the architectural details of the proposed orchestration framework. Section IV describes the implementation and testbed setup. Section V discusses experimental evaluations and results. Finally, Section VI concludes the paper and outlines future research directions.

## II. RELATED WORK

Research on UAV-enabled communication systems spans a wide range of domains, from emergency response applications to distributed task orchestration across aerial platforms. In this section, we first review recent advances in using UAVs to support connectivity and situational awareness in emergency scenarios. We then focus on works addressing task orchestration in UAV-based systems, highlighting the limited availability of real-world implementations.

### A. UAVs for Emergency Scenarios

Unmanned Aerial Vehicles (UAVs) have become central to the development of resilient emergency communication systems, especially in the aftermath of natural disasters, large-scale accidents, or infrastructure failures. Their rapid deployability, aerial flexibility, and independence from terrestrial infrastructure make UAVs a viable option for ensuring basic connectivity and data dissemination under harsh and unpredictable conditions.

Recent efforts have explored both architectural frameworks and real-world deployments of UAVs in disaster response. For example, the WiND project demonstrated the feasibility of UAV-mounted Raspberry Pi nodes forming Wi-Fi-based mesh networks across affected regions, supporting captive portals and synchronization mechanisms to assist survivors [5]. Such hands-on deployments reveal the practical potential of UAVs to form ad hoc networks in post-disaster scenarios. At the architectural level, Zhao et al. [6] proposed a layered design

in which UAVs act as relays, dynamically adjusting trajectories and device-to-device (D2D) routing to interconnect isolated ground areas. This dynamic control is further expanded in recent deep reinforcement learning (DRL)-based studies. For instance, Liu et al. [7] introduced the HybridComm framework, which concurrently optimizes the UAV's position, bandwidth, and uplink/downlink ratios while enforcing communication prioritization for rescuers. Their DRL-based optimization approach substantially increases ground coverage and transmission quality compared to static reward models. In terms of protocol and connectivity, Liu et al. [8] delve into UAV-IoT convergence in disaster environments, proposing a layered architecture that incorporates lightweight communication protocols and intelligent relaying. This complements the adaptive deployment strategy described by Jeong et al. [9], where UAV placement is governed by ground user density and service urgency, resulting in optimized coverage and reduced access delays. Trajectory control has also emerged as a key enabler for efficient UAV emergency networking. Christy et al. [10] apply reinforcement learning techniques to derive safe and energy-aware UAV paths, particularly when end-user devices are power-constrained. Their approach highlights how navigation decisions can directly influence service latency and mission reliability.

The role of UAVs extends beyond communication, as demonstrated by Surman et al. [11], who evaluate UAV reliability and integration in emergency medical scenarios, and by Yang et al. [12], who present SatellStitch, a UAV-based image stitching method that removes the need for GNSS or GCPs by leveraging satellite references for geospatial accuracy. Such imaging systems can assist in situational awareness and damage assessment when conventional mapping tools are unavailable.

### B. UAV Task Orchestration and Real Deployments

Despite UAVs being increasingly envisioned as mobile compute nodes, most task orchestration research remains confined to simulations or architectural models. Few contributions validate real-world deployments capable of dynamic containerized task migration.

FOCUS [13] offers one of the earliest implementations, leveraging fog computing and SDN to manage tasks across UAV mesh networks, yet without using modern container orchestration tools like Kubernetes. AirEdge [14] proposes dependency-aware scheduling for federated aerial computing but stops short of deployment on resource-constrained UAVs. Other works emphasize mission-driven orchestration. Suomalainen et al. [15] target public safety services across heterogeneous platforms, while Cao et al. [16] focus on end-to-end slice management for UAV-assisted 6G networks. However, both remain conceptual or simulator-based. Several studies integrate trajectory and task scheduling, as in [17], which uses multi-objective reinforcement learning for emergency coordination. Soliman et al. [18] introduce Kubernetes-based orchestration for agricultural drone swarms using blockchain, yet their design is domain-specific. Aerialist [19] simulates

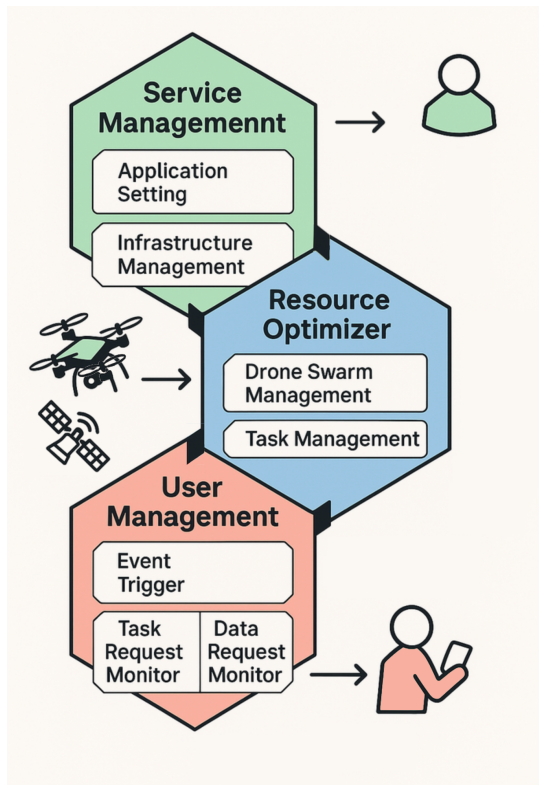


Fig. 2: Framework architecture

Kubernetes orchestration over UAVs but lacks physical validation.

Altogether, these contributions reveal a growing interest in UAV task orchestration but also highlight a research gap: the limited number of practical, modular, and portable orchestration implementations validated in real UAV platforms. Our work addresses this gap by integrating Kubernetes-based orchestration directly onto drones and validating live task migration triggered by UAV mobility, demonstrating both feasibility and robustness in dynamic aerial networks.

### III. MOBILITY-AWARE ORCHESTRATION FRAMEWORK

The proposed mobility-aware orchestration framework has been proposed in [4], and it integrates UAV swarms into IoT networks, specifically designed to address the dynamic demands and stringent performance requirements of emergency scenarios. The architecture is structured into three synergistic layers: the User Management Layer (UML), Service Management Layer (SML), and Resource Optimizer Layer (ROL), each addressing distinct operational responsibilities.

#### A. User Management Layer (UML)

The UML focuses on monitoring, managing, and predicting the behaviors of users and IoT devices. It continuously captures real-time telemetry, including user mobility patterns, data demand dynamics, and conditions of IoT sensors. Advanced data analytics modules integrated within UML leverage machine learning algorithms to predict future connectivity

demands and user locations, providing proactive input to the higher orchestration layers.

#### B. Service Management Layer (SML)

The SML oversees strategic resource allocation and system-wide configuration, leveraging Kubernetes' orchestration capabilities to optimize resource utilization. It utilizes Prometheus for collecting and managing telemetry data and Grafana for visualization and real-time monitoring. This centralized management layer dynamically configures UAV swarm behaviors and adjusts computational workloads, caching strategies, and QoS policies to rapidly adapt to evolving scenario requirements and resource constraints.

#### C. Resource Optimizer Layer (ROL)

The ROL represents the intelligent decision-making core, optimizing UAV swarm positioning, IoT data caching, and computational task allocation in real-time. Specifically regarding this work, it integrates two specialized modules:

a) *Drone Swarm Management:* This module optimally positions UAVs considering energy constraints, communication demands, and network coverage requirements. Using real-time telemetry from Prometheus, it computes UAV trajectories that enhance coverage and extend operational longevity.

b) *Task Management:* Efficiently allocates computational tasks between UAV nodes and terrestrial infrastructure based on real-time resource availability, UAV battery status, and computational load metrics. The orchestration strategy employs a hybrid centralized-distributed decision-making process to maintain scalability and responsiveness.

This comprehensive mobility-aware orchestration framework ensures efficient resource utilization, robust communication, and high adaptability, thereby significantly enhancing the performance of UAV-enabled IoT networks in critical emergency scenarios.

### IV. IMPLEMENTATION AND TESTBED SETUP

The practical realization of our mobility-aware orchestration framework involves an intricate integration of advanced software tools and robust hardware components, designed to demonstrate feasibility, scalability, and real-time responsiveness under realistic conditions. We established a testbed comprising multiple UAV nodes, terrestrial edge infrastructure, and IoT sensor devices to replicate a representative emergency communication scenario.

Each UAV platform is based on the Holybro X650 quadcopter frame, equipped with a Pixhawk 6C flight controller and a Raspberry Pi 5.0 companion computer. The Pixhawk handles low-level navigation and control tasks using autopilot firmware<sup>1</sup>, while the Raspberry Pi is responsible for executing higher-layer logic, including containerized applications and orchestration agents. Communication between the flight controller and the companion board is established over a UART interface, with the MAVLink protocol enabling application-layer coordination and telemetry exchange.

<sup>1</sup><https://ardupilot.org>

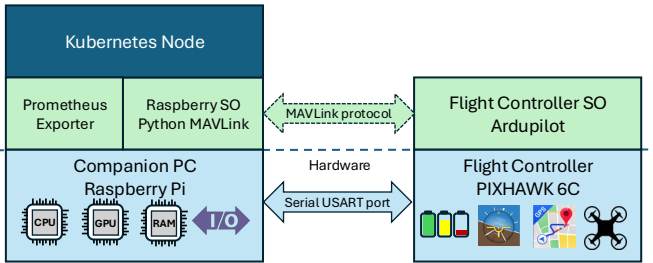


Fig. 3: Deployed modules and onboard software architecture. The Raspberry Pi companion computer is connected to the Pixhawk flight controller via UART. MAVLink is used at the application layer to manage telemetry and control commands, enabling coordination between low-level navigation and high-level orchestration tasks.

Figure 3 illustrates the structure of the deployed modules and the software-hardware integration adopted in our implementation. The hardware connection between the flight controller (e.g., Pixhawk 6C) and the companion computer (e.g., Raspberry Pi) is realized through a serial UART interface. At the software layer, the MAVLink protocol is used for real-time communication, enabling the exchange of telemetry data and navigation commands. This setup allows the Raspberry Pi to monitor flight status, adjust trajectories, and issue control directives when needed, while simultaneously handling orchestration logic and service deployment. The onboard stack includes a Kubernetes agent that enables each UAV to join the orchestration cluster. Resource usage, positioning data, and battery metrics are continuously monitored using Prometheus exporters deployed on each companion node. These metrics are tagged with UAV identifiers and transmitted to a central Prometheus server, where they are stored as time-series data.

Prometheus’s pull-based architecture and flexible label-based query model facilitate efficient aggregation of UAV status information. The orchestration controller queries this database periodically to inform scheduling decisions. Kubernetes is extended with a custom scheduler plugin that uses real-time metrics to rank available UAVs based on energy availability, geographic proximity to the task area, and CPU load. This hybrid centralized-distributed control allows global visibility without sacrificing local autonomy.

## V. PERFORMANCE EVALUATION

To validate the proposed mobility-aware orchestration framework in realistic conditions, we conducted a detailed experimental campaign using physical UAV nodes and edge hardware. This evaluation complements the system design by measuring orchestration responsiveness during node migration events triggered by mobility constraints.

The testbed consists of two Raspberry Pi 5 devices, each acting as a UAV companion node running a Kubernetes-managed microservice platform. The Raspberry Pis communicate over WiFi with an access point emulating a static base station. Each device runs containerized services (pods), with a




Fig. 4: The deployed UAV platform with companion Raspberry Pi and Pixhawk flight controller.

controller enforcing service migration whenever a node moves beyond a preconfigured geographic threshold. The migration simulates the UAV exiting the operational area and requiring its workloads to be rescheduled on an available nearby UAV.

The experimental setup focused on measuring the downtime incurred during such service migrations under different system loads. We varied the number of concurrently deployed pods across four configurations: 4, 8, 16, and 32 pods. For each configuration, we simulated different distances from the access point by altering the transmission power of the WiFi adapters (20 dBm, 10 dBm, and 5 dBm), corresponding approximately to effective ranges of 0–10m, 30–50m, and 100–120m, respectively. The goal was to emulate degradation in link quality as the UAV moves away from the communication hub.

Each configuration was repeated 30 times, and the average downtime—defined as the time between a node becoming unreachable and the service being fully restored on a different node—was measured. Results are presented with a 99% confidence interval.

Figure 4 shows the deployed UAV used in the experimental campaign, including the onboard Raspberry Pi and Pixhawk flight controller. Figure 5 illustrates the measured service downtimes across the different pod counts and signal levels. As expected, increasing the number of pods on the node leads to higher reallocation delays. At 5 dBm (representing the farthest simulated range), the average downtime with 32 pods exceeded 10 seconds, highlighting the impact of network degradation on orchestration responsiveness.

Conversely, when operating at full transmission power (20 dBm), all configurations achieved sub-second recovery, validating the framework’s ability to adapt quickly in proximity-limited deployments. These results confirm the sensitivity of the orchestration logic to both resource availability and link quality, supporting its use in time-critical emergency operations.

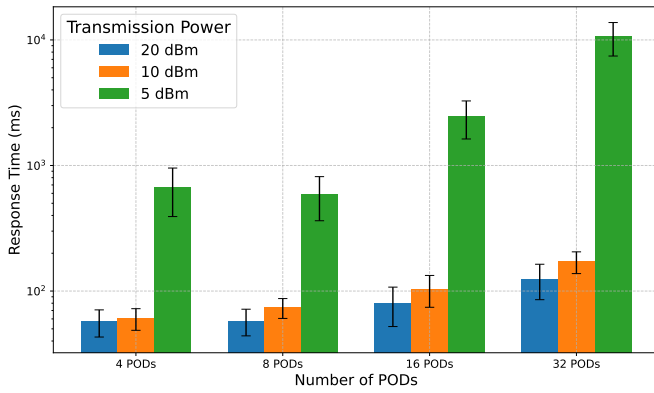


Fig. 5: Average service downtime during pod migration vs. number of allocated pods. Results include 99% confidence intervals for three transmission power levels.

## VI. CONCLUSIONS AND FUTURE WORK

This paper presented a mobility-aware orchestration framework tailored for UAV-enabled IoT networks deployed in emergency scenarios. By integrating Kubernetes-based service orchestration with Prometheus telemetry monitoring, we enabled intelligent coordination of containerized services across mobile aerial nodes. Our architecture addresses key challenges such as constrained communication, limited onboard resources, and the dynamic nature of UAV mobility. Through a combination of real-world experiments and rigorous testing, we demonstrated the ability of our system to dynamically reallocate services in response to network conditions and mobility thresholds. The evaluation highlighted the sensitivity of orchestration latency to both pod density and signal degradation, validating our approach for mission-critical deployments where continuity of service is essential. The modularity of our solution ensures it can scale across heterogeneous deployments, supporting multi-layer architectures involving ground users, UAVs, and satellite backhaul. Its compatibility with standard orchestration technologies also makes it suitable for integration with existing edge/cloud infrastructures.

Future work will focus on three main directions. First, we aim to extend our orchestration logic by incorporating predictive analytics to anticipate UAV mobility and proactively schedule service migrations. Second, we plan to integrate federated learning capabilities to support collaborative decision-making across the UAV swarm, even under intermittent connectivity. Finally, we will expand the scale of our testbed to include a fleet of UAVs with heterogeneous sensing and actuation capabilities, enabling a broader range of emergency response applications such as real-time video processing, multi-hop communication, and adaptive coverage optimization.

## ACKNOWLEDGMENT

The work has been funded by the European Union – Next Generation EU - Call PRIN 2022 - D.D. n. 104 of 02/02/2022 - Project RAIN4C - “Reliable Aerial and satellite Networks:

joint Communication, Computation, Caching for Critical scenarios” (Project id: 20227N3SPN, CUP: J53D23007020001).

## REFERENCES

- [1] A. Khan, S. Gupta, and S. K. Gupta, “Emerging uav technology for disaster detection, mitigation, response, and preparedness,” *Journal of Field Robotics*, vol. 39, no. 6, pp. 905–955, 2022.
- [2] M. Erdelj and E. Natalizio, “Uav-assisted disaster management: Applications and open issues,” in *2016 international conference on computing, networking and communications (ICNC)*. IEEE, 2016, pp. 1–5.
- [3] B. Anuraj, D. Calvaresi, J.-M. Aerts, and J.-P. Calbimonte, “Dynamic swarm orchestration and semantics in iot edge devices: A systematic literature review,” *Ieee access*, 2024.
- [4] A. Trotta, A. Heideker, I. Zyrianoff, G. Interdonato, and S. Pizzi, “(poster) advancing non-terrestrial networks for critical scenarios with the rain4c framework,” in *2024 20th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*. IEEE, 2024, pp. 780–782.
- [5] K. G. Panda, S. Das, D. Sen, and W. Arif, “Design and deployment of uav-aided disaster emergency network,” *IEEE access*, vol. 7, pp. 102 985–102 999, 2019.
- [6] N. Zhao, W. Lu, M. Sheng, Y. Chen, J. Tang, F. R. Yu, and K.-K. Wong, “Uav-assisted emergency networks in disasters,” *IEEE Wireless Communications*, vol. 26, no. 1, pp. 45–51, 2019.
- [7] L. Zhao, X. Liu, and T. Shang, “Maximizing coverage in uav-based emergency communication networks using deep reinforcement learning,” *Signal Processing*, vol. 230, p. 109844, 2025.
- [8] H. Hydher, D. N. K. Jayakody, K. T. Hemachandra, and T. Samarasinghe, “Uav deployment in wsn system for emergency/remote area applications,” *Computer Networks*, vol. 257, p. 110977, 2025.
- [9] N. Lin, Y. Liu, L. Zhao, D. O. Wu, and Y. Wang, “An adaptive uav deployment scheme for emergency networking,” *IEEE Transactions on Wireless Communications*, vol. 21, no. 4, pp. 2383–2398, 2021.
- [10] T. Zhang, J. Lei, Y. Liu, C. Feng, and A. Nallanathan, “Trajectory optimization for uav emergency communication with limited user equipment energy: A safe-dqn approach,” *IEEE Transactions on Green Communications and Networking*, vol. 5, no. 3, pp. 1236–1247, 2021.
- [11] K. Surman and D. Lockey, “Unmanned aerial vehicles and pre-hospital emergency medicine,” *Scandinavian journal of trauma, resuscitation and emergency medicine*, vol. 32, no. 1, p. 9, 2024.
- [12] Z. Wei, C. Lan, Q. Xu, L. Wang, T. Gao, F. Yao, and H. Hou, “Satellstitch: Satellite imagery-assisted uav image seamless stitching for emergency response without gcp and gnss,” *Remote Sensing*, vol. 16, no. 2, p. 309, 2024.
- [13] G. Secinti, A. Trotta, S. Mohanti, M. Di Felice, and K. R. Chowdhury, “Focus: Fog computing in uas software-defined mesh networks,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 6, pp. 2664–2674, 2019.
- [14] U. Awada, J. Zhang, S. Chen, and S. Li, “Airedge: A dependency-aware multi-task orchestration in federated aerial computing,” *IEEE Transactions on Vehicular Technology*, vol. 71, no. 1, pp. 805–819, 2021.
- [15] J. Suomalainen, K. Ahola, M. Sallio, G. Kiss, G. Megyaszi, R. Asif, P. Jehkonen, and J. Rivalan, “Tactical orchestration: Network, security, and drone intelligence for mission-critical operations,” in *2024 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*. IEEE, 2024, pp. 1061–1066.
- [16] H. Cao, N. Kumar, L. Yang, M. Guizani, and F. R. Yu, “Resource orchestration and allocation of e2e slices in software-defined uavs-assisted 6g terrestrial networks,” *IEEE Transactions on Network and Service Management*, vol. 21, no. 1, pp. 1032–1047, 2023.
- [17] S. Akter, D. V. A. Duong, and S. Yoon, “Joint optimization of uav trajectory, task offloading, and resource allocation in uav-aided emergency response operations,” *IEEE Internet of Things Journal*, 2025.
- [18] S. Soliman, A. Bendary, and H. Dahshan, “Enhancing agricultural efficiency with blockchain-orchestrated drone swarms and kubernetes,” in *2024 6th Novel Intelligent and Leading Emerging Sciences Conference (NILES)*. IEEE, 2024, pp. 567–570.
- [19] S. Khatiri, S. Panichella, and P. Tonella, “Simulation-based testing of unmanned aerial vehicles with aerialist,” in *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings*, 2024, pp. 134–138.