

Swarm Network-as-a-Service (SNaaS)

Balsam Alkouz, Osama Amin, and Basem Shihada

Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division

King Abdullah University of Science and Technology (KAUST), Saudi Arabia

Email: {balsam.alkouz, osama.amin, basem.shihada}@kaust.edu.sa

Abstract—Emerging on-demand connectivity scenarios increasingly require networking solutions with stringent service-level guarantees. We propose Swarm Network-as-a-Service (SNaaS), a service-oriented framework that leverages fleets of drones to provide on-demand connectivity at scale. SNaaS explicitly models drone-to-device and drone-to-drone interactions as composable services, enabling consumers to request connectivity through Service-Level Agreements (SLAs). We formalize atomic and composite SNaaS services, present an SDN-inspired architecture that integrates the service-oriented triad of provider, consumer, and registry. We introduce a composition framework that orchestrates drones into end-to-end services. Within this framework, we define and analyze three composition strategies, i.e., direct, clustered, and parallel, and propose a queuing-theory-based heuristic for selecting the most suitable strategy under varying load conditions. A dedicated enforcement module continuously monitors queue stability and SLA latency, adaptively reconfiguring the swarm when violations occur. Experiments using real air-to-ground measurements show that the framework consistently outperforms fixed compositions, achieving lower latency, fewer SLA violations, and smoother adaptation as load and swarm size increase.

Index Terms—Swarm Network-as-a-Service (SNaaS), Drone-as-a-Service, UAV Assisted Networks, Service Composition, 6G, Queuing Theory.

I. INTRODUCTION

The rise of smart and cognitive cities, coupled with the rollout of 6G networks enabling ultra-low latency and massive connectivity, is driving unprecedented data growth and demanding networking infrastructures that can meet extreme service-level requirements at scale [1]–[3]. Importantly, this need for on-demand, high-capacity networking is not limited to futuristic cities. It is already evident in scenarios such as *temporary high-density events*. For example, the FIFA World Cup create connectivity demand that spikes far beyond the capacity of fixed infrastructure. Current industry practice relies on Cells on Wheels (COWs), which are costly, rigid, and limited in scalability [4]. Similarly, in *infrastructure-less environments* such as regional construction sites, offshore oil rigs, or disaster zones, connectivity is typically provided by satellites. These are slow to deploy, expensive, and unsuitable for low-latency, high-throughput requirements [5].

Drones or Unmanned Aerial Vehicles (UAVs) provide a compelling alternative. They are agile, quickly deployable, and capable of forming airborne networks on-demand. Yet, due to bandwidth, hardware, and scheduling constraints, a single UAV can typically support only a limited number of concurrent devices (on the order of a few tens), making it insufficient for high-demand scenarios [6]. To provide meaningful capacity,

drones must operate as a swarm, collectively serving thousands of devices. Furthermore, in infrastructure-less environments, drones can act as relays, extending connectivity across vast or obstructed terrains. This introduces a new paradigm: swarm-based, on-demand networking infrastructure.

Research in Non-Terrestrial Networks (NTNs) has explored the role of drones, balloons, and satellites in providing connectivity [7], [8]. However, existing efforts have largely focused on algorithmic or architectural optimizations without framing the problem in terms of services [9]. We argue that the *service paradigm* provides a natural abstraction by separating the *functional property*, which defines the provision of connectivity services within a target area, from *non-functional properties*, which describe the Quality of Service (QoS) requirements. In this context, key QoS requirements such as end-to-end latency, aggregate throughput, and signal reliability can be explicitly specified in Service-Level Agreements (SLAs). Modeling **Swarm Network-as-a-Service (SNaaS)** allows consumers to request connectivity based on these guarantees, without concerning themselves with the underlying deployment, coordination, or management of the UAVs.

Viewing SNaaS through the service paradigm introduces challenges in *allocation*, *composition*, *scaling*, and *robustness*, all of which require rethinking traditional networking under a service-oriented lens. In particular, *composition* concerns assembling drone-to-device and drone-to-drone links into end-to-end services that satisfy global QoS guarantees such as latency, throughput, and resilience. *This work focuses on the composition problem, optimizing how individual assignments are orchestrated into reliable, QoS-aware connectivity services.*

Prior work on Swarm-based Drone-as-a-Service (SDaaS) modeled swarms for delivery tasks, optimizing routes in skyway networks under payload and charging constraints [10]. SDaaS assumes fixed paths and focuses on delivery deadlines [11]. In contrast, SNaaS addresses *connectivity*, operating in *open, dynamic airspace* where drones must compose communication links to devices and each other. The optimization goals also shift from minimizing delivery time to ensuring QoS guarantees such as throughput, latency, and resilience.

This paper lays the groundwork for SNaaS by formally defining the paradigm, presenting a service-oriented architecture, and addressing the fundamental challenge of composition in a confined area. The goal is to establish SNaaS as a foundation for a new line of research at the intersection of service computing and non-terrestrial networking.

The contributions of this work are as follows:

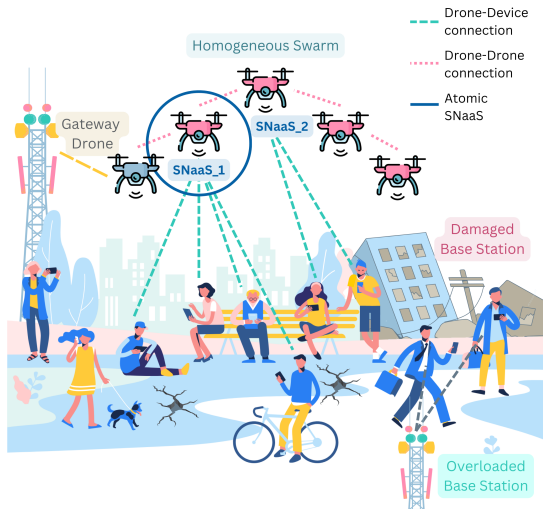


Fig. 1. Illustration of an SNaaS deployment in a disaster recovery scenario. Devices connect to nearby drones (drone-to-device links), drones forward data among themselves (drone-to-drone links), and traffic exits through a gateway drone or ground station to reach the core network.

- We introduce SNaaS for UAV-assisted connectivity, formally defining atomic and composite services that characterize its functional and non-functional attributes.
- We design an integrated Service-Oriented Architecture and Software-Defined Networking architecture that enables discovery, composition, and enforcement of swarm connectivity services while abstracting away swarm coordination from consumers.
- We propose and evaluate a queueing theory driven composition and enforcement framework that dynamically selects and adapts service compositions to preserve latency and stability guarantees under varying load.
- We validate the framework using real air-to-ground measurements, demonstrating efficient performance and smooth adaptation across varying load and swarm scales.

A. Motivating Scenario

Consider a disaster recovery operation where terrestrial communication infrastructure has been destroyed. First responders and civilians in the affected area require immediate, reliable connectivity to coordinate rescue activities, transmit critical data, and maintain communication with external agencies. Terrestrial solutions are too slow to deploy, and satellites are costly and unable to provide the required low-latency, high-throughput connections. To bridge this gap, a swarm of drones is deployed to act as an airborne communication infrastructure.

In this setting, three layers of connectivity emerge, as shown in Fig. 1. At the first layer, *devices* (e.g., mobile phones, sensors, laptops) connect to nearby drones through *drone-to-device links*. At the second layer, drones communicate with each other via *drone-to-drone links*, where drones forward traffic among themselves forming a multi-hop aerial relay network. At the third layer, one or more *gateway drones* act as the *bridge to the core network*, forwarding traffic from the

swarm to terrestrial backhaul infrastructure such as fiber, 5G base stations, or satellite uplinks.

For tractability, we assume a fleet of homogeneous drones with identical communication range, energy capacity, and throughput. The drones are positioned at fixed hovering locations in the affected area, ensuring predictable coverage. Each drone can serve only a limited number of devices at once, and every device must be assigned to exactly one drone. We also assume that drones approaching battery depletion are proactively replaced by standby drones that take over their hovering positions, with a brief overlap period during which both drones operate concurrently to prevent service disruption. Energy optimization is outside the scope of this paper and will be considered in future work.

We define an *atomic SNaaS service* as the bundle of one drone serving its allocated devices together with the drone-to-drone relay links necessary to forward data toward the gateway. To achieve end-to-end connectivity across the deployment, *multiple atomic services must be composed into a single, cohesive service*. The central challenge therefore lies in designing efficient *composition strategies* that *orchestrate* these atomic services into an end-to-end network. Unlike allocation, which assigns devices to individual drones, composition determines how the resulting assignments are connected and coordinated to *meet global QoS requirements*. In this paper, we focus specifically on optimizing service composition to minimize end-to-end *latency*, subject to constraints on drone capacity and coverage.

The difficulty of composition in SNaaS arises from the trade-offs between different ways of orchestrating atomic services. A *direct composition*, where each drone forwards traffic straight to the gateway, minimizes hop count but quickly creates congestion at the gateway drone. *Clustered compositions* reduce congestion by aggregating traffic at cluster-head drones, but require careful coordination to prevent bottlenecks at the heads. *Parallel compositions* distribute traffic across multiple relay paths, but each additional hop increases end-to-end latency, making long relay chains costly despite their load-balancing benefits. These conflicting trade-offs make it difficult to determine which composition style is best suited for a given deployment.

II. RELATED WORK

UAVs have emerged as core enablers in NTN, acting as airborne base stations, relays, or edge nodes to extend coverage and capacity in 5G/6G environments [12]. Prior studies focus on optimizing placement, routing, and spectrum reuse to improve throughput, latency, and energy efficiency [13]–[16]. For example, joint drone–user association and 3D positioning have been studied using optimal, greedy, and distributed learning-based schemes [17], [18]. Some recent efforts integrate UAVs into 6G architectures to enable dynamic control and network slicing [19]–[21]. However, these works treat UAVs as network components rather than *services*.

The concept of Drones-as-a-Service (DaaS) has been applied mainly to delivery and mission execution [22]. Existing

frameworks in this domain define UAV services for logistics, package routing, and task assignment, often using cloud-based orchestration [23]. Service composition and optimization in these delivery frameworks typically aim to minimize distance, energy consumption, and delivery delay [24], [25].

Swarm-based UAV research demonstrates how coordinated fleets outperform individual drones in coverage, fault tolerance, and scalability [26]–[28]. Cooperative control, formation flight, and multi-agent decision making have been extensively studied using consensus and reinforcement learning frameworks [29], [30]. Our approach leverages these coordination insights but applies them to a *service composition context*: drones are orchestrated into composable, SLA-compliant connectivity chains rather than purely cooperative formations.

To situate our work within this line of research, it is important to distinguish between delivery swarm services and connectivity-oriented swarm networking services. In SDaaS, the service is explicitly defined as *the delivery of packages over a skyway segment*, and composition occurs when a package cannot be delivered via a single segment [31]. SDaaS therefore composes *segments* in a skyway network [10]. Crucially, the swarm membership is fixed because the number of drones is determined by the package set in the request. In contrast, SNaaS composes *swarm members themselves*. Here, the service is connectivity, and when a single drone cannot satisfy QoS (e.g., latency requirements) additional drones are incorporated as infrastructure elements rather than as carriers. This makes the two paradigms fundamentally distinct in both service definition and composition logic.

Several studies label their architectures as “UAV-as-a-Service” or “Networking-as-a-Service” frameworks. For example, D3S [32] introduces a four-phase pipeline (Demand–Decision–Deployment–Service) for adaptive aerial coverage, while other efforts propose service-based UAV architectures [33] or integrate SDN with queueing models to improve control plane efficiency [34]. Although these works describe themselves as *service-centric*, they *operate primarily at the infrastructure or control layer*: UAVs are optimized as network components, not provisioned as consumer-facing services with measurable guarantees. In contrast, our SNaaS framework adopts a service-computing perspective [22], [35], exposing swarm connectivity as an SLA-governed service.

Queueing theory has been widely used to analyze delay and stability in wireless and UAV networks [36]. Models such as M/M/1 and M/M/m have characterized SDN controller delays, link utilization, and UAV scheduling [34]. Our work extends this foundation by embedding queueing models into the *service composition process* itself. Rather than evaluating network latency post hoc, SNaaS uses queueing parameters to guide composition selection, enforce stability, and satisfy consumer SLAs dynamically.

III. SWARM NETWORK-AS-A-SERVICE MODEL

We now formalize the SNaaS paradigm. Let $D = \{d_1, d_2, \dots, d_m\}$ denote the set of entry drones and $U =$

$\{u_1, u_2, \dots, u_n\}$ the set of devices in the operating environment. Each device is assumed to be connected to exactly one drone, following a fixed allocation policy (e.g., assignment to the nearest drone within range). A set of gateway drones $G = \{g_1, g_2, \dots, g_\ell\}$ provides the bridge between the swarm and the terrestrial backhaul, with each gateway capable of forwarding traffic out of the aerial network.

We distinguish three functional roles for drones in the swarm. *Entry drones* are directly connected to devices, serving as the access points for traffic entering the swarm. *Relay drones* forward traffic toward gateways; they may either act purely as forwarders or simultaneously serve as entry drones for their own allocated devices. *Gateway drones* act as the bridge between the swarm and the terrestrial backhaul (e.g., fiber, satellite, or 5G base station).

A. Atomic SNaaS Service

An *atomic SNaaS service* represents the smallest unit of service delivery in the system. It consists of a single drone $d_j \in D$, the set of devices allocated to it, and the relay links required for d_j to forward its traffic toward the gateway g . Formally, we denote:

$$SNaaS_j = \underbrace{\{(u_i, d_j) \mid u_i \in U_j\}}_{\text{device-to-drone links}} \cup \underbrace{\{(d_j, d_k) \mid d_k \in N_j\}}_{\text{drone-to-drone links}}.$$

where $U_j \subseteq U$ is the set of devices allocated to d_j , and N_j is the set of neighbor drones to which d_j can forward data.

B. Composite SNaaS Service

While single SNaaS instances capture local connectivity, end-to-end networking requires composing multiple instances into a cohesive service. A *composite SNaaS service* is therefore defined as the union of all SNaaS instances across the deployed swarm D , such that all devices are connected to the gateway:

$$SNaaS_D = \bigcup_{d_j \in D} SNaaS_j,$$

A composite service is *valid* if, for every device $u_i \in U$, there exists a path (u_i, d_j, \dots, g) consisting of device-to-drone and drone-to-drone links that terminates at the gateway.

C. QoS Objective

The quality of a composite SNaaS service is measured by global QoS metrics such as end-to-end latency, coverage, and throughput. In this work, we focus on minimizing latency while ensuring that capacity constraints are respected:

$$C_d = \frac{\mu}{\lambda},$$

where μ is the drone’s service rate and λ is the per-device arrival rate.

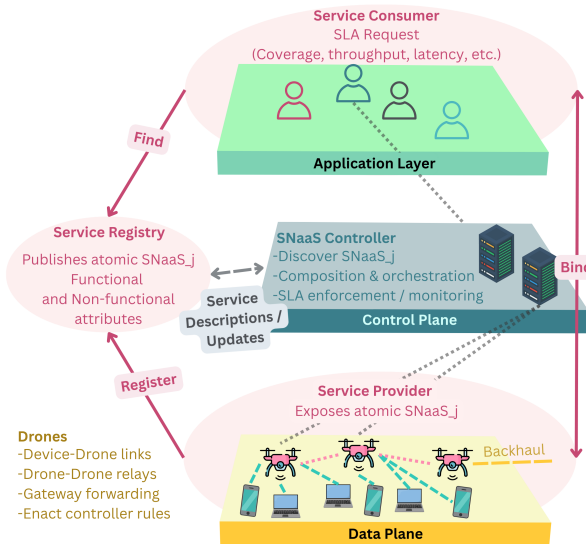


Fig. 2. SNaaS architecture integrating the Service-Oriented Architecture (SOA) triad with the Software-Defined Networking (SDN) paradigm.

IV. SNAAS ARCHITECTURE

We model the architecture of SNaaS by integrating the *Service-Oriented Architecture (SOA) triad* (provider, registry, consumer) [37] with the *Software-Defined Networking (SDN) paradigm* (application layer, control plane, data plane) [38]. Integrating SOA with SDN enables SNaaS to be both consumable and enforceable: SOA exposes connectivity as a service through SLA-based abstractions, while SDN provides the control mechanisms to compose and enforce these services over the swarm. Together, they bridge the gap between high-level service requests and low-level drone orchestration.

A. SOA Triad in SNaaS

In the SNaaS context, the *service provider* corresponds to the swarm operator. Each drone exposes an atomic SNaaS instance, which consists of the devices allocated to it and the relay links it maintains with its neighbors. These atomic services are published to the *service registry*, which functions as the directory of available connectivity services. The registry holds both functional properties, such as neighbor relations and gateway reachability, as well as non-functional attributes, such as latency, throughput, and coverage. The service consumer interacts only with the registry, submitting connectivity requests that include desired QoS parameters or SLA requirements.

B. SDN Layer Mapping

The SOA roles map directly onto the SDN layers: at the *application layer*, consumers submit SLA-driven connectivity requests while remaining agnostic to swarm details; the *control plane* functions as the SNaaS management system, discovering atomic services, composing them into end-to-end services, and enforcing QoS and capacity constraints through continuous monitoring and adaptation. While physically distributed across drones, the control plane is logically centralized, in the same

spirit as SDN controllers. The *data plane* is realized by the drones themselves, which execute the control decisions by establishing device-to-drone and drone-to-drone links and forwarding traffic to deliver the requested connectivity.

C. Interaction Flow

The interaction follows a simple flow: a consumer submits an SLA request, the control plane retrieves atomic SNaaS descriptions from the registry, computes a composition, and installs the resulting configuration across the drones. The registry serves as a passive catalog, while the control plane actively composes, enforces, and updates services based on swarm state. The data plane executes these decisions, and monitoring feedback enables continuous SLA validation and re-composition when needed. This architecture departs from prior cloud-edge-drone models [22] by embedding SDN-style plane separation within an SOA framework, enabling a clear distinction between service abstraction, orchestration, and execution. As a result, SNaaS is well suited to infrastructure-less scenarios where swarms must self-organize to deliver connectivity.

V. QUEUEING-BASED SNAAS COMPOSITION FRAMEWORK

The composition framework defines how atomic SNaaS services are orchestrated into end-to-end connectivity services (Fig. 3). It consists of three core modules: (i) **Composition strategies** (direct, clustered, and parallel) that define different orchestration patterns; (ii) **A queueing-theory-based heuristic** that evaluates each strategy under current load and topology conditions to select the one that best satisfies latency and stability requirements; and (iii) **An SLA compliance and stability enforcement module** that ensures contractual guarantees are maintained by adaptively reconfiguring the swarm when queues become unstable or latency targets are violated.

The input to the framework is the current swarm state, including the set of drones, their capacities, device-to-drone allocations, and gateway availability, together with the SLA specified by the consumer. The expected output is a composed SNaaS service: a plan that defines the drone-to-drone forwarding paths from entry drones to gateways and the scheduling policy for traffic forwarding, such that the consumer's SLA constraints are satisfied.

A. SNaaS Composition Strategies

We introduce the composition strategies that define how drones can be orchestrated to deliver end-to-end connectivity under different load and topology conditions. We identify three fundamental ways of composing atomic SNaaS services. Each type is motivated by a distinct deployment scenario, realized by a concrete algorithm for constructing the composition, and characterized by specific performance tradeoffs. Figures 4–6 illustrate the resulting service graphs.

1) *Direct Composition:* In direct composition, all entry drones forward traffic directly to the gateway without intermediate relays. This configuration minimizes hop count and provides the lowest latency under light traffic. However,

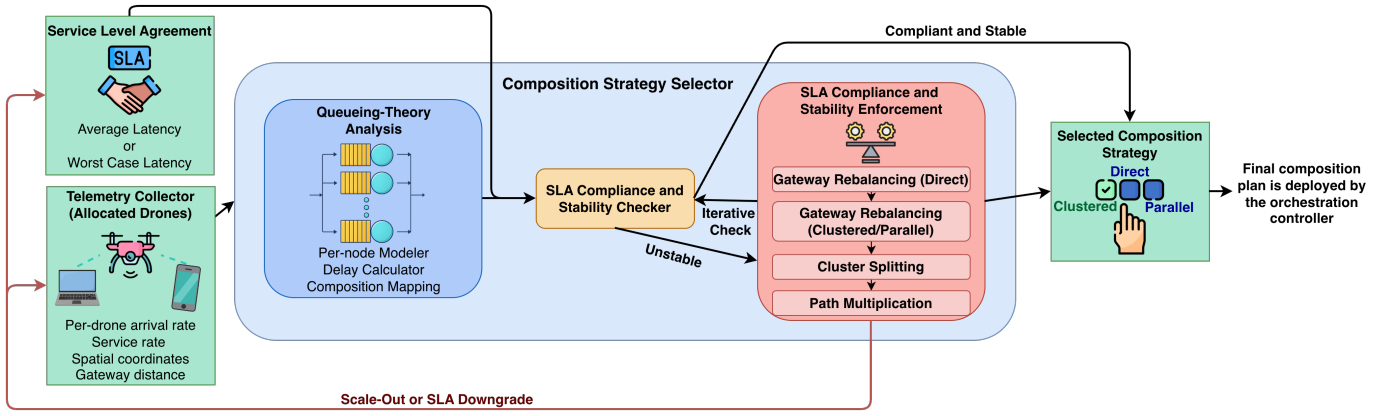


Fig. 3. Queuing-Based SNaas Composition Framework

because every entry drone injects traffic into the gateway, congestion quickly builds up under moderate to heavy loads.




Fig. 4. Direct composition: each entry drone forwards traffic directly to the gateway.

Algorithm 1 describes how entry drones are assigned to gateways in direct composition. The key decision occurs in the assignment step [Line 3], where each entry drone d_j selects a gateway g_k based on a weighted rule that combines proximity and current gateway load. Setting $\alpha = 1$ reduces the rule to nearest-gateway selection, while $\alpha = 0$ corresponds to pure load balancing. Different values of α enable hybrid policies that balance both factors.

Algorithm 1 Direct Composition Assignment

- 1: Identify the set of gateway drones $G = \{g_1, g_2, \dots, g_\ell\}$.
- 2: **for** each entry drone d_j **do**
- 3: Assign a gateway g_k by solving:

$$g_k = \arg \min_{g \in G} (\alpha \cdot \text{dist}(d_j, g) + (1 - \alpha) \cdot \text{load}(g)).$$

- 4: Forward all device traffic assigned to d_j along (d_j, g_k) in FIFO order.
- 5: **end for**

Tradeoffs. Direct composition is efficient for small-scale or lightly loaded deployments, as it minimizes delay by avoiding relays. Its main limitation is scalability: the gateway quickly becomes a bottleneck as traffic increases. When multiple gate-

ways exist, performance depends heavily on how effectively traffic is distributed among them.

2) *Clustered Composition:* In clustered composition, multiple drones act as *cluster heads* that aggregate traffic from surrounding drones before forwarding to a gateway. This reduces congestion at the gateway by distributing load across intermediate heads. The main challenge is selecting suitable cluster heads that balance proximity and traffic load, and choosing an appropriate gateway per cluster.

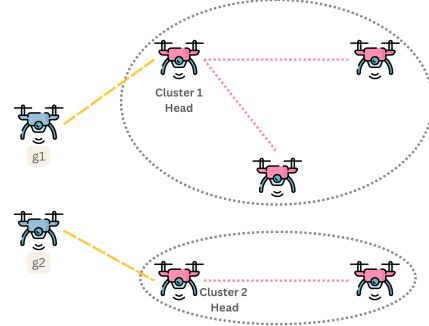


Fig. 5. Clustered composition: traffic is aggregated at cluster heads before reaching the gateway.

Algorithm 2 first fixes the number of clusters from capacity [Line 2] and forms spatial clusters via k -means [Line 3]. For each cluster, the *gateway selection* step [Line 5] chooses g^* using a weighted rule that balances the centroid's proximity to each gateway and the gateways' current load. Given g^* , the *head selection* step [Line 6] picks d^* that is both well-placed relative to g^* and not overloaded. Finally, members attach to d^* , and d^* connects directly to g^* [Line 7].

Tradeoffs. Clustered composition aggregates traffic and alleviates gateway congestion by introducing an intermediate layer of aggregation. However, the cluster heads themselves may become bottlenecks if their assigned load grows too large. The choice of k directly influences this balance: too few clusters increase the burden on individual heads, while too many clusters dilute the benefits of aggregation and may

Algorithm 2 Clustered Composition Assignment

- 1: Identify the set of gateway drones $G = \{g_1, g_2, \dots, g_\ell\}$.
- 2: Determine the number of clusters

$$k = \max\left(|G|, \left\lceil \frac{|U|\lambda}{\mu} \right\rceil\right),$$

where $|U|$ is the total number of devices, λ is the per-device arrival rate, μ is the per-drone service rate.

- 3: Apply k -means to the spatial positions of non-gateway drones, yielding clusters $\{C_1, C_2, \dots, C_k\}$ with centroids $\{c_1, \dots, c_k\}$.
- 4: **for** each cluster C with centroid c **do**
- 5: **Gateway selection:** choose the gateway for this cluster as

$$g^* = \arg \min_{g \in G} (\alpha \cdot \text{dist}(c, g) + (1 - \alpha) \cdot \text{load}(g)).$$

- 6: **Head selection:** choose the cluster head as

$$d^* = \arg \min_{d \in C} (\alpha \cdot \text{dist}(d, g^*) + (1 - \alpha) \cdot \text{load}(d)).$$

- 7: Connect all members of C to d^* , and connect d^* directly to g^* in FIFO order.
 - 8: **end for**
-

reintroduce gateway congestion.

3) *Parallel Composition:* Parallel composition establishes multiple disjoint relay chains from entry drones to gateways. We assume that each drone forwards to only one successor and does not split its traffic. Under this assumption, parallelism arises from the swarm as a whole: different groups of drones form separate chains operating concurrently. A long sequential relay chain is therefore a special case of parallel composition, where only one path is active.




Fig. 6. Parallel composition: traffic is distributed across multiple relay paths.

Algorithm 3 proceeds in four stages. First, it computes the required number of parallel paths from the total demand, respecting drone capacities, while enforcing a lower bound equal to the number of gateways [Lines 1–2]. Second, each gateway initializes one path by attaching the entry drone that minimizes a weighted cost combining proximity and current load, identical to the cost structure used in the clustered strategy. These gateway-anchored choices form the initial set of paths [Lines 3–7]. If the required number of paths exceeds the number of gateways, the algorithm selects additional starting points.

Algorithm 3 Parallel Composition Assignment

- 1: Compute required number of paths:

$$k = \max\left(|G|, \left\lceil \frac{|U|\lambda}{\mu} \right\rceil\right),$$

where $|U|$ is the total number of devices, λ is the per-device arrival rate, μ is the per-drone service rate.

- 2: Initialize an empty set of paths P^* and mark all drones as unassigned.

(Gateway-anchored initialization)

- 3: **for** each gateway $g \in G$ **do**
- 4: Select entry drone

$$d^* = \arg \min_{d \in D_{\text{entry}}} \alpha \cdot \text{dist}(d, g) + (1 - \alpha) \cdot \text{load}(d),$$

over unassigned drones.

- 5: Start a new path $p = [d^*, g]$, add it to P^* .
- 6: Mark d^* as assigned.
- 7: **end for**

(Additional path creation if $k > |G|$)

- 8: **if** $k > |G|$ **then**
- 9: Build list \mathcal{L} of all pairs (d, g) where d is unassigned:

$$\text{cost}(d, g) = \alpha \cdot \text{dist}(d, g) + (1 - \alpha) \cdot \text{load}(d).$$

- 10: Sort \mathcal{L} by cost ascending.
- 11: **for** the top $(k - |G|)$ pairs in \mathcal{L} **do**
- 12: Create new path $p = [d, g]$ and add it to P^* .
- 13: Mark d as assigned.
- 14: **end for**
- 15: **end if**

(Assign remaining drones to paths)

- 16: **while** there exists an unassigned drone **do**
- 17: For each unassigned drone d and each path $p \in P^*$ compute:

$$\text{cost}(d, p) = \alpha \cdot \text{dist}(d, p) + (1 - \alpha) \cdot \text{load}(p),$$

where $\text{load}(p)$ is updated dynamically as members join.

- 18: Select the pair (d^*, p^*) with minimum cost.
 - 19: Append d^* to the end of path p^* .
 - 20: Update $\text{load}(p^*)$.
 - 21: Mark d^* as assigned.
 - 22: **end while**
-

It evaluates every unassigned drone against every gateway, computes the same weighted cost, sorts all drone–gateway pairs, and selects the best pairs until the target number of paths is reached [Lines 8–15]. Finally, all remaining drones are assigned to the existing paths. Assignment is performed iteratively and load-aware: at each step, the algorithm connects the drone–path pair that yields the lowest updated cost, where the path’s load increases dynamically as drones join it. This ensures that later decisions account for the actual, evolving utilization of each path [Lines 16–22].

Tradeoffs. Parallel composition enhances performance by

distributing traffic across multiple relay chains, reducing congestion on any single path. The number of parallel paths plays a central role: too few paths concentrate load and increase end-to-end delay, while too many paths fragment traffic and raise coordination and control complexity.

B. Queuing-Theory Based Heuristic for Composition Selection

The core challenge in SNaaS composition is selecting drone orchestrations that meet latency and stability guarantees under dynamic traffic and deployment conditions. To address this, we use an M/G/1 priority queueing framework, which captures congestion-dominated multi-hop aerial forwarding, supports heterogeneous drones and traffic classes, and enables end-to-end delay evaluation by aggregating per-hop delays [39]. The framework maps real-time swarm telemetry to queueing parameters to analytically assess candidate compositions.

a) *Modeling assumptions:* Each drone d_j (entry, relay, or gateway) is modeled as an M/G/1 queue with two traffic classes:

- **High-priority (control) packets:** arrival rate $\lambda_j^{(c)}$, service-time distribution with mean $\bar{X}_j^{(c)}$ and second moment $\mathbb{E}[(X_j^{(c)})^2]$.
- **Low-priority (data) packets:** arrival rate $\lambda_j^{(d)}$, service-time distribution with mean $\bar{X}_j^{(d)}$ and second moment $\mathbb{E}[(X_j^{(d)})^2]$.

Let $\lambda_j = \lambda_j^{(c)} + \lambda_j^{(d)}$ be the total arrival rate. The total load at drone d_j is

$$\rho_j = \lambda_j^{(c)} \bar{X}_j^{(c)} + \lambda_j^{(d)} \bar{X}_j^{(d)}, \quad \text{stability requires } \rho_j < 1.$$

The server uses *non-preemptive priority*, with control packets always served ahead of data packets.

b) *Per-node delay:* Define for node d_j the residual service time

$$R_j = \frac{\lambda_j^{(c)} \mathbb{E}[(X_j^{(c)})^2] + \lambda_j^{(d)} \mathbb{E}[(X_j^{(d)})^2]}{2(1 - \rho_j)}.$$

High-priority (control) delay. Control packets wait only for residual work in service:

$$W_j^{(c)} = \frac{R_j}{1 - \rho_j^{(c)}}, \quad D_j^{(c)} = W_j^{(c)} + \bar{X}_j^{(c)},$$

where

$$\rho_j^{(c)} = \lambda_j^{(c)} \bar{X}_j^{(c)}.$$

Low-priority (data) delay. Data packets wait for: (i) residual time, (ii) all queued control packets, and (iii) control arriving during the waiting period:

$$W_j^{(d)} = \frac{R_j}{(1 - \rho_j^{(c)})(1 - \rho_j)}, \quad D_j^{(d)} = W_j^{(d)} + \bar{X}_j^{(d)}.$$

c) *Mapping strategies to queueing networks:* Each composition induces a routing pattern and therefore a set of arrival rates $\{\lambda_j^{(c)}, \lambda_j^{(d)}\}$ at each drone:

- *Direct:* multiple entry drones send control and data directly to gateways (parallel M/G/1 servers).
- *Clustered:* entries \rightarrow cluster heads \rightarrow gateways, creating a two-layer priority queueing network.
- *Parallel:* several disjoint relay chains toward gateways; each path is a series of M/G/1 priority queues.

d) *End-to-end latency:* Let π be a path from an entry drone to a gateway:

$$\pi = (d_s, d_1, \dots, d_k, g).$$

The end-to-end latency for traffic class $x \in \{c, d\}$ is

$$L^{(x)}(\pi) = \sum_{\ell \in \pi} D_\ell^{(x)}.$$

Given a composition c with path set Π_c and traffic fractions ω_π :

$$L_{\text{avg}}(c) = \sum_{\pi \in \Pi_c} \omega_\pi L^{(d)}(\pi), \quad L_{\text{max}}(c) = \max_{\pi \in \Pi_c} L^{(d)}(\pi).$$

Mission-critical SLAs typically evaluate worst-case latency; consumer-grade tasks may use average latency.

e) *Decision rule:* Given measured $\lambda_j^{(c)}, \lambda_j^{(d)}$ and service distributions, select the composition

$$\begin{aligned} c^* &= \arg \min_{c \in \mathcal{C}} L_{\text{SLA}}(c) \\ \text{s.t. } &\rho_j < 1, \forall j, \\ &|U_j| \leq C_d \text{ (coverage/capacity constraint)}. \end{aligned}$$

Algorithm 4 describes the main steps in the selection process. [Line 1] collects the current swarm telemetry (per-node control and data arrival rates and service-time statistics) and enumerates the set of candidate compositions \mathcal{C} (e.g., direct, clustered, and parallel variants instantiated for the current topology). For each candidate composition $c \in \mathcal{C}$ [Lines 2–7], [Line 3] derives the resulting routing pattern by mapping device traffic to end-to-end paths Π_c and computing the induced arrival rates at every drone along those paths. Given these per-node rates, Line 4 evaluates the priority M/G/1 queueing model at each node, computing the residual service time R_j and the class-specific delays $D_j^{(c)}$ and $D_j^{(d)}$. [Line 5] then aggregates these per-node delays into an end-to-end latency metric $L_{\text{avg}}(c)$ or $L_{\text{max}}(c)$, depending on whether the SLA is defined in terms of average or worst-case latency. In [Line 6], any composition whose utilization violates the stability condition $\rho_j < 1$ at any node is marked infeasible and excluded from further consideration. Finally, [Line 8] selects the composition c^* that minimizes the SLA-relevant latency metric $L_{\text{SLA}}(c)$ among all feasible candidates, and returns it as the chosen orchestration plan.

Algorithm 4 Priority M/G/1 Guided Strategy Selection

- 1: Collect telemetry: arrival rates $\lambda_j^{(c)}, \lambda_j^{(d)}$, service statistics, and enumerate \mathcal{C} .
 - 2: **for** each composition $c \in \mathcal{C}$ **do**
 - 3: Derive routing $\Rightarrow \Pi_c$ and per-node arrival rates.
 - 4: Compute $R_j, D_j^{(c)}, D_j^{(d)}$ for all nodes.
 - 5: Evaluate $L_{\text{avg}}(c)$ or $L_{\text{max}}(c)$ depending on SLA.
 - 6: Mark c infeasible if $\rho_j \geq 1$ for any j .
 - 7: **end for**
 - 8: Return $c^* = \arg \min_{c \in \mathcal{C}} L_{\text{SLA}}(c)$.
-

C. SLA Compliance and Stability Enforcement

If queueing analysis shows that no candidate composition satisfies queue stability ($\rho_j < 1$) or meets the SLA latency bound, the controller activates a stability enforcement (SE) stage, since both conditions are mandatory contractual requirements. The SE module iteratively applies corrective actions, recomputing utilizations and end-to-end latency after each step, and halts as soon as all queues satisfy $\rho_j \leq \rho_{\text{max}} < 1$ and latency compliance is restored, thereby minimizing reconfiguration overhead while ensuring stable and SLA-compliant service delivery.

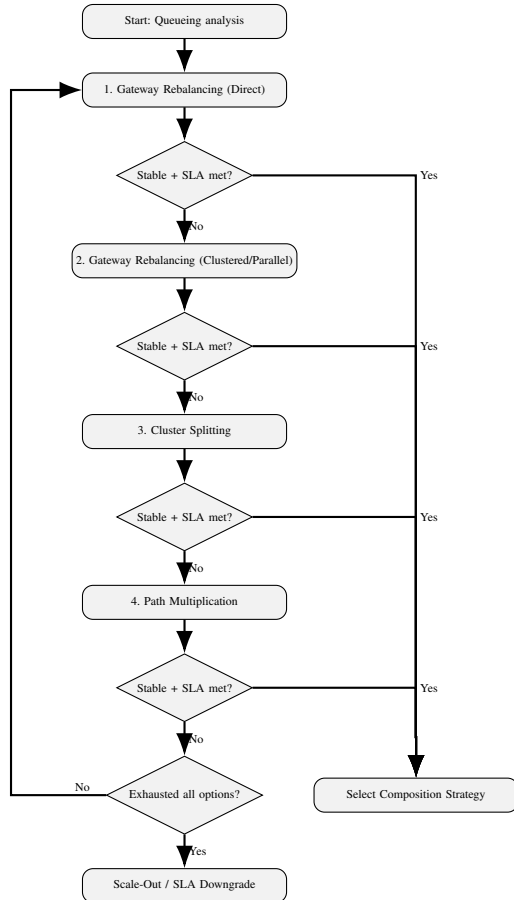


Fig. 7. Flowchart of the SLA Compliance and Stability Enforcement procedure.

Procedure: The enforcement logic, summarized in Fig. 7, proceeds through four ordered edits, where each edit is applied once in sequence; if SLA compliance is not achieved, the procedure then iteratively cycles through the same edits until all corrective options are exhausted.

1) **Gateway rebalancing (direct).** The controller first applies direct composition adjustments, as they are minimally disruptive and affect only routing, not the swarm structure. Entry drones are reassigned to alternate gateways by adapting the weighting factor α_g to discourage heavily loaded gateways and favor lighter ones. This allows connections to farther gateways when needed. If the resulting mapping satisfies gateway stability but still violates the SLA, or remains unstable, the controller escalates to the next enforcement step.

2) **Gateway rebalancing (clustered/parallel).** This step redistributes aggregate traffic by reassigning cluster heads or relay chains to alternate gateways when direct reassignment is insufficient. The controller adapts the gateway weighting factor α_g to penalize saturated gateways and favor underloaded ones while preserving internal composition. If the resulting mapping satisfies gateway stability and SLA latency constraints, enforcement stops; otherwise, it proceeds to the next step.

3) **Cluster splitting.** If gateway rebalancing fails, the controller splits overloaded clusters by dividing any cluster with $\lambda_h^{\text{agg}} > \mu_h$ into two subclusters to reduce per-head utilization. The split increases the number of clusters ($k \rightarrow k + 1$) using the same k-means spatial and load-aware criteria as in composition, followed by local head reselection. If recomputed stability or latency constraints are still violated, the controller proceeds to the next step.

4) **Path multiplication.** Cluster splitting is applied before path multiplication, as it addresses overload at the aggregation level, while path multiplication restructures relay paths. If congestion persists, the controller splits highly utilized relay chains and recomposes them into parallel paths, dividing traffic to reduce per-relay load and improve stability. The updated paths are then evaluated for queue stability and end-to-end latency, with further enforcement applied if violations remain.

Scale-out or SLA downgrade. If neither stability nor the SLA latency target can be achieved after all preceding corrections, the controller initiates a scale-out action or triggers SLA renegotiation. In a scale-out action, new drones or gateways are provisioned to expand service capacity and restore compliance. If resource addition is infeasible, the controller requests an SLA downgrade, i.e., loosening the latency bound or reducing the traffic demand, to reestablish a feasible operating region under the current infrastructure.

VI. EXPERIMENTS AND RESULTS

This section evaluates the proposed framework implemented on Google Colab (Python) using a high-RAM CPU runtime (≈ 50 GB RAM) on a Google Compute Engine backend.

A. Dataset

Experiments employ the AERPAW UAV-based Signal Dataset [40], which provides real-world air-to-ground (A2G)

received-signal strength (RSS) measurements at a 3.3 GHz carrier frequency for three UAV altitudes (40 m, 70 m, 100 m). The dataset was collected at five fixed ground nodes, supplying spatial diversity in the measured link quality. For each altitude, RSS samples collected across all ground nodes and time were aggregated to compute a single mean RSS value. During simulation, channel bandwidth was treated as a control parameter and swept between 5–20 MHz. For each altitude–bandwidth configuration, signal-to-noise ratios (SNRs) were computed using a bandwidth-dependent thermal noise model, and per-UAV service rates (μ) were deterministically estimated using a Shannon-capacity-based throughput proxy scaled by a PHY efficiency factor ($\eta = 0.6$). These μ values were used as fixed service parameters in the M/G/1 priority queues.

B. Experimental Parameters

Unless stated otherwise, simulations assume a transmit power of $P_t = 20$ dBm, a data packet size of 1 kB, a control packet size of 32 B, a carrier frequency of 3.3 GHz, and a duration of 120 s. Thermal noise power [41] is computed as:

$$N = -174 + 10 \log_{10}(B) + NF,$$

where B denotes the channel bandwidth and NF is the receiver noise figure [41]. We fix the UAV altitude at 40 m and the channel bandwidth at 5 MHz.

C. Request Generation

A total of 100 synthetic consumer requests are generated per evaluation bin, each specifying an SLA with average latency L_{SLA} and stability threshold $\rho < 1$. Each request also defines the number of user devices to be served and their allocation to preassigned entry drones, respecting each drone’s capacity C_d . User devices are uniformly randomly distributed within a $100\text{m} \times 100\text{m}$ service area, while UAV positions are also randomized. To preserve controlled comparisons across operating points, results are grouped into bins along the x-axis (e.g., SLA targets or number of devices), and the 100 requests within each bin are varied by a small perturbation of $\pm 10\%$ around the bin’s nominal value. This captures realistic variations in user density and load distribution across the swarm under consistent spatial constraints. The SNaaS controller dynamically composes UAV services to satisfy these requirements. Per-user traffic arrivals are modeled as a Poisson process, reflecting independent user activity and aligning with standard assumptions in M/G/1 service models.

D. Scaling and Topology

Real AERPAW link distributions are used to parameterize the simulations by mapping link quality to per-link service rates. We assume interference is mitigated through standard wireless resource management techniques (e.g., TDMA/FDMA-based orthogonal channel allocation, power control, frequency reuse, and directional antennas), and detailed interference-aware optimization in SNaaS is left for future work. We evaluate three scales representing disaster-response scenarios:

- **Small:** 10 entry UAVs, 3 gateway
- **Medium:** 25 entry UAVs, 5 gateways
- **Large:** 100 entry UAVs, 10 gateways

E. Experimental Evaluation

a) **Experiment 1: Violation Rate vs. SLA Latency:** We study the impact of SLA strictness on system compliance by measuring the violation rate as a function of the requested latency bound (Fig. 8). For each requested SLA latency, 100 requests are generated on a small-scale swarm. We compare our proposed queueing-theory-based enforcement against two brute-force baselines: brute-force Direct and brute-force Clustered compositions, each enumerating all feasible combinations up to a cap of 200 candidates; a brute-force Parallel baseline is omitted due to its prohibitive computational cost, and parallel composition is not generally useful for small-scale swarms, as further demonstrated in subsequent experiments. The results show that relaxing the SLA (higher allowable latency) consistently reduces violations for all methods. Brute-force Direct exhibits the highest violation rate, as directly connected devices overload gateways and create queueing bottlenecks. In contrast, brute-force Clustered and our approach achieve nearly identical and significantly lower violation rates, indicating that clustered composition is the dominant structure at this scale and that our method successfully selects the same solutions identified by brute force. Importantly, as shown in the accompanying runtime figure (Fig. 9), our approach reaches this performance with substantially lower computation time, highlighting its efficiency advantage.




Fig. 8. Violation rate versus requested SLA latency for a small-scale swarm.

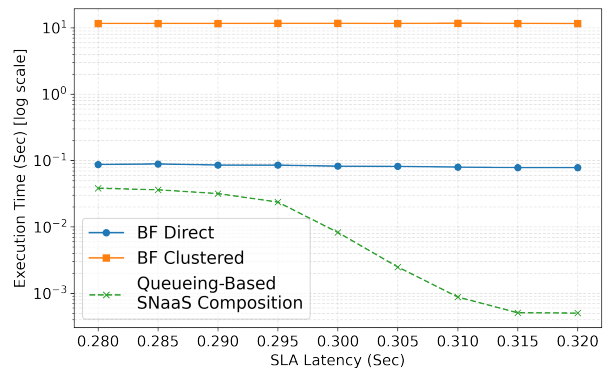


Fig. 9. Average execution time per request versus SLA latency (log scale).

b) Experiment 2: Latency vs. Number of Devices:

This experiment highlights the importance of the stability enforcement step by reporting the average latency as the number of connected devices increases in a medium-scale swarm (Fig. 10). We fix the requested SLA latency and progressively increase the offered load by adding devices, averaging results over multiple Monte Carlo runs. We compare direct, clustered, and parallel compositions without enforcement against the proposed framework with enforcement enabled. Minor local fluctuations in the curves may appear due to the controlled randomness in request generation as explained earlier, but the overall trend remains the key observation. The results show that parallel composition without enforcement achieves the highest latency; although additional hops distribute load and maintain stability as the number of devices increases, this comes at the cost of significantly higher end-to-end delay. In contrast, direct and clustered compositions without enforcement become unstable early as device count grows, due to excessive load on gateways or cluster heads; clustered fails earlier than direct because the number of clusters are not optimized as the weight for the cost function is not optimized ($\alpha = 0.5$). Our proposed framework consistently outperforms all baselines by dynamically switching between composition strategies as load increases while enforcing stability constraints, and by fine-tuning the cost function weight (α) to achieve the lowest average latency. This adaptive behavior and strategy selection are further analyzed in the following experiment.

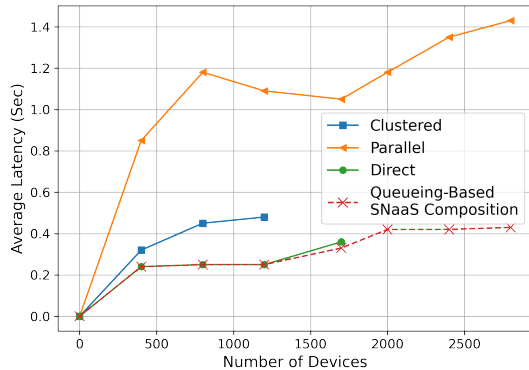


Fig. 10. Average latency versus number of devices for a medium-scale swarm.

c) Experiment 3: Strategy Selection Behavior Under Increasing Load:

This experiment analyzes how the SLA-aware enforcement mechanism selects composition strategies as device load increases (Fig. 11). At low load, direct composition dominates due to minimal latency. As load grows and stability degrades, the framework shifts to clustered composition, and at high load increasingly selects parallel composition to distribute traffic. This progression demonstrates adaptive strategy selection that balances latency and stability to satisfy SLA constraints. The small residual direct selection at high load is attributable to randomness.

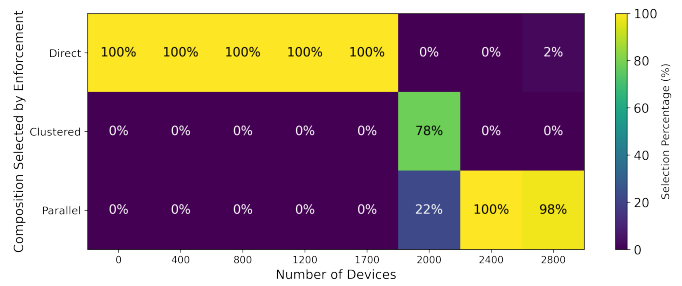


Fig. 11. Heatmap of composition strategy selection frequency versus number of devices.

d) Experiment 4: Impact of Swarm Size on Latency Under Increasing Load:

This experiment evaluates how swarm size affects scalability and latency as the number of connected devices increases under a fixed SLA latency requirement (Fig. 12). Unlike the earlier definition of small, medium, and large scenarios, where swarm scale and device counts were jointly varied, here the same device load is applied across all swarm scales to isolate the effect of swarm size alone. Using the proposed enforcement mechanism, we measure the average queuing latency while progressively increasing device counts for three swarm scales (small, medium, large), averaging over Monte Carlo trials with randomized device placements. The results show that larger swarms sustain stable operation for substantially higher device counts and achieve lower latency, since additional drones provide more service capacity and enable better load distribution across gateways and relays. This trend highlights the performance–cost tradeoff in swarm deployment: increasing the number of drones improves feasibility and reduces latency under higher device demand, but at the expense of greater deployment cost, enabling operators to determine the minimum swarm size required to satisfy a given SLA.




Fig. 12. Average latency versus number of devices under SLA enforcement for different swarm scales (small/medium/large).

VII. CONCLUSION AND FUTURE WORK

This paper introduces Swarm Network-as-a-Service (SNaaS) as a service-oriented abstraction for on-demand, SLA-driven connectivity using UAV swarms. It models drone-to-device and drone-to-drone interactions as composable

services, formalizes atomic and composite SNaas services, and presents an SOA–SDN-inspired architecture that addresses service composition under latency and stability constraints.

The paper proposes direct, clustered, and parallel composition strategies together with a queueing-theory-based selection and enforcement framework that analytically evaluates stability and end-to-end latency using a priority M/G/1 model. Experimental results based on real AERPAW data show near brute-force performance at significantly lower computational cost, smooth adaptation across load regimes, and clear performance–cost tradeoffs as swarm size scales. Future direction can extend SNaas with swarm-driven coordination, where drones adapt routes using local queue and latency feedback to enable self-organization, faster response, and resilient, SLA-compliant operation. In this setting, each drone pursues local optimization goals that collectively converge toward globally efficient and SLA-compliant network states.

ACKNOWLEDGMENTS

The work is partly funded by the King Abdullah University of Science and Technology (KAUST) Global Fellowship Program - Award No. RFS-KGFP2024-6784.

REFERENCES

- [1] H. Galal, R. Chowdhary, K. Vessali, V. Paul, G. Jaisinghania, and M. Kabbara, "Cognitive cities: A journey to intelligent urbanism," 2023.
- [2] S. Dang, O. Amin, B. Shihada, and M.-S. Alouini, "What should 6g be?" *Nature Electronics*, vol. 3, no. 1, pp. 20–29, 2020.
- [3] C. Zhang, S. Dang, M.-S. Alouini, and B. Shihada, "Big communications: Connect the unconnected," *Frontiers in Communications and Networks*, vol. 3, p. 785933, 2022.
- [4] M. Ghoshal, I. Khan, Z. J. Kong, P. Dinh, J. Meng, Y. C. Hu, and D. Koutsonikolas, "Performance of cellular networks on the wheels," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, 2023, pp. 678–695.
- [5] O. Kodheli, E. Lagunas, N. Maturo, S. K. Sharma, B. Shankar, J. F. M. Montoya, J. C. M. Duncan, D. Spano, S. Chatzinotas, S. Kisseleff, J. Querol, L. Lei, T. X. Vu, and G. Goussetis, "Satellite communications in the new space era: A survey and future challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 1, pp. 70–109, 2021.
- [6] W. Xu, Y. Sun, R. Zou, W. Liang, Q. Xia, F. Shan, T. Wang, X. Jia, and Z. Li, "Throughput maximization of uav networks," *IEEE/ACM Transactions on Networking*, vol. 30, no. 2, pp. 881–895, 2021.
- [7] S. Ammar, C. P. Lau, and B. Shihada, "An in-depth survey on virtualization technologies in 6g integrated terrestrial and non-terrestrial networks," *IEEE Open Journal of the Communications Society*, vol. 5, pp. 3690–3734, 2024.
- [8] J. Zhou, S. Dang, B. Shihada, and M.-S. Alouini, "On the outage performance of space-air-ground integrated networks in the 3d poisson field," *IEEE Transactions on Vehicular Technology*, vol. 73, no. 3, pp. 4401–4406, 2023.
- [9] S. Javaid, N. Saeed, Z. Qadir, H. Fahim, B. He, H. Song, and M. Bilal, "Communication and control in collaborative uavs: Recent advances and future trends," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 6, pp. 5719–5739, 2023.
- [10] B. Alkouz, A. Bouguettaya, and A. Lakhdari, "Density-based pruning of drone swarm services," in *2022 IEEE International Conference on Web Services (ICWS)*. IEEE, 2022, pp. 302–311.
- [11] —, "Failure-sentient composition for swarm-based drone services," in *2023 IEEE International Conference on Web Services (ICWS)*. IEEE, 2023, pp. 493–503.
- [12] 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.
- [13] M. Mozaffari, A. T. Z. Kasgari, W. Saad, M. Bennis, and M. Debbah, "Beyond 5g with uavs: Foundations of a 3d wireless cellular network," *IEEE Transactions on Wireless Communications*, vol. 18, no. 1, pp. 357–372, 2018.
- [14] X. Sun and N. Ansari, "Latency aware drone base station placement in heterogeneous networks," in *GLOBECOM 2017-2017 IEEE Global Communications Conference*. IEEE, 2017, pp. 1–6.
- [15] Y. Su, H. Zhou, Y. Deng, and M. Dohler, "Energy-efficient cellular-connected uav swarm control optimization," *IEEE Transactions on Wireless Communications*, vol. 23, no. 5, pp. 4127–4140, 2023.
- [16] L. Zhang, A. Celik, S. Dang, and B. Shihada, "Energy-efficient trajectory optimization for uav-assisted iot networks," *IEEE Transactions on Mobile Computing*, vol. 21, no. 12, pp. 4323–4337, 2021.
- [17] H. El Hammouti, D. Hamza, B. Shihada, M.-S. Alouini, and J. S. Shamma, "The optimal and the greedy: Drone association and positioning schemes for internet of uavs," *IEEE Internet of Things Journal*, vol. 8, no. 18, pp. 14 066–14 079, 2021.
- [18] H. El Hammouti, M. Benjillali, B. Shihada, and M.-S. Alouini, "Learn-as-you-fly: A distributed algorithm for joint 3d placement and user association in multi-uavs networks," *IEEE Transactions on Wireless Communications*, vol. 18, no. 12, pp. 5831–5844, 2019.
- [19] F. Wei, G. Feng, S. Qin, Y. Peng, and Y. Liu, "Hierarchical network slicing for uav-assisted wireless networks with deployment optimization," *IEEE Journal on Selected Areas in Communications*, 2024.
- [20] S. Ebrahimi, F. Bouali, and O. C. Haas, "Resource management from single-domain 5g to end-to-end 6g network slicing: A survey," *IEEE Communications Surveys & Tutorials*, vol. 26, no. 4, pp. 2836–2866, 2024.
- [21] S. Ammar, W. Abderrahim, and B. Shihada, "Maritime-oriented network slicing in o-ran integrated aerial-terrestrial networks," *IEEE Transactions on Mobile Computing*, 2025.
- [22] A. Hamdi, B. Alkouz, B. Shahzaad, A. Bouguettaya, A. G. Neiat, F. Salim, and D. Y. Kim, "Drone-as-a-service: Research challenges and directions," *Proceedings of the IEEE*, 2025.
- [23] B. Alkouz, B. Shahzaad, and A. Bouguettaya, "Service-based drone delivery," in *2021 IEEE 7th International Conference on Collaboration and Internet Computing (CIC)*. IEEE, 2021, pp. 68–76.
- [24] W. Lee, B. Alkouz, B. Shahzaad, and A. Bouguettaya, "Package delivery using autonomous drones in skyways," in *Adjunct Proceedings of the 2021 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2021 ACM International Symposium on Wearable Computers*, 2021, pp. 48–50.
- [25] W. Lee, B. Shahzaad, B. Alkouz, and A. Bouguettaya, "Reactive composition of uav delivery services in urban environments," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 10, pp. 13 453–13 466, 2024.
- [26] B. Alkouz, A. Bouguettaya, A. Lakhdari, and S. Yangui, "Signal-based approach for reliable drone swarm delivery," in *2024 IEEE International Conference on Web Services (ICWS)*. IEEE, 2024, pp. 1238–1244.
- [27] S. Javed, A. Hassan, R. Ahmad, W. Ahmed, R. Ahmed, A. Saadat, and M. Guizani, "State-of-the-art and future research challenges in uav swarms," *IEEE Internet of Things Journal*, vol. 11, no. 11, pp. 19 023–19 045, 2024.
- [28] P. Cao, L. Lei, S. Cai, G. Shen, X. Liu, X. Wang, L. Zhang, L. Zhou, and M. Guizani, "Computational intelligence algorithms for uav swarm networking and collaboration: A comprehensive survey and future directions," *IEEE Communications Surveys & Tutorials*, vol. 26, no. 4, pp. 2684–2728, 2024.
- [29] B. Alkouz and A. Bouguettaya, "Formation-based selection of drone swarm services," in *MobiQuitous 2020-17th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, 2020, pp. 386–394.
- [30] F. Venturini, F. Mason, F. Pase, F. Chiariotti, A. Testolin, A. Zanella, and M. Zorzi, "Distributed reinforcement learning for flexible and efficient uav swarm control," *IEEE Transactions on Cognitive Communications and Networking*, vol. 7, no. 3, pp. 955–969, 2021.
- [31] B. Alkouz, A. Bouguettaya, and S. Mistry, "Swarm-based drone-as-a-service (sdaas) for delivery," in *2020 IEEE International Conference on Web Services (ICWS)*. IEEE, 2020, pp. 441–448.
- [32] F. Nait-Abdesselam, A. Alsharoa, M. Y. Selim, D. Qiao, and A. E. Kamal, "Towards enabling unmanned aerial vehicles as a service for heterogeneous applications," *Journal of Communications and Networks*, vol. 23, no. 3, pp. 212–221, 2021.

- [33] O. Bekkouche, K. Samdanis, M. Bagaa, and T. Taleb, "A service-based architecture for enabling uav enhanced network services," *IEEE Network*, vol. 34, no. 4, pp. 328–335, 2020.
- [34] M. A. B. S. Abir, M. Z. Chowdhury, and Y. M. Jang, "A software-defined uav network using queueing model," *IEEE Access*, vol. 11, pp. 91 423–91 440, 2023.
- [35] A. Bouguettaya, M. Singh, M. Huhns, Q. Z. Sheng, H. Dong, Q. Yu, A. G. Neiat, S. Mistry, B. Benatallah, B. Medjahed *et al.*, "A service computing manifesto: the next 10 years," *Communications of the ACM*, vol. 60, no. 4, pp. 64–72, 2017.
- [36] W. Saad, Z. Han, T. Basar, M. Debbah, and A. Hjørungnes, "A selfish approach to coalition formation among unmanned air vehicles in wireless networks," in *2009 International Conference on Game Theory for Networks*. IEEE, 2009, pp. 259–267.
- [37] C. M. MacKenzie, K. Laskey, F. McCabe, P. F. Brown, R. Metz, and B. A. Hamilton, "Reference model for service oriented architecture 1.0," *OASIS standard*, vol. 12, no. S18, pp. 1–31, 2006.
- [38] D. Kreutz, F. M. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmoly, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2014.
- [39] D. Gross, *Fundamentals of queueing theory*. John wiley & sons, 2008.
- [40] C. Dickerson, A. H. Fahim Raouf, O. Ozdemir *et al.*, "Aerpaw uav-based signal data collected at varying altitudes and sampling rates for wireless communication studies," 2025, dataset. [Online]. Available: <https://doi.org/10.5061/dryad.2z34tmpvv>
- [41] A. Goldsmith, *Wireless communications*. Cambridge university press, 2005.