Distributed Serverless Function Scheduling in Ad-Hoc Drone Networks

Giuseppe De Palma^{a,b}, Saverio Giallorenzo^{a,b}, Alexandre Heideker^a, Matteo Trentin^{a,b,c}, Angelo Trotta^a, Gianluigi Zavattaro^{a,b}

> ^aAlma Mater Studiorum - Università di Bologna, Italy ^bINRIA, Sophia Antipolis, France ^cUniversity of Southern Denmark, Sophia Antipolis, Denmark

Abstract

The increasing use of Unmanned Aerial Vehicles (UAVs) in critical applications, such as disaster response, compels efficient communication and computation frameworks for highly dynamic ad-hoc networks. We present an interpretation of Function-as-a-Service serverless computing within the distributed settings of drone swarms to address their peculiar challenges in functionality deployment, resource allocation, and mission adaptability. We propose a novel two-layer network overlay architecture, combining a gossipbased topology management layer with a function scheduling layer to support distributed function scheduling. Our system introduces a declarative language of Ad-Hoc Allocation Priority Policies (AHAPP), tailored for adhoc drone networks, enabling flexible function deployment based on resource constraints and operational needs. The resulting combination addresses the volatility of UAVs networks by supporting execution semantics for stable- and dynamic-topology scenarios, function offloading, and resilience to network disruptions. We present experiments confirming that the features provided by our proposal support the efficient execution of serverless functions in adhoc drone networks, effectively handling their dynamic and heterogeneous nature, while achieving strong performance in terms of reliability, scheduling

Email addresses: giuseppe.depalma2@unibo.it (Giuseppe De Palma),

saverio.giallorenzo2@unibo.it (Saverio Giallorenzo), alexandre.heideker@unibo.it (Alexandre Heideker), matteo.trentin2@unibo.it (Matteo Trentin),

angelo.trotta5@unibo.it (Angelo Trotta), gianluigi.zavattaro@unibo.it (Gianluigi Zavattaro)

time, and communication overhead.

1 1. Introduction

The proliferation of unmanned aerial vehicles (UAVs), or drones, in civilian and commercial applications has sparked significant interest in developing efficient communication and computation frameworks for such decentralised, ad-hoc wireless networks, where each node participates in routing by forwarding data for others [1]. As these networks become increasingly complex, traditional computing architectures struggle to meet the demanding requirements of real-time data processing, dynamic resource allocation, and mission-critical operations [2].

Case Study and Motivating Example. The case study we present in this section is one of the possible use cases that our proposal can host. While we use the case throughout the article to explain our approach, in Section 6, we present evaluations regarding generic simulations with different topologies and number of elements, besides the ones that simulate the case study.

We illustrate an example of such highly dynamic systems with a running case study, modelling a multiphase Fire Detection and Response scenario. We represent the phases in Figure 1, and comment on them below.

As part of its responsibilities, the agency managing a natural area actively monitors various potential threats to the environment. Among these, fire detection is a critical component of their surveillance efforts. As show in panel a) of Figure 1, the drones fly across a monitored area, using cameras and sensors to conduct continuous aerial surveillance.

To ensure early detection and a swift response, the agency conducts hourly fire monitoring routines using drones, performing real-time heat signature analysis, detecting and precisely geolocating fire origins. Once a drone detects a fire, it transmits comprehensive geospatial data to the command centre, providing critical first-response information that enables rapid and targeted intervention.

Following initial fire identification, the agency commands the drones in fire locations to execute a survivor detection routine, depicted in panel b) of Figure 1. Equipped with multi-spectral imaging sensors and using machine learning algorithms, they conduct scans through smoke, debris, and challenging terrain [3].



Figure 1: Depiction of the Fire Detection and Response case study main phases: a) Fire Detection, b) Survivor Location, and c) Real-Time Emergency Response Coordination.

The last phase of the case is that of survivor tracking, shown in panel c) 34 of Figure 1. Given the information on the identified groups of survivors, the 35 agency instructs the drones in the survivors' locations to follow a continuous 36 tracking and information transmission routine [4]. Precision GPS equipment 37 and real-time video streaming enable persistent monitoring of survivor lo-38 cations. The drones can generate comprehensive data reports that include 39 exact coordinates, visual tracking information, estimated survivor conditions, 40 and recommended access routes for rescue teams. 41

Ad-Hoc Networks Serverless Computing. In contexts like the one described in 42 the above case study, Function-as-a-Service (FaaS) Serverless Computing [5] 43 emerges as a promising paradigm to address the challenges of lightweight, 44 scalable, and distributed management of computational workloads in drone 45 ad-hoc networks — we provide background information on FaaS for the un-46 familiar reader in Section 2, along with a general introduction to ad-hoc 47 networks and issues with workload distribution, drone swarms in particular. 48 FaaS is a cloud computing execution model where cloud providers auto-49 matically manage the infrastructure required to run application code. This 50 paradigm shifts from traditional server-based architectures to event-driven, 51 stateless functions that the platform runs in response to user invocations. 52 We argue that scenarios that require rapid deployment, minimal overhead, 53

and efficient resource usage like the one presented above, hardly fit into the
existing, "traditional" application deployments practices, where all the drones
already preload all functionalities that one can trigger.

Indeed, we see at least two main limitations with the preloading ap-57 proach. First, the drones have limited storage capacity, hence, loading many 58 disparate software modules would unlikely be a viable course. Second, and 59 more important, it is unrealistic that the designers of the system can fore-60 see and program all the functionalities that one can need deployed on the 61 drones. In the preload setting, missing functionalities require the drones to 62 waste precious time to go back to base, update their firmware/programming, 63 and return to the interested areas. 64

To make the above observations more concrete, let us consider our case 65 study. While one might preset the drone swarm for detecting fire alerts, this 66 specialisation would limit the capabilities of the swarm to only that kind of 67 emergency. However, emergency scenarios are unpredictable and may require 68 rapid deployment of new capabilities. For instance, if the fire happens near 69 in an industrial area, the emergency protocol could require specific chem-70 ical detection algorithms, another example is extreme weather conditions, 71 which would require specific flight patterns and sensing protocols. Focusing 72 even more specifically on the case study, e.g., the machine learning models 73 for heat signature analysis and survivor detection might need frequent up-74 dates to improve accuracy and adapt to new conditions. With preloaded 75 functions, updating these models would require taking drones out of service 76 for reprogramming, creating dangerous gaps in coverage. Moreover, from 77 an efficiency perspective, the multiphase nature of the operation (detection, 78 survivor location, tracking) means that most functions are not needed con-79 tinuously. Preloading all functions would waste valuable onboard resources. 80 Of course, considering circumstances much different from fires, like floods 81 and earthquakes, would require the drones to host even more heterogeneous 82 sets of functionalities. 83

Hence, preloading all possible functionalities for these events hardly fits
the exceptionality and time pressure of each situation.

Similarly, more sophisticated machanisms, like dynamic loading approaches, while an improvement over preloading, face important limitations when compared a distributed serverless architecture. The requirement to pre-allocate storage space for potential modules remains a constraint [6], whereas a serverless approach enables drones to share computational resources and dynamically offload functions. The update process is also cumbersome with dynamic

loading, requiring individual updates to each drone, while we propose an in-92 terpretation of serverless that can make functions available across the network 93 of valid drones. Furthermore, dynamic loading systems must still anticipate 94 which modules might be needed and include their dependencies, limiting true 95 flexibility. Serverless architectures overcome these limitations by allowing de-96 ployment of arbitrary functions without pre-planning module dependencies. 97 Moreover, most critically, in centralised approaches, when network partitions 98 occur drones that cannot communicate with the central orchestrator cannot 99 obtain the new functionalities. In contrast, our distributed serverless ar-100 chitecture makes function scheduling available through peer-to-peer sharing 101 even when disconnected from a central infrastructure, making it more re-102 silient for ad-hoc drone networks operating in challenging environments 103 indeed, the inherent scalability of the serverless approach facilitates the dy-104 namic addition/removal of network nodes, a common requirement in drone 105 swarm operations. 106

The idea of marrying ad-hoc drone networks and serverless comes from 107 recent advancements in serverless technologies, which allows the efficient ex-108 ecution of both the functions on edge devices and, more importantly for 109 the ad-hoc context, the components of the serverless platform itself, like the 110 scheduler of the functions. For instance, FunLess [7, 8] is a recent serverless 111 platform tailored for private edge cloud systems that leverages WebAssembly¹ 112 (Wasm) as its runtime environment. This feature offers several advantages, 113 among which portability and consistent development and deployment across 114 heterogeneous devices (as found in ad-hoc networks) and a reduced memory 115 footprint. 116

While technologies like FunLess provide support for building a serverless platform for ad-hoc drone networks, its implementation presents peculiar challenges that we address in this article.

Specifically, we look at both the architecture of the serverless platform and the protocols that support its functionalities in an environment characterised by the latency and intermittent connectivity of mobile aerial communication. In doing so, we strive to provide reliable function execution and enhance resource allocation across heterogeneous drone configurations.

¹²⁵ Supporting Ad-Hoc Serverless Computing. In Section 3, we present our ap-¹²⁶ proach, hinged on a two-layer network overlay system designed for drone

¹https://webassembly.org/.

The architecture we propose uses a "gossiping" layer, built on networks. 127 top of the physical plane of drones and base station, to manage the net-128 work topology and drone status tracking. The gossiping layer plays a crucial 129 role in disseminating information about drone status, available resources, 130 and network topology. On top of this layer, we find a "function scheduling" 131 one, which handles function placement and routing. The system implements 132 both *hierarchical* and *progressive* scheduling approaches, which capture the 133 high volatility of ad-hoc drone networks. Specifically, under the hierarchi-134 cal approach, the base station maintains complete network information and 135 directly assigns functions, useful in more static situations, with broadband, 136 low-latency communication. In the progressive approach, the base station 137 and the drones collaborate in making routing decisions to identify the vi-138 able targets for function execution, based on local information — this sec-139 ond approach fits scenarios with higher latencies and narrower bandwidth. 140 which prevent the base station from having an up-to-date and comprehensive 141 knowledge of the status of the network. 142

Moreover, while traditional serverless settings consider a 1:1 ratio between 143 a user invocation and the execution of the related function — logically, the 144 invocations existentially quantify the functions — we note that, in ad-hoc 145 network settings, the user might require the execution of a function on all 146 viable nodes (e.g., the fire detection routine found in our case study). For this 147 reason, we introduce a *universal quantifier* that allows users to indicate the 148 execution of the same function on all valid drones. Recognising that the fluid 149 nature of ad-hoc drone networks presents challenges in assuring the correct 150 execution of functions — invocations could more likely fail than classical 151 serverless deployment due to network latencies and loss of connectivity — 152 we also introduce two types of execution semantics for functions, a "strong" 153 and a "weak" one, with their related interpretation of the above quantifiers. 154 Briefly, the strong semantics requires the exact execution of a function on 155 either one specific drone or all viable drones, while the weak semantics offers 156 more flexibility by allowing execution on at least one drone (so the function 157 can run on a set amount of alternative nodes, while the user receives the 158 successful response from the drones) or *almost all* drones (setting a timeout 150 after which the user receives all provided responses). Since drones have 160 limited hardware, we also introduce the possibility of offloading the execution 161 of functions to external computation nodes, like cloud and edge nodes, when 162 necessary and indicated by the user. 163

Governing Decentralised Function Scheduling. A hallmark of classical server-164 less computing is that all nodes that make up the computation units available 165 to the platform are functionally equivalent, which allows the users to ignore 166 the underlying topology and leave to the discretion of the platform the task of 167 deciding where to allocate function instances. While all computation nodes 168 are functionally equivalent, performance-wise, they might differ depending 169 on e.g., their geographical position or their status. For instance, effects like 170 $code \ locality \ [9] \ -$ due to latencies in loading function code and runtimes 171 or session locality [9] — due to the need to authenticate and open new 172 sessions to interact with other services — can substantially increase the run 173 time of functions. This observation sparked extensive research into strate-174 gies [10, 11, 12, 13, 14, 15] to enhance function execution efficiency. However, 175 using a single, universal scheduling policy is inadequate for addressing diverse 176 requirements, which led to the development of Allocation Priority Policies 177 (APP) [16], a declarative language that for the definition of custom schedul-178 ing policies, so that the serverless platform can support multiple scheduling 179 approaches simultaneously, each tailored for specific function groups. 180

In the UAVs context, nodes are not only performance- but also functionallywise different, e.g., if we need to run a video scan of a specific area, only the drones equipping a camera and hovering that area are valid execution nodes.

For this reason, we introduce and integrate in our platform a new lan-184 guage, inspired by APP for Ad-Hoc Networks, called AHAPP (Ad-Hoc APP), 185 which we present in Section 4. We deem AHAPP a new language because it 186 mainly takes inspiration from APP regarding the definition of policies. How-187 ever, in APP, the serverless scheduler holds a copy of a script that describes 188 the scheduling of all functions; on the contrary, since in our platform all 189 nodes can act as a scheduler, we pair a function execution invocation with 190 its AHAPP script that describes its scheduling logic — i.e., the function 191 "travels" with its scheduling script. 192

At its core, AHAPP allows users to specify whether functions should 193 run on all drones or just one — with variations for the weak semantics. 194 The language specifies resource-based targeting through its "attributes" sys-195 tem, which handles both physical properties, like geographical positions, and 196 hardware requirements, such as available energy, memory, and storage but 197 also access to rotors, cameras, or GPUs. AHAPP supports workflow man-198 agement through the expression of *affinity* among functions (inspired by a 199 recent dialect of APP [17]), allowing functions to specify dependencies on 200 other functions that must (have) run on the same drone. Since more than 201

²⁰² one drone can qualify for executing the function, AHAPP adopts, from APP, ²⁰³ the concept of selection *strategy*, where the user can specify the priorities ²⁰⁴ for choosing among alternative nodes. This feature allows users to refine ²⁰⁵ scheduling based on factors like geographical proximity, load balancing, and ²⁰⁶ energy reserves.

In Section 5, we delve into the details of the network protocol underlying 207 our system. This protocol is essential for maintaining an accurate and up-to-208 date view of both the network topology and the status of individual drones 209 within the distributed system and support the scheduling of functions. Rely-210 ing on periodic message exchanges, the protocol addresses key challenges in 211 mobile ad-hoc networks, particularly within drone swarms, where dynamic 212 topologies, intermittent connectivity, and diverse hardware capabilities re-213 quire robust distributed communication mechanisms. 214

Following the above design, we devise experiments to validate our ap-215 proach, presenting them in Section 6. To assess the performance of dis-216 tributed function deployment algorithms, we conduct extensive simulations 217 across various application scenarios. Our evaluation highlights how the pro-218 posed decentralised scheme improves upon centralised scheduling models, 219 maintaining a high service deployment ratio close to the optimal value while 220 reducing reliance on global network state knowledge. Moreover, the interac-221 tion between knowledge dissemination, scheduling decisions, and deployment 222 success will be carefully analysed, as stale or incomplete network information 223 can significantly impact function placement and execution efficiency. The ex-224 periments confirm that the functionalities provided by our proposal support 225 the efficient execution of serverless functions in ad-hoc drone networks, effec-226 tively adapting to their dynamic and heterogeneous nature. Our approach 227 demonstrates strong performance in terms of reliability, scheduling time, and 228 communication overhead, showing that decentralised function scheduling is 229 both scalable and resilient under varying network conditions. 230

We conclude, in Section 8, by drawing final remarks and discussing future development.

233 2. Background

Before presenting our approach, in Section 3, we provide context for the key technologies and concepts underlying this research: serverless computing and ad-hoc (drone) networks.

237 2.1. Background on Serverless

Serverless computing has emerged as a response to the growing complexity 238 and management challenges in modern cloud applications [5]. The core of the 239 Serverless model is a Function-as-a-Service (FaaS) platform, which enables 240 developers to build applications as compositions of stateless, event-driven 241 functions that automatically scale according to demand. This model proves 242 particularly cost-effective for applications with variable traffic patterns, as 243 developers pay only for actual function execution time rather than maintain-244 ing constantly running servers. The key characteristic of these platforms is 245 their infrastructure transparency. The underlying resources that power func-246 tion execution (concretely, e.g., the cluster of nodes) are abstracted away, 247 and developers need to only upload functions in any (supported) program-248 ming language and connect them to a triggering event, e.g., HTTP requests. 249 When triggered by such events, the platform executes the relative function by 250 selecting one of the available resources designed for execution and transmit-251 ting the invocation request to it. The function is then executed in ephemeral 252 environments that can be instantiated and removed on demand. These en-253 vironments, which may be implemented using virtual machines, containers, 254 or specialised runtimes, handle all aspects of resource allocation and exe-255 cution. Our research leverages this infrastructure transparency property of 256 FaaS platforms in a novel context: ad-hoc networks. 257

258 2.2. Background on Ad-Hoc Networks

Ad-hoc networks enable decentralised communication without fixed in-259 frastructure, making them suitable for dynamic environments, such as dis-260 aster response, military operations, and remote sensing [18]. Unlike conven-261 tional networks, these systems continuously adapt to changing conditions, 262 requiring robust mechanisms for routing, resource allocation, and fault toler-263 ance. Aerial ad-hoc networks, composed of drones, introduce additional com-264 plexities due to their three-dimensional mobility, limited energy resources, 265 and highly dynamic topology [1]. Unlike ground-based networks, drone-based 266 systems must continuously adapt to environmental conditions, network dis-267 ruptions, and intermittent connectivity. 268

Traditional ad-hoc drone networks rely on pre-programmed behaviours and functions loaded at a base station, requiring drones to return for reprogramming when new functionality is needed [19]. This approach proves especially problematic in emergency scenarios like disaster response, where returning to base interrupts critical operations and wastes valuable response time. Each round trip consumes significant battery power that could otherwise support mission-critical tasks, while the travel time reduces effective coverage and delays deploying new capabilities [20].

A more efficient approach would enable remote reprogramming by sending 277 new functions directly to drones in the field, maintaining continuous opera-278 tion, preserving battery life, and allowing rapid adaptation to changing mis-279 sion requirements without physical recall. This remote function deployment, 280 which aligns with our interpretation of serverless for ad-hoc networks, also 281 enables dynamic load balancing and task redistribution across the network. 282 critical for maintaining service when drones may fail or become unreachable 283 in emergency situations. 284

However, bringing serverless computing to ad-hoc networks presents pe-285 culiar challenges. Traditional cloud serverless platforms rely on centralised 286 schedulers found in complete topologies, i.e., the centralised scheduler can 287 directly reach all nodes, and with complete system visibility, enabling pre-288 cise, globally-controlled function allocation and resource management. In 289 contrast, it is difficult to efficiently impose global coordination on ad-hoc net-290 works, which frequently take (e.g., routing) decisions relying on (bandwidth-291 efficient but partial) local, aggregated knowledge. 292

Hence, implementing an ad-hoc network version of serverless requires 293 scheduling to be lightweight, adaptive, and resilient to outdated or partial in-294 formation. In particular, when dealing with drones, the resource constraints 295 impose strict limits on function execution, demanding efficient scheduling 296 routines that optimise energy and computational costs. Knowledge dissem-297 ination further complicates function deployment. In a cloud environment, 298 function scheduling benefits from stable, high-bandwidth communication. In 299 an ad-hoc network, global state propagation is impractical due to bandwidth 300 and energy constraints, requiring decentralised approaches such as gossip-301 based information exchange. Function execution also faces unpredictability, 302 as network disruptions can lead to execution failures, necessitating adaptive 303 mechanisms. 304

Despite these challenges, serverless computing aligns well with the eventdriven nature of drone networks. Modern, lightweight serverless execution environments, e.g., based on Wasm [8, 7], offer the possibility to both run functions and the serverless platform services across heterogeneous devices. By integrating decentralised scheduling, adaptive knowledge dissemination, and fault-tolerant execution, serverless computing can enhance the flexibility and efficiency of aerial ad-hoc networks. This work introduces a decentralised ³¹² function scheduling framework tailored to these constraints.

313 3. An Approach for Ad-hoc Serverless Computing

The approach that we propose consists of a two-layer network overlay among the drones and the Base Station, built on top of the physical plane, where the connections among the nodes in the network depend on the range and strength of the wireless signal and position of the drones w.r.t. each other and the Base Station.

In this section, we present the design and principles behind our approach. We concretely evaluate our proposal, e.g., the application of the different scheduling modalities, in Section 6.

We draw a schema of the overlay network we propose in Figure 2. In 322 the figure, we depict, at the top, the physical plane, i.e., that of the physical 323 drones and Base Station, under which we find the "gossiping" layer, dedicated 324 to keeping track of the topology of the ad-hoc network, as well as the sta-325 tus of each drone. Specifically, this kind of information includes knowledge 326 about the neighbours of each element (i.e., immediately reachable nodes) for 327 each drone, as well as status information such as its available memory and 328 storage. The drones spread information about their and their neighbours' 329 status, e.g., whether their camera or GPU are available for use, and their 330 current workload. In the figure, we denote these exchanges as $g(D_1, \dots, D_n)$ 331 to represent a gossip g about the drones with identifiers D_1, \dots, D_n — to il-332 lustrate the key mechanics of our approach protocol without too much detail. 333 in the figure, we sketch a simplified sequence of gossiping exchanges among 334 the drones and describe the protocol in full detail in Section 5. Gossip-335 based protocols are widely used in mobile ad-hoc networks thanks to their 336 ability to minimise network congestion and provide robust, scalable commu-337 nication. Unlike explicit point-to-point communications, which often require 338 extensive control messages to maintain routing and scheduling state, gossip-339 ing relies on periodic, neighbour-only exchanges of compact status updates. 340 Each node transmits its local state to immediate neighbours, who in turn dis-341 seminate this information further, enabling the network to converge toward 342 a consistent, global view over time. This decentralised mechanism eliminates 343 the need for coordinated broadcasts or central controllers, making it highly 344 suitable for dynamic and infrastructure-less environments. As demonstrated 345 by Shah [21], gossiping effectively reduces redundant message transmissions, 346 lowering the likelihood of network congestion. More recent work by Fang 347



Figure 2: An example of the proposed overlay network.

et al. [22] reinforced this observation, showing that asynchronous gossipbased algorithms enable scalable, low-overhead coordination in distributed optimisation problems. Thus, in the context of ad-hoc networks serverless computing, gossiping can provide a lightweight and fault-tolerant substrate for disseminating state information, enabling each node to make informed scheduling decisions without relying on global knowledge or infrastructure.

In Figure 2, we find, at the bottom, the "function scheduling" layer, dedi-354 cated to the placement of functions and concerning the messages on the rout-355 ing and allocation of functions, along with complementary messages such as 356 the response to the user after the execution of a function. In the Figure, 357 we represent these exchanges with the message (f, a, D_1) , which we read as 358 "deliver function f on drone D_1 following the scheduling policy a". As an 359 example, we depict in the figure an illustrative flow that, starting from the 360 Base Station, goes to drone D_4 , which decides to forward the message to 361 drone D_3 , which finally relays the message to the addressee, D_1 , that will 362 execute the function. To inform these placement and routing decisions, the 363 components at this layer use the information gathered from the gossiping one. 364

Closing the example, when the target drone (D_1) terminates the execution of the function, it sends back to the Base Station the result of the execution as a response. Similarly to the routing of the execution request message, the response message from the drone to the Base Station uses information from the gossiping layer to determine the path of the delivery. Since the responserelay part is standard for ad-hoc networks, we omit its description in the rest of our specification and focus on the traits distinctive of our approach.

As noticeable by the pairing of function execution requests with their scheduling policy, the function scheduling layer supports the definition of user-defined, per-function scheduling information expressed in the AHAPP language, discussed in Section 4. On the contrary, the gossiping layer runs a predefined behaviour on all nodes to disseminate status information within the network.

The schema in Figure 2 also illustrates another distinctive feature of our 378 proposal. The functions that execute on the drones are not already present 379 on the drones themselves. On the contrary, the functions "travel" with the 380 request to run them. From a programming perspective, the creator and 381 maintainer of the functionalities implemented by the functions is any user 382 who has access to the platform and can issue the request to execute them 383 at the Base Station. More practically, the user directly interacts only with 384 the Base Station, by sending application-level messages, e.g., using REST 385 or other protocols (the user-Base Station communication is an orthogonal 386 detail to our proposal) and they do not directly interact with the function 387 scheduling layer. 388

The gossiping and function-scheduling layers run in parallel and without 389 coordination. At the gossiping level, the drones regularly exchange unso-390 licited messages about their status and that of their neighbourhood. At the 391 function-scheduling level, users trigger the execution of functions on drones 392 by sending requests to the Base Stations. The function-scheduling layer uses 393 the information gathered from the gossiping layer available at the moment 394 of scheduling to determine the allocation of functions on drones. Hence, the 395 gossiping layer just informs the scheduling at each node (Base Station and 396 drones), but it does not interact directly with the function-scheduling layer 397 to take decision. 398

399 3.1. Rationale behind our two-layer approach

The reasoning behind our two-layer approach comes from the fact that we want to separate the scheduling concerns from the operations for network 402 connectivity, which are automatic and invisible to the users.

The gossiping layer builds directly on the physical plane, to construct a view of all the drones, performing distributed monitoring of their status. Given the low level of abstraction of this protocol and its relative independence w.r.t. the task of scheduling functions, we avoid embedding scheduling requests information into the messages exchanged at this level.

In the design of the scheduling layer, we capture the fact that different 408 functions can require different amounts and kinds of resources, depending on 409 the nature of the task they need to accomplish. This characteristic is a strong 410 departing point between classical serverless computing and our proposal. 411 Indeed, most serverless platforms define an arbitrary, hard-coded logic for 412 the scheduling of functions. This approach assumes that all the nodes that 413 make up the cluster of available targets for the execution of functions are 414 functionally equal. We argue that, in ad-hoc networks, this assumption is 415 unrealistic, since elements such as the position of the drone and the hardware 416 it carries (e.g., some functions might require a GPU while others might need 417 camera access, etc.) can determine the successful execution of a function. 418 Therefore, we associate each function with a specific policy that informs the 419 Base Station and the intervening drones on the requirements of the function 420 under scheduling. 421

422 3.2. Hierarchical vs Progressive Scheduling

While associating a function with its scheduling policy helps users target the correct drones for executing the functions, lowering errors and increasing efficiency, we notice that, depending on the state of the network, the system can decide on which drone(s) to place a function in two main ways: a *hierarchical* and a *progressive* one.

In this section, we discuss the main trade-offs between these two scheduling modalities. We empirically explore the conjectures presented in this section in Section 6, through simulations.

In the first case, hierarchical, one can assume that the network topol-431 ogy is relatively stable — e.g., drones move slowly w.r.t. the propagation of 432 messages — so that the Base Station has up-to-date information about the 433 status of the drones to specify which drone(s) shall run the function under 434 scheduling, leaving to the other nodes that connect the target drone to the 435 Base Station (if any) the task of (efficiently) delivering this request. We rep-436 resent this case in Figure 3. In the figure, we use the notation (f, a, D) to 437 indicate a message for scheduling function f on drone D, following the policy 438



Figure 3: Example of hierarchical scheduling, scheduling layer.

a. Notice that the request sent by the platform's user — marked, in this and 439 the subsequent figures in this section, using a dashed line to indicate the 440 usage of a different communication protocol stack than the one among the 441 Base Station and the Drones — that reaches the Base Station (from the left, 442 in Figure 3) does not indicate a target. Indeed, while the user is unaware 443 of the underlying composition and configuration of the drone network, it de-444 fines properties in the companion policy useful to identify valid executors of 445 the function. As part of the hierarchical scheduling, it is the Base Station 446 that decides on which drone to schedule the function. More precisely, while 447 the typical use case for function execution is individual (one function, one 448 execution node), in ad-hoc networks the users can require the execution of 449 the same function on multiple drones, like, as presented in our case study 450 from Section 1, when scheduling a function that needs to check for fires on all 451 the areas covered by the drones of the network. As discussed in Section 3.3, 452 the language for policies that we provide to users allows them to express the 453 possibility of targeting multiple drones. 454

For simplicity, above and in the reminder of the section, we preserve the same notation f for functions in messages. However, we assume that the messages between the Base Station and the drones not only carry the code of the function as sent by the user but also a function instance identifier f_{id} defined by the Base Station to keep track of the function's lifecycle.

In the second case, progressive, one can assume that the network is highly
 dynamic, so that the Base Station cannot construct a reliable representation



Figure 4: Example of progressive scheduling, scheduling layer.

of the status of the network based on the gossiping information. Since the 462 Base Station cannot reliably identify the targets to execute the function, it 463 uses the information it has to "push" the request in a *direction* that likely 464 includes a target able to run the function. Hence, the Base Station does not 465 indicate which drone shall run the function and rather forwards the request 466 to the drones it can directly contact that lay in the target direction. In turn, 467 the drones that receive the progressive scheduling consider whether they can 468 execute the function or forward it to other drones in their neighbourhood 469 that might satisfy the scheduling requirements. We illustrate the logic of 470 progressive scheduling in Figure 4, reproducing the same function scheduling 471 result of Figure 3. In particular, notice that, differently from Figure 3, the 472 message transmitted by the Base Station and the other nodes has the same 473 structure of the one sent by the user, i.e., it does not include an indication 474 of a specific target drone, but rather uses the generic (f, a) form. 475

476 3.3. Execution Targets: Strong vs Weak Semantics

As mentioned, in traditional serverless platforms, users invoke the piecemeal execution of functions and ignore which target runs it². However, in our context, users might need to run the same function on multiple targets, e.g., to gather information about the different areas they are in.

²This abstraction over the topology of execution nodes constitutes one of the hallmark traits of classical serverless computing, i.e., since nodes are all functionally equivalent, users can safely ignore which node executes their function.

Strong Semantics: "one" and "all". When our users specify a scheduling policy, they can indicate whether to run the function on "one" valid target which captures the binding with an existential quantifier, i.e., we select one drone among those that can execute the function — or on "all" valid targets — which captures the binding with a universal quantifier, i.e., we select all drones among those that can execute the function.

Notably, the two invocation modalities do not only imply different tar-487 geting behaviours of the Base Station (and the intervening drones, in the 488 progressive case), but also different timings for forwarding the response back 480 to the user. Indeed, under the "one" scheduling request modality, the Base 490 Station directly forwards back to the user the response it receives from the 491 target executor (possibly dispatched through the other drones that connect 492 them). On the contrary, under the "all" scheduling request modality, the 493 Base Station implements a scatter-gather pattern, where it responds to the 494 user only after it received all responses to the requests it sent to the target 495 drones. 496

Of course, failures can happen. For instance, in the "one" case, the Base 497 Station might not receive a response from the target drone, e.g., because 498 the target correctly received the request but could not send back a response 499 (e.g., because it became disconnected or due to malfunction). To prevent 500 the Base Station from indefinitely waiting for a response, we pair a timeout 501 to each scheduling request and, if the Base Station received no response af-502 ter the given timeout, it reports the failed execution back to the user. By 503 extension, under the "all" execution modality, we consider failed the exe-504 cution of a function invocation when at least one of the requests did not 505 return its response within the timeout. These issues include events such 506 as communication problems (e.g., lost messages during both the scheduling 507 and the response message deliveries) and node failures (e.g., drones that be-508 come irresponsive due to hardware malfunction). In all these cases, the only 509 component that tracks the status of function executions is the Base Station, 510 which follows the above failure-handling logic to make sure that user requests 511 eventually terminate and report back to the user a response. 512

Since, in the above scheduling semantics, we impose strict requirements on the scheduling of functions, we call it *strong*, i.e., where the "one" scheduling modality entails the execution on exactly one target and the "all" modality requires all targets to return a response. Practically, a system can support this strong semantics if the underlying gossiping network can provide all the information about the drones, i.e., if the network enjoys a broad bandwidth ⁵¹⁹ so that the drones exchange all the necessary information with the Base ⁵²⁰ Station and the network can support the efficient and timely delivery of the ⁵²¹ coordination messages.

Weak Semantics: "at least one" and "almost all". When the connections 522 among the drones are weaker, the network experiences narrow-bandwidth 523 communications, which challenge the timely delivery of coordination mes-524 sages and the effective propagation of the full amount of information about 525 the status of the drones. In these cases, imposing the strong semantics pre-526 sented above could be counterproductive, since its strict requirements would 527 likely fail many invocations, due to network latencies. For this reason, we 528 introduce a second type of semantics, called *weak*, that modifies the interpre-529 tation of the "one" and "all" scheduling modalities. Under the weak semantics, 530 we soften the exactly "one" requirement to an "at least one", where the Base 531 Station (and intervening drones) sends a set number of concurrent requests 532 (defined by the user, in the policy) to possible targets and responds with a 533 positive answer as soon as it receives the first response from one of the tar-534 gets. Similarly, we weaken the "all" universality requirement to an "almost 535 all" interpretation that specifies a timeout for the gathering of the responses. 536 If all responses reach the Base Station within the set timeout (by the user, in 537 the policy), it behaves as in the "all" modality; otherwise, the Base Station 538 sends back to the user as many responses as it received at the triggering of 530 the timeout (and discards all possible successive ones). 540

The above logic defines how our proposal handles fault tolerance in the 541 highly dynamic context of drone networks. Essentially, fault tolerance is 542 implicitly and effectively managed at the software-component level through 543 the continuous operation of the gossiping knowledge dissemination protocol. 544 Specifically, nodes dynamically react based on the latest disseminated infor-545 mation. If a node fails to send periodic updates or responses (due to compo-546 nent failures or intermittent connectivity), other nodes inherently detect this 547 occurrence by the absence of fresh data and consequently ignore or remove 548 this node from their local views. Additionally, nodes actively monitor their 549 software components and, upon detecting local component failures, imme-550 diately stop advertising their resource availability through the gossip-based 551 heartbeat messages. This rapid dissemination of status changes ensures that 552 the entire network quickly adapts to node-level failures and maintains robust 553 scheduling decisions. While all elements maintain a view of the status of the 554 network, the system does not need these views to be consistent since each 555

drone determines whether it can execute a function for which it receives an 556 allocation request. Regarding the handling of failures in function schedul-557 ing, the Base Station is the only component that maintains the status of the 558 ongoing executions and registers an execution failure when a given function 559 allocation fails to respond within a set timeout. The system does not imple-560 ment any recovery/reallocation policies to handle faults and rather leaves to 561 the user the decision on what to do in response to a failed function execution 562 attempt. 563

Strong ang Weak Semantics, Exemplified. To better illustrate how each execution target presented in this section works in practice, let us consider a situation where we have several drones around an arbitrary point in space p, and we want to schedule a function f on them.

In the strong semantics case, using the "all" target, we would send the 568 function to each suitable drone in the area (whether the delivery of the ex-569 ecution message happens in a single hop or through multiple hops depends 570 on the topology of the network), and wait for all of their responses. If just 571 one of these executions happens to fail, or the response does not reach the 572 Base Station within a given timeout, the entire invocation is considered un-573 successful. The user receives no result, except for a notification of the failed 574 execution attempt. 575

Still in the strong semantics, using the "one" target, we would send the function to the "most suitable" (as determined by the selection policy, discussed in Section 4) drone available, e.g., the one with the most available resources, and wait for its response. Similarly to "all", if the drone fails to respond, the invocation is considered unsuccessful, except the failure is only imputable to the single targeted drone.

Under the weak semantics, the "all" target defaults to its "almost all" counterpart. We would still send the functions to all the suitable drones in the area. However, failure handling changes significantly: after a set timeout, the Base Station sends back to the user all the responses it received for this execution request, even if some target drones have not responded yet. In other words, unlike with the 'all' target, we tolerate partial results and consider it a successful invocation.

The weak counterpart of the "one" target is "at least one", which would send invocation requests to a set number of suitable drones in the area. In this case, we consider the invocation successful unless all targeted drones fail to respond within a given timeout. The first response the Base Station receives ⁵⁹³ is the one sent back to the user and all others are discarded. Compared to ⁵⁹⁴ the "one" target, we lose the guarantee that the function is only run once; ⁵⁹⁵ however, the use of redundant requests increases the invocation's tolerance ⁵⁹⁶ to failure.

⁵⁹⁷ 3.4. Mixing Scheduling Modalities and Network Semantics

The attentive reader might have already spotter the fact that the hierarchical and progressive scheduling modalities and the weak and strong semantics are orthogonal ingredients that one can mix to obtain different scheduling logics.

While this observation is sound, we notice that, e.g., pairing the hierarchical scheduling with the weak semantics, could produce an inefficient configuration, i.e., one where the Base Station does not have enough information to choose a target node for executing the function or the network does not have enough capacity to support the effective completion of the flow of function execution.

Hence, the interplay between the scheduling modalities and the network semantics generates a design/combination space that deserves dedicated work to explore.

⁶¹¹ Since this proposal regards the general presentation of our approach, in ⁶¹² the reminder, we fix a combination of these dimensions and delegate their ⁶¹³ proper exploration to future work (e.g., with specific case studies and exper-⁶¹⁴ iments).

Concretely, we fix the combination of the strong semantics with the hi-615 erarchical scheduling modality and the weak semantics with the progressive 616 modality, which, we conjecture, are the most reasonable pairings of this space 617 w.r.t. the QoS of ad-hoc networks. Our rationale for this choice is that 618 the strong semantics requires broad network bandwidth and timely updates, 619 which is also an important requirement of the hierarchical modality. Sim-620 ilarly, pairing the weak semantics with the progressive scheduling modality 621 mediates the limitations of narrow-bandwidth, lagging networks. 622

Given that the decision on whether to follow the strong or weak semantics is a system-level choice — specifically, taken by the Base Station depending on the information it receives from the drones — when the users define the policy for their functions, we ask them to indicate which strong semantics scheduling configuration they want for their function ("one" or "all"), followed by the parameter of the corresponding weak interpretation (the number of repetitions for the "at least one" and the timeout for the "almost all"). In



Figure 5: Example of function execution offloading, scheduling layer.

this way, depending on the current configuration of the network, the Base
Station follows the corresponding interpretation and can enact the related
scheduling modality.

Function Execution Offloading. Another important trait of our proposal is 633 that users can indicate whether a drone can offload the execution of a function 634 to some computation node outside the drone network. When a user specifies 635 that the target drone can offload function execution, they can define a piece 636 of data that might be on the drone and that the function requires for its 637 execution — this behaviour is typical of workflows where chains of functions 638 manipulate the data found in some repository, which, in this context, is the 639 local cache of drones. Thus, if the drone decides to offload the execution 640 of a function, e.g., due to limited execution capacity, it retrieves the data 641 indicated by the user (e.g., querying its local key-value data storage) and 642 sends an "offload" message to the Base Station that carries both the code of 643 the function and the data it needs to run on. 644

As an example, depicted in Figure 5, consider function f, which needs 645 to run within a certain location to gather some sensory data, like humidity 646 and light irradiation. In the figure, using hierarchical scheduling, the Base 647 Station determines that D is the target drone, issuing the placement of that 648 function with the message (f, a, D). However, when the message reaches 649 D, the drone realises it cannot run the function, e.g., because its current 650 workload level does not afford it. Since the scheduling policy a specifies that 651 D can offload the execution of this function, D sends back to the Base Station 652

a message $(f, a, \delta)^3$, where δ carries the possible data found on D useful to run f, as indicated by the offload parameter in a — of course, one can also introduce refinements of the offloading logic, e.g., having the Base Station preserve a copy of f and a so that the drone can just send back a message (f_{id}, δ) carrying the function identifier and the data.

Notice that, in Figure 5, we do not indicate which device runs f outside the ad-hoc network (and belonging to the Edge-Cloud domain). Since the execution of the offloaded functions is an external ingredient to our proposal — i.e., the ad-hoc drone serverless network works irrespective of how the Base Station decides to handle the execution of these functions — we abstract away the specification of the execution of offloaded functions, and we focus only on the specification of the offload functionality.

⁶⁶⁵ 3.5. Software Components: the Base Station and the Drones

We close this section by presenting the software components that support function scheduling, routing, execution, offloading, and response in the elements that make up the ad-hoc network: the Base Station and the drones. Starting from the Base Station, depicted in Figure 6, we find three main software components: the function scheduler, the state sentinel, and the function execution tracker. The first two modules deal respectively with the function scheduling and the gossiping layers.

The function scheduler is the component that handles the request for function execution from users (represented in the figure with the inbound message (f, a) on the left of the component), their allocation on the proper drones (exemplified with the outbound message (f, a, D) on the right of the component), and the relay of responses (e.g., (f_{id}, r)) and offloading (e.g., (f, a, δ)).

The state sentinel module deals with the status information of the drones disseminated through gossiping. In the Base Station, this component acts as a "sink" of information, used by the function scheduler to gather information on the status of the drones present in the system.

The last component, the function execution tracker, records the status of the execution of functions, including their cardinalities (e.g., a function can

³From brevity, we assume that drones identifiers $D_1, ..., D_n$ and data, as contained in δ , belong to distinct domains, so that the drones and Base Station can distinguish the nature of the message they receive. In implementations, there is no need for such assumption because messages usually carry a label to categorise them.



Figure 6: Software components found within the Base Station in our approach.

have a cardinality equal to the number of nodes if the scheduler issued its execution on "all" nodes) and their timeouts.

While slightly more complex, note how the architecture of the drones, depicted in Figure 7, recalls that of the Base Station. Indeed, while the drones are execution nodes for the functions, they also work as routers/schedulers of the requests, which requires the inclusion of a dedicated function scheduler and a state sentinel.

The main difference between the function scheduler found in the Base 692 Station and the one found in the drones is in their behaviour w.r.t. the 693 two scheduling modalities, hierarchical and progressive. When dealing with 694 a hierarchical request, the drone's function scheduler merely works as an 695 executor, if the receiver of the message is the target drone, or forwarder, 696 in charge of choosing the next hop of the message. Under the progressive 697 scheduling, the function scheduler assumes a much more "active" role, by 698 evaluating both whether its host can run the function and/or if it needs to 699 select which other drones to forward the execution message to. 700

To inform these decisions, the function scheduler uses the information about the other nodes in the network gathered by the state sentinel (similarly to its Base Station counterpart) along with the status of the drone's hardware and the current function execution workload (provided by the func-



Figure 7: Software components found within a drone in our approach.

⁷⁰⁵ tion execution runtime).

As mentioned, the state sentinel exchanges messages with the nodes in direct connection with the host drone to advertise its status information (gathered from the hardware and function execution runtime modules) and assemble an overall picture of the configuration of the network.

The function execution runtime is the component that, embedding a We-710 bAssembly engine, can run the functions. Note that, through the execution 711 runtime, the functions can both access the hardware of the drone (e.g., cam-712 era, rotors, etc.) and its local runtime APIs, among which its key-value stor-713 age, useful when running function workflows where some function "leaves" 714 some data on the drone and a successive one retrieves it for computation. 715 These APIs are particularly useful when implementing the offloading feature 716 — where the policy that accompanies the function can specify some value 717 that the function uses and the offloading command pairs with the request of 718 running the function on some external component. 719

Regarding failures, which mainly concentrate on the drones components,

we note that, since the network topology is highly dynamic, the software com-721 ponents react w.r.t. the information that reaches each node, so, if a node fails 722 to send updates or responses, the other nodes just ignore that node. In par-723 ticular, the Base Station, following its failure-handling logic (cf. Section 3.3) 724 handles the absence of responses from targeted drones through timeouts. If 725 a drone node realises that some of its components fail (e.g., a malfunction-726 ing camera), it stops advertising the availability of those components in the 727 gossiping phase, which updates all nodes about its changed state. 728

729 4. Ad-Hoc Allocation Priority Policies

We now present the language, called Ad-Hoc Allocation Priority Policies (AHAPP), that we provide to users for expressing the scheduling policies of their functions. As seen in Section 3, these policies inhabit the term a in message exchanges, like (f, a) and (f, a, D), for the scheduling of functions, and they are interpreted by the function scheduling modules of the Base Station and the drones to implement the scheduling logic of each function.

We start by introducing the constructs and parameters of the language and then apply them for capturing the functionalities of our case study (cf. Section 1).

739 4.1. The AHAPP language

The AHAPP language takes inspiration from the APP language and its dialects [17, 23, 24, 25, 26, 27, 28, 16]. As such, AHAPP is a subset of YAML [29], a widespread configuration language used in popular software for Cloud Computing like Docker and Kubernetes.

Unlike APP and its variants, AHAPP introduces a fundamental architec-744 tural shift. In previous approaches, a single script governed the scheduling 745 of all functions, using tags to specify different policies for functions (i.e., one 746 tag can specify the scheduling of many functions, each carrying that tag) 747 with their scheduling logic. An important novelty introduced by AHAPP is 748 that function scheduling requests see the function under scheduling paired 740 with its policy — essentially the policy "travels" with the function within the 750 system. This design innovation stems directly from the peculiar character-751 istics of our drone network, where individual drones are capable of taking 752 scheduling decisions. Consequently, the messages transmitted must not only 753 carry the function but also embed the logic determining the scheduling of 754 that function. While, in AHAPP, we do not use tags to associate functions 755

```
targets:
 all: { almost_all: n } | one: { at_least_one: n }
offload: id | b
affinity: { id_1, ..., id_j }
attributes:
 position: { lat: d, lon: d, alt: d, range: n }
 loadavg: { max: n }
 energy: { min: n }
 memory: { min: n }
 storage: { min: n }
 device: { h_1, ..., h_l }
  where h_{i,\ldots,l} \in \{
     rotors: { lock: b },
     camera: { lock: b, quality: id },
     gpu: { lock: b }
  }
strategy: [r_1, ..., r_k] where r_{1,\ldots,k} \in \{\text{position}, \text{ energy}, \ldots\}
```

Figure 8: The AHAPP language grammar.

to policies, we retain the functionality of tagging functions, useful to express
affinity constraints (detailed below).

In Figure 8, we show the structure of AHAPP policies and the values that each option can assume, using | to indicate alternatives, id for identifiers (tags), b for Booleans, n for natural numbers, and d for decimals. We comment on the constructs from top to bottom. In Figure 8, we slightly stylise the code for readability, in particular, we represent YAML inline sets as {...} rather than using the **!!set**-prefix notation.

Targets. The targets key encodes the execution modality of the function, i.e., either on "all"/"almost all" drones or on "one"/"at least one" of them, as introduced in Section 3.3. Recalling that section, we pair the configurations of the strong and weak semantics so that selecting the all option requires the user to also indicate the timeout of the almost_all clause; similarly, selecting the one option requires the specification of the number of requests the system is allowed to replicate under the at_least_one modality.

In case a policy selects the all-almost_all option, the Base Station inter-771 prets the instruction by sending to all drones the request to run its companion 772 function. The difference between the two modalities is that, under the strong 773 semantics, the Base Station considers as failed the scheduling of a function 774 when at least one of the requests to the drones did not return its response 775 within a set timeout (defined by the platform). Under the weak semantics, 776 the Base Station returns all responses received within the timeout indicated 777 by the user as parameter of the almost_all option — i.e., receiving at least 778 one response within the timeout configures as a successful execution of the 770 function. On the drones' scheduler, this option entails the execution of the 780 function (if the drone can run it) and/or the forwarding of the request. 781

When the policy selects the one-at_least_one option, the Base Station 782 interprets one by selecting a specific drone that shall run the function, while, 783 under at_least_one, the Base Station sends multiple requests and returns 784 to the user the first response it receives. To send these requests, the Base 785 Station splits the number of available requests n among the nodes that it 786 can contact directly (in particular, those in the direction where some nodes 787 can run the function). Practically, the Base Station achieves this "splitting" 788 by sending out requests that carry the same policy it received but where 789 the at_least_one parameter indicates a partition of the function instances 790 granted by the initial request. As an example, consider a Base Station that 791 receives a policy with at_least_one equal to 5, and it decides to send out two 792 requests. These requests would carry the same policy the Base Station re-793 ceived but would differ in the value of the at_least_one option, which could 794 be resp. e.g., 2 and 3 (or any non-zero 2-partition that sums is 5). In turn, 795 the drones that forward this request reduce the value of the at_least_one 796 by one if they run the function and split that value according to the number 797 of messages it forwards for function execution. 798

When targets is omitted, we consider one and at_least_one: 2 as defaults.

Offload. The offload key defines whether a drone supposed to run a function (e.g., when directly selected by the Base Station) and unable to do it, e.g., due to a lack of energy or memory capacity (cf. Section 3.3) can redirect the execution of that function to the "edge-cloud" part of the platform. Notice that users can specify two values for this option, either an identifier or a Boolean value. The boolean values simply indicate whether the offload option is enabled or not. When providing an identifier, the user indicates that the offload option is enabled and the offload request shall carry the piece of data (if any) pointed by that identifier as a key in the drone's local key-value cache. When omitted, the offload option is disabled.

Affinity. The affinity clause specifies a list of tags of functions that are 811 required to be in execution or have run (within a fixed timeout) on the drone 812 to run the function under scheduling. This option is particularly useful when 813 implementing workflows of functions, e.g., where function f runs on a certain 814 drone and then q has to run on the same drone because f left some data that 815 q needs for its execution. The default value of this option is the empty 816 list. Note that our interpretation of affinity does not imply synchronous 817 execution of functions, e.g., a function can express affinity to indicate that 818 it needs to run on a drone where some function ran and left some data 819 for the former to work on. To implement this mechanism, we assume that 820 the platform sets an "execution window" where the nodes keep track of the 821 function tags they have executed. In this window, nodes maintain a list of 822 tags of completed functions, and after a timeout period, they remove these 823 "stale" tags. This timeout mechanism works in tandem with the drone's 824 cache memory eviction policy, ensuring proper garbage collection of data left 825 behind on the drone. 826

Attributes. The attributes key encodes properties and physical resources required by the functions to run on a target drone.

Properties include position and loadavg. The position, with the in-829 dication of the latitude, longitude, altitude, expressed as decimals, and a 830 range that specifies a sphere of "tolerance" from the coordinates with radius 831 *n* metres. The loadavg specifies the maximal percentage of capacity occupa-832 tion (e.g., CPU) within a given period (set by the platform, e.g., 1 minute). 833 When omitted, the default value of the position sets all coordinates to 0 834 and the radius to its maximum definable value; the default value of loadayg 835 is 50. 836

We divide the overview of the resources into two groups. The first group, that of energy, memory, and storage specifies the minimal amount of resources necessary for effectively hosting the computation of the function. Specifically, the percentage of energy (default to 50), the amount of memory and storage capacity, in MB (resp. default to 8 and 32). The second group, that of devices, indicate a set of hardware resources found on the drone that the function requires accessing for its execution. The default value for this

option is the empty set. The possible elements in the set are (self-descriptive) 844 rotors, camera, and gpu. When indicating these requirements, the user can 845 express whether the function has to lock — Boolean and default to false 846 — that hardware exclusively during its execution, e.g., to move the drone 847 or to run compute-intensive parallel algorithms. During scheduling, if some 848 other function is "locking" the specified hardware, the drone cannot run that 849 function. The camera option can also specify a quality literal, like FHD and 850 4K; the default value is HD. 851

Strategies. The last option, that of the strategy, specifies an ordering among 852 the level of the attributes of the drones — except for the devices, which 853 are binary — that the scheduler can use to choose a target when it identifies 854 multiple valid ones. In practice, the strategy option allows the users to 855 refine the deployment of their functions by specifying the prioritisation logic 856 the scheduler should follow when choosing among equally valid targets. In 857 this way, users can orient the deployment of their functions according to cus-858 tom rules, e.g., by prioritising low-load or high-energy devices or minimising 859 latency through geographically-aware scheduling. 860

The values for this option are **position** (closer), **loadavg** (lower) and 861 available (higher) energy, memory, and storage. When defining the strategy, 862 the user provides an exclusion order for ranking the possible candidates. The 863 option also has as its trailing, default value the parameter random. Hence, if 864 the user omits the definition of the strategy option, the scheduler randomly 865 chooses among the possible candidates, otherwise it first filters out the can-866 didates considering the options in the list in order of appearance and then 867 chooses randomly among the remaining candidates, if any. 868

4.2. Modelling the Case Study with AHAPP

To illustrate the usage of AHAPP, we model the case study presented 870 in Section 1 as a serverless architecture of four functions and their related 871 AHAPP policies. For clarity, in the description below, we call the functions 872 by their associated tag. In the policies, we omit to specify the at_least_one 873 and almost_all options, which would require us to further characterise the 874 usage and topological context of the ad-hoc network. Thus, for brevity, we 875 omit their definition, which makes the policies use their respective default 876 values. 877

We call the first function ScanFireAlert. This function runs on all drones and processes real-time sensor data for fire aerial surveillance. To this aim,

Listing 1: ScanFireAlert policy.	Listing 3: ScanPeople policy.		
targets: all	targets: all		
attributes:	attributes:		
<pre>device: { camera, gpu }</pre>	<pre>position: { range: 100,</pre>		
	<pre>lat: X, lon: Y, alt: Z }</pre>		
	<pre>device: {</pre>		
Listing 2: AnalysePeopleScan policy.	camera		
targets: all	gpu		
<pre>affinity: { ScanPeople }</pre>	<pre>rotors: { lock: true }</pre>		
offload: ScanPeopleData	}		

the function analyses heat signatures and geospatial information to detect 880 potential fire outbreaks. When a fire is identified, the function runs an AI 881 model for fire recognition and categorisation, working on heat intensity and 882 smoke patterns. The response it produces is a fire alert, containing the 883 geolocation data, fire characteristics, and severity assessment. We show in 884 Listing 1 the code of the AHAPP policy for this function. In the policy, we 885 target all drones and impose the execution of the function on only those 886 nodes that equip a camera to acquire visual data on the territory hovered by 887 the drone and a gpu for the efficient execution of the fire-detection AI model. 888 After the response from the ScanFireAlert function(s), we invoke the 880 execution of the second function, called ScanPeople. The function runs on 890 drones in areas where ScanFireAlert detected fires. ScanPeople processes 891 data from the drone's camera to detect human heat signatures and movement 892 patterns. This analysis requires the repeated acquisition of visual data from 893 the same area, which implies that the function must have (exclusive) access 894 to the drone's rotors to move it according to its detection routine. Like 895 ScanFireAlert, also this function uses an AI model for categorisation. This 896 function does not provide a contextual response to the user — it just confirms 897 its successful execution — but it rather streams a continuous feed of data 898 which other functions can use to refine the detection of potential survivors. 899 We provide, in Listing 3, the policy for this function. In the policy, we 900 specify that the function shall run on all drones within a range of 100m 901 from the X, Y, Z position (we use these symbols as placeholders for the actual 902 coordinates) where ScanFireAlert found fire. Besides requiring drones with 903 a camera and gpu, the function needs to acquire a lock on the drone's rotors 904

⁹⁰⁵ to implement its detection routine.

Launched in close succession after ScanPeople, AnalysePeopleScan per-906 forms advanced analysis on the data produced and stored on the drones by 907 ScanPeople, to confirm the presence of survivors. The function employs 908 machine learning algorithms to distinguish human signatures from other 909 heat sources and movement patterns. The function evaluates the confidence 910 level of each detection and generates, as a response, detailed information 911 about potential survivor locations, including their estimated condition. The 912 AnalysePeopleScan allows us to illustrate two other features of the AHAPP 913 language: affinity and offloading. Indeed, rather than giving a location 914 of the drones that should run the AnalysePeopleScan function, we indicate 915 that they shall run on (all) drones where the ScanPeople function is running 916 (or has run). The reason for this indication is that AnalysePeopleScan runs 917 analyses on the data produced by ScanPeople, stored in the local cache of 918 the drone. Since this piece of data is fundamental for its execution, but the 919 function can otherwise run on the Edge-Cloud, we set the policy to allow the 920 offload of the function using the data stored, by ScanPeople, under the key 921 ScanPeopleData in the local cache of the drone. This configuration allows us 922 to handle the case where the target drone appears available to the Base Sta-923 tion, but it is unable to host the function because of lack of resources, which 924 might happen if, e.g., the Base Station is relying on stale status data and the 925 drone is hosting more functions than what the Base Station knows about. 926 With the offload option, the drone can react to this situation by sending 927 the request back to the Base Station — along with the data the function is 928 supposted to work on, found in the local cache of the drone, indicated as the 929 value of the offload clause. Once the Base Station receives an offloading 930 message, it forwards a request, carrying the data the function has to work 931 on, to a companion edge-cloud platform that acts as a "computational fall-932 back". When offloading the execution to the companion cloud-edge platform, 933 the Base Station takes care of interacting with the latter to also receive the 934 response after the execution of the function and notify the user. Since the 935 function runs in a domain not directly controller by ours, it is the companion 936 one that determines the execution of the function, including the handling of 937 possible errors that arise during its execution. 938

The last function, called TrackSurvivor, maintains continuous monitoring of confirmed survivors. It processes real-time GPS coordinates and video feeds to track survivor movements and status. The function generates a stream of data that include precise location updates, visual confirmation

```
Listing 4: TrackSurvivor policy.
targets: one
attributes:
position: { range: 25,
  lat: X, lon: Y, alt: Z }
device: {
  camera: { lock: true }
  gpu: { lock: true }
  rotors: { lock: true }
  }
strategy: [ energy, position ]
```

data, and dynamic access route calculations for rescue teams [30]. Since, 943 now, we have precise information about the survivors, we issue the individ-944 ual (one) execution of a TrackSurvivor function instance for each group of 945 survivors in a given location. The tracking routine requires the exclusive 946 access to the drone's **rotors** and **camera** to follow the survivors and of the 947 drone's gpu for the efficient processing of data for streaming. Since more 948 than one drone can be valid for the execution of the function, we specify 949 the strategy for choosing among multiple alternatives, i.e., first, we favour 950 the drones with the highest energy level (since the function needs to con-951 tinuously track the survivors for as long as possible), second, we favour the 952 drones closer to the required position (within the 25m range, imposed by 953 the policy) — as per definition (cf. Section 4.1) if these two selection steps 954 fail to weed down the options to one drone, we default to randomly select one 955 of the alternatives. 956

On the Complexity of AHAPP Scheduling. Computationally, the complexity of scheduling functions interpreting APP scripts is generally linear in the size of the nodes that execute the functions [17]. AHAPP is not an exception.

We break down the above statement considering the weak and strong semantics, which we pair respectively with the hierarchical and progressive scheduling — hence, when referring to the scheduling approach we also implicitly refer to the related semantics.

⁹⁶⁴ Under the hierarchical scheduling, the Base Station has precise knowledge ⁹⁶⁵ of all the drones in the network. Irrespective of the targets value, the Base

Station considers the application of the policy against all drones. Indeed, 966 under the all option, the Base Station needs to evaluate the policy against 967 the complete set of available nodes to select the ones that can execute the 968 function. Similarly, under the one option, the Base Station has to evaluate 969 the policy against all nodes to find the highest-scoring one which can execute 970 the function (or randomly choose among the tied highest-scoring ones, if any). 971 In the case of multi-hop function request messages, the intervening drones 972 act as forwarders of the message. This task is computable in constant time, 973 assuming the drones maintain routing tables computed when they receive 974 information from the gossiping protocol. 975

When running under the progressive scheduling, the Base Station does 976 not have a complete view of the elements in the network, but rather a local 977 view limited to the one-hop drones it can directly reach. Similarly to the 978 hierarchical scheduling (and, as argued above, irrespective of the target op-970 tion), the Base Station has to compare the policy against all drones, although 980 limited to the one-hop ones in its neighbourhood, making the complexity of 981 the task linear in the number of drones in the network, in the worst-case 982 scenario — a degenerate case where all drones are in the neighbourhood of 983 the Base Station, which can directly reach them. In the general case, the 984 Base Station considers the applicability of the policy against a fraction of 985 the drones in the network, to decide to which one-hop drones to forward the 986 request to. Once a drone in the network receives a progressive function exe-987 cution request, it essentially follows the same (linear) algorithm of the Base 988 Station, considering whether it can execute the function itself and whether 980 it shall forward the execution to other one-hop nodes in its neighbourhood. 990

⁹⁹¹ 5. Scheduling and Gossiping Layer Network Protocol

After having described the specification of the programming and be-992 haviour of the function scheduling layer, we present the details of the un-993 derlying network protocol, useful to maintain an accurate and up-to-date 994 view of both the network topology and individual drone status across the dis-995 tributed system and support function scheduling. The protocol addresses key 996 challenges in mobile ad-hoc networks, particularly drone swarms, where dy-997 namic topologies, intermittent connectivity, and varied hardware capabilities 998 require robust distributed communication mechanisms. Since this protocol 999 both provides information to the upper layer of function scheduling and sup-1000 ports the efficient relay of function allocation messages, we discuss the main 1001

technical details that characterise how the scheduler chooses its allocation 1002 targets (the nodes where functions shall run) and how the network deliv-1003 ers the related messages to destination. In addition to supporting efficient 1004 function scheduling, the proposed gossip-based knowledge dissemination in-1005 herently provides robust fault tolerance. Due to the highly dynamic nature 1006 of drone swarm topologies, the protocol is designed to react immediately 1007 to changes in network conditions. Nodes detect failures implicitly through 1008 the absence of periodic updates, promptly removing unreachable nodes from 1009 their local views. Additionally, if a failure occurs, each neighbour node in-1010 stantly ceases advertising its availability through the gossiping mechanism, 1011 ensuring rapid propagation of updated state information. Consequently, the 1012 function allocation module can dynamically adapt to network disruptions, 1013 maintaining robust and resilient scheduling decisions. 1014

Elements of the Network Protocol. We model the system as a set of n_{II} mo-1015 bile or static nodes forming a multihop mobile ad-hoc network. To enable 1016 serverless function scheduling in such environments, we propose a *function*-1017 allocation-aware protocol that abstracts the complexity of the underlying 1018 network from the scheduling layer. This protocol ensures that function allo-1019 cation requests are handled efficiently by disseminating network knowledge 1020 and facilitating function scheduling decisions. The protocol consists of two 1021 main modules: *Knowledge Dissemination*, which spreads network informa-1022 tion, including topology, hardware capabilities, and node positions (cf. Sec-1023 tion 5.1); and Function Allocation Module, which determines how to allocate 1024 function execution requests based on scheduling (AHAPP) instructions and 1025 the available knowledge (cf. Section 5.2). These two modules allow a clear 1026 separation between knowledge dissemination and function allocation. Knowl-1027 edge dissemination operates as a background process, proactively updating 1028 all nodes with real-time system status information like network topology and 1029 node hardware availability. In contrast, function allocation is triggered on 1030 demand when a function request arises. This dual-phase approach ensures 1031 that the system remains continuously updated and capable of immediate 1032 function allocation, boosting both responsiveness and resource efficiency. 1033

¹⁰³⁴ 5.1. Knowledge Dissemination

Efficient function scheduling in mobile ad-hoc networks depends on the ability of nodes to exchange real-time information regarding network topology, resource availability, and mobility patterns. To address this challenge,

we propose a gossip-based knowledge dissemination protocol that propagates 1038 relevant information through periodic *heartbeat* messages. The primary ob-1039 jective of this protocol is to maintain an up-to-date network view while mini-1040 mizing communication overhead, which is essential in dynamic environments 1041 where excessive message exchange can degrade performance. The proto-1042 col supports three levels of knowledge dissemination, each representing a 1043 trade-off between information completeness and communication efficiency: 1044 full dissemination, full dissemination with change detection, and aggregated 1045 dissemination. 1046

Each node periodically transmits a Heartbeat message to its one-hop neighbours at fixed intervals of t_H seconds. This message consists of two components: local information and global information. The local information includes all parameters necessary for function scheduling, as detailed in Section 4. Specifically, let s_H byte be the size of the local information packet where each node reports data like its position, CPU load, residual energy, available memory and storage, camera quality, and device locks.

The global information component varies depending on the selected knowledge dissemination method, which determines how network-wide data is shared. This aspect, along with specific periodicities of the Heartbeat message, differentiates the three dissemination strategies discussed below.

Full Dissemination. In the full dissemination model, each node periodically 1058 transmits its own status along with the status of all its known neighbours. 1059 Over time, this approach ensures that every node acquires a complete global 1060 view of the network. The primary advantage of this method is its ability 1061 to support global decision-making for hierarchical scheduling, allowing nodes 1062 to make globally-informed placement and routing decisions. Indeed, sim-1063 ilarly to the BATMAN protocol [31], this technique enables each node to 1064 discover and store the best next-hop toward a given destination based on 1065 received broadcast messages. The protocol leverages hop-count metrics to 1066 determine the optimal next-hop, ensuring efficient packet forwarding. Since 1067 only local next-hop information is maintained rather than full routing tables, 1068 the approach remains highly adaptive to mobility and topology changes. If a 1069 next-hop node becomes unreachable, the system dynamically selects an alter-1070 native best next-hop based on newly received Heartbeat message, allowing 1071 for rapid adaptation to network disruptions. In a static network, repeated 1072 iterations eventually disseminate each node's information to all others, en-1073 suring network-wide synchronization. However, in a dynamic environment, 1074

frequent topology changes can result in the propagation of stale information, rendering it obsolete before reaching all nodes. Consequently, nodes
may base scheduling decisions on outdated network views, leading to suboptimal or incorrect function placements, which ultimately degrades system
performance.

Furthermore, as each Heartbeat message eventually reaches a size of 1080 $n_{U} \cdot s_{H}$ bytes, this approach introduces significant scalability challenges. The 1081 global information contained in the message includes data from $n_U - 1$ nodes, 1082 increasing the transmission burden. As a result, the overhead scales propor-1083 tionally with the number of nodes, leading to excessive communication costs. 1084 In large and highly dynamic networks, continuously propagating exhaustive 1085 state information can cause network congestion, increase latency, and de-1086 grade overall system efficiency. 1087

Full Dissemination with Change Detection. To reduce unnecessary transmis-1088 sions while preserving a global network view, we introduce an optimised 1089 variant of full dissemination in which nodes transmit updated information 1090 only when a state change occurs. Even though this technique uses the same 1091 periodic Heartbeat message in each interval of t_H seconds, the message will 1092 only be sent if there are changes to be reported. Moreover, only the changed 1093 attribute is sent, reducing unnecessary transmissions while preserving essen-1094 tial network-wide synchronization. 1095

This approach retains the benefits of complete dissemination while mini-1096 mizing redundant transmissions, thereby reducing communication overhead. 1097 The effectiveness of this method depends on the rate of topology and system 1098 changes: in highly dynamic scenarios, frequent updates may still result in sig-1099 nificant overhead, whereas in more stable conditions, the protocol achieves 1100 a near-optimal balance between completeness and efficiency. To refine up-1101 date propagation, we introduce threshold-based triggers for non-binary node 1102 attributes such as position, CPU load, residual energy, available memory, 1103 and storage. Specifically, an offset threshold is defined for each parameter, 1104 specifying the minimum value change required to report it. These thresholds 1105 are denoted as follows: th_p for the position, th_c for the CPU load, th_e for the 1106 residual energy, th_m for available memory and th_s for the available storage. 1107

Another improvement is an alternative Heartbeat comprising a set of change blocks. Each block contains the essential node identification, a sequence number, an attribute identifier, and its new value. Despite these optimizations, similar to the full dissemination model, the Heartbeat message will eventually reach a size of $n_U \cdot s_H$ bytes, introducing scalability concerns in large networks.

Aggregated Knowledge Dissemination. In contrast to the previous methods, 1114 the aggregated knowledge model restricts information exchange to local node 1115 states and summarised network metrics. Each node transmits every t_H sec-1116 onds its own status, reduced to only position information, along with an ag-1117 gregated representation of its neighbourhood, such as minimum, maximum, 1118 or average resource availability. This approach significantly reduces com-1119 munication overhead by preventing the propagation of detailed node-level 1120 data across the network. Specifically, the network overhead is drastically 1121 minimised, as each Heartbeat message does not exceed s_H bytes. However, 1122 this method introduces a trade-off between efficiency and decision-making 1123 accuracy, as nodes make scheduling choices based on partial and abstracted 1124 network knowledge rather than complete global awareness. Consequently, 1125 the aggregated knowledge model is most suitable for progressive allocation 1126 strategies, where function scheduling relies on localised rather than global 1127 decision-making. 1128

More in detail, the global information component of the Heartbeat mes-1129 sage starts from the own status information; the approach chooses the best 1130 metrics of each neighbour to compose the Heartbeat message, intending to 1131 represent its potential. Less CPU load from the neighbour nodes is used 1132 instead of its own, as well as greater residual energy, available memory, and 1133 storage to compose the Heartbeat. Regarding non-binary node attributes, 1134 like camera and GPU availability, the aggregated information considers the 1135 node's and its neighbours' availability. 1136

Note that position information is not included in the global information. This omission is intentional, as only one-hop location data is used to forward function allocation requests based on positional constraints (more detail in Section 5.2). This approach is analogous to geographic routing, as employed in GPSR (Greedy Perimeter Stateless Routing) [32], where GPS coordinates are leveraged to forward packets toward a specific destination.

Table 1 outlines the advantages and limitations of the different dissemination strategies, highlighting their impact on communication overhead and their suitability for various network dynamics, i.e., the expected rate of topological changes over time. The effectiveness of each method is contextdependent, influenced by network size, stability, and the computational requirements of the deployed functions. In our evaluation (cf. Section 6), we

Dissemination Method	Knowledge Scope	Communication Overhead	Application Context
Full	Global (all nodes)	High	Stable topologies
Full w/ Change	Global (updated states)	Medium	Stable and Moderate dynamics
Aggregated	Local + aggregated metrics	Low	Dynamic topologies

Table 1: Comparison of Knowledge Dissemination Strategies.

analyse the performance trade-offs among these approaches.

1150 5.2. Function Allocation Module

Assuming the dissemination of network knowledge, when the system re-1151 ceives the request to execute a function, it must determine the most suitable 1152 node where to allocate the function, taking into account network constraints, 1153 resource availability, and scheduling efficiency. The function allocation mod-1154 ule governs this process, ensuring that function allocation works based on 1155 scheduling (AHAPP) instructors and the available knowledge at each node. 1156 As described in Section 3, the system supports two scheduling strategies: 1157 hierarchical scheduling, which relies on a global network view, and progressive 1158 scheduling, which operates under localised knowledge. We remind that, in 1159 this work, we make these two approaches respectively correspond to strong 1160 and weak execution semantics (cf. Section 3.4). The hierarchical schedul-1161 ing assumes complete network knowledge, enabling it to achieve efficiency 1162 under the strong semantics, where function allocation is directly decided by 1163 the Base Station and the intervening nodes work are routers. Conversely, 1164 progressive scheduling is paired with the weak semantics, where a best-effort 1165 approach governs scheduling decisions based on neighbourhood-level informa-1166 tion — due to incomplete network knowledge. In the following, we assume 1167 the hierarchical and progressive scheduling modalities paired with their re-1168 spective strong and weak semantics. 1169

Hierarchical Scheduling. The hierarchical scheduling strategy assumes that nodes have access to global knowledge obtained through either the *Full* or *Full with Change* dissemination methods. When a function allocation request for function f with scheduling policy a is initiated, the scheduling node (i.e., the Base Station) has a complete view of the network and can determine the optimal execution target. The request is then transmitted directly to the selected node using multi-hop forwarding. This approach ensures precise function placement and avoids allocation overhead — w.r.t. the time/redundancy thresholds of the weak semantics. However, its efficiency depends on the accuracy and timeliness of the global knowledge, making it more suitable for stable topologies with low-latency communication (and broadband connectivity).

The scheduling procedure depends on the target value, which, in the hierarchical methodology, can be set to either all or one.

- 1184 1185
- if all is selected, the system identifies all nodes that satisfy the scheduling constraints and allocates the function execution on each of them;
- 1186 1187
- if one is selected, a selection process determines a single "suitable" node for execution.

To facilitate the selection process, driven by the strategy clause, we 1188 introduce ranking classes for requested resources, categorising them into three 1189 discrete levels: high, medium, and low. This categorisation simplifies the 1190 selection process, as a ranking based on continuous numerical values may be 1191 impractical when distinguishing between groups of nodes with comparable 1192 capabilities. To achieve this categorisation, we define two threshold values 1193 for each parameter: 11, which differentiates between low and medium, and 12, 1194 which separates medium from high. The thresholds are defined as follows: 1195 l_{1_c} and l_{2_c} for CPU load, l_{1_e} and l_{2_e} for residual energy, l_{1_m} and l_{2_m} for 1196 available memory, and l_s and l_s for storage. 1197

Unlike these real-valued parameters, the binary constraints, such as locks and position, do not require threshold-based categorisation. Locks are either present or absent, directly determining whether a node is eligible for execution. Similarly, position-based constraints are treated as boolean conditions, where a node is considered inside or outside the requested deployment area.

Progressive Scheduling. The progressive scheduling modality is designed for 1204 dynamic networks, where the nodes have only localised and aggregated knowl-1205 edge. Instead of making a global decision, the scheduling process follows a 1206 best-effort heuristic, progressively forwarding the request toward regions with 1207 successful execution capability. Each node receiving the request, evaluates 1208 its own resources and those of its direct neighbours, determining whether to 1209 execute the function locally or forward the request further — except the Base 1210 Station, which can choose only the second option. This method introduces 1211

greater adaptability and reduces overhead, as nodes do not need to maintain 1212 an exhaustive global state. However, suboptimal placements may occur due 1213 to the incomplete network view available to each node. Unlike the hierar-1214 chical scheduling, the progressive approach relies solely on local information, 1215 meaning that each node knows only its direct one-hop neighbours and aggre-1216 gated global information. Consequently, the scheduling cannot be globally 1217 assigned (like in the hierarchical scheduling, by the Base Station); instead, 1218 the scheduling follows a best-effort approach, dynamically deciding the next 1219 best hop (if needed) at each step. 1220

To implement the progressive scheduling logic, we employ a combination of ranking and probabilistic selection. Let $\mathcal{N} = \{n, n_1, n_2, ...\}$ be the set of potential allocation destinations, that includes node n and its one-hop neighbours. For each $\hat{n} \in \mathcal{N}$, we define a score function $v(\hat{n})$ as follows:

$$v(\hat{n}) = \alpha \cdot d(\hat{n})_{\beta} + (1 - \alpha) \cdot h(\hat{n}) \tag{1}$$

where $\alpha \in [0..1]$ is a pre-defined weight factor, $d(\hat{n})_{\beta}$ represents the direction factor, which depends on the parameter β that defines the maximum angular displacement, and the relative angular direction of \hat{n} (having position $(x_{\hat{n}}, y_{\hat{n}})$) with respect to the allocation destination position X and position Y, if specified in the AHAPP request. This value is computed as follows:

$$\vec{nd} = (\text{position} \setminus X - x_n, \text{ position} \setminus Y - y_n), \quad \vec{nn} = (x_{\hat{n}} - x_n, y_{\hat{n}} - y_n)$$
$$\theta = \arccos\left(\frac{\vec{nd} \cdot \vec{nn}}{\|\vec{nd}\| \|\vec{nn}\|}\right)$$

1232

1231

$$d(\hat{n})_{\beta} = \begin{cases} 1, & \theta \leq 0^{\circ}, \\ 1 - \frac{\theta}{\beta}, & 0^{\circ} < \theta < \beta, \\ 0, & \theta \geq \beta. \end{cases}$$
(2)

1233 In turn, $h(\hat{n})$ captures attribute constraints and is defined as:

$$h(\hat{n}) = \frac{\operatorname{lock} + \sum_{a \in \operatorname{ATT}} \min(1, (\frac{\hat{n}_a}{a})^2)}{|\operatorname{ATT}| + 1}$$
(3)

where ATT = {loadavg, energy, memory, storage} is the set of attributes, \hat{n}_{a} represents the aggregated value of attribute a received via the Heartbeat message from node \hat{n} , and lock is 1 if the aggregated device constraints in \hat{n} fulfil the request, and 0 otherwise.

After computing $v(\hat{n})$ for each node in \mathcal{N} , node n forwards the function request to node \hat{n} with probability:

$$p_{\gamma}(v(\hat{n})) = \min\left(1, \left(\frac{v(\hat{n})}{\gamma}\right)^2\right) \tag{4}$$

where γ is determined based on the targets value, with specific parameters γ_{almost_all} and $\gamma_{at_least_one}$ — we remind that in the progressive scheduling, due to its best-effort nature, the all and one targets of the hierarchical scheduling are respectively mapped to their almost_all and at_least_one counterparts, reflecting the system's inability to guarantee a strong-semantics scheduling. If $\hat{n} \equiv n$, the function is executed locally at node n.

Scheduling					
Strategy	Knowledge	Scheduling	Semantic	Best Suited For	
	Global	Direct selection		Stable topologies	
Hierarchical	(Full / Full w/ Change)	optimal placement	Strong	low-latency networks	
	Local	Forwarding-based		Highly dynamic networks	
Progressive	(Aggregated)	adaptive heuristic	Weak	bandwidth-constrained systems	

Table 2: Comparison of Function Allocation Strategies.

Semantics/Scheduling Selection. The selection of an appropriate function allocation semantics depends on multiple factors, including network stability, resource constraints, and application requirements. Table 1 provides an overview of the advantages and drawbacks of the two function allocation semantics. The observations align with those in the previous table on knowledge dissemination, as the effectiveness of scheduling methods is inherently tied to the availability and accuracy of network state information.

In Section 6, we conduct quantitative analyses that compare the above selection w.r.t. different network conditions.

In the future, we plan to explore more sophisticated approaches on a hybrid model that dynamically transitions between hierarchical and progressive strategies based on network conditions and resource availability. One such hybrid, adaptive mechanisms would autonomously switch among different dissemination and allocation methods to optimise allocation efficiency while minimizing communication overhead.

1261 6. System Evaluation

We present an evaluation of the proposed serverless function scheduling 1262 framework for ad-hoc drone networks through extensive simulations. Note 1263 that, although we use the case study introduced in Section 1 throughout the 1264 article to illustrate the potential and practical applicability of our approach, 1265 our evaluation is more general and comprehensive. In addition to simula-1266 tions based on scenarios from the case study, we conduct extensive generic 1267 simulations varying topologies, mobility patterns, and numbers of nodes, to 1268 thoroughly assess the performance and robustness of our proposed system 1269 under diverse network conditions. 1270

1271 6.1. Simulation Setup

To ensure a comprehensive evaluation, we conduct simulations in a controlled environment using predefined network configurations. For the simulation we use $OMNeT++^4$ coupled with the INET library⁵. We run the simulations with the default parameters defined in Table 3, which remain constant unless otherwise specified. In the following, we use dissemination method names and scheduling modalities interchangeably due to their tight coupling.

Table 3: Default Simulation Parameters			
Parameter	Value		
Number of Nodes n_U	12		
Function request generation period n_U	20s		
Node Speed (dynamic scenario)	$15 \mathrm{mps}$		
Message Transmission Interval (t_H)	5s		
Local Information Size (s_H)	176B		
Ranking Weights (α, β)	$\alpha=0.7,\beta=90^\circ$		
Probability Factors Heartbeat (p_H)	0.1		
Almost all probability (γ_{almost_all})	0.7		
At least one probability $(\gamma_{at_least_one})$	0.9		

¹²⁷⁹ The simulation considers both static and mobile scenarios as well as a re-¹²⁸⁰ alistic case study based on the fire detection and survivor tracking scenario

⁴https://omnetpp.org/

⁵https://inet.omnetpp.org/

presented in Section 1. We generate random function requests, associated
with scheduling policies compliant with the AHAPP language (cf. Figure 8).
Additionally, we randomly assign the parameter values within each function
request, ensuring a diverse set of test cases for evaluating system performance.

1286 6.2. Performance Metrics

¹²⁸⁷ We evaluate our framework based on the following key performance indi-¹²⁸⁸ cators (KPI):

- Average Function Deployment Delay: The average time taken for a function to be successfully deployed.
- Average Function Deployment Number: The number of successfully deployed functions for each request generated in the system.
- Network Overhead: The size/amount of control messages exchanged in the network to achieve service deployment.

¹²⁹⁵ These metrics provide insights into the efficiency, reliability, and scalability ¹²⁹⁶ of the proposed approach under different conditions.

1297 6.3. Simulation Scenarios

¹²⁹⁸ We consider three simulation scenarios to evaluate the system under di-¹²⁹⁹ verse network conditions:

- Static Grid-Based Deployment: in this scenario, nodes are statically placed in a grid topology. This setup allows us to assess performance in a well-structured network, where topology remains constant.
- Dynamic Mobile Nodes: where nodes move dynamically following a predefined mobility model. This scenario simulates a realistic UAV network, where nodes frequently change positions, impacting connectivity and deployment decisions. In this scenario, the nodes move at a fixed speed of 15mps.
- Case Study: Fire Detection and Rescue: this scenario follows the realworld use case described in Section 1, where a drone swarm is deployed for fire detection, survivor localization, and tracking. The goal is to evaluate system performance in a multi-phase, real-world environment.





the static scenario.



Figure 11: Deployment delay in the static Figure 12: Deployment delay in the dynamic scenario.

Figure 9: Number of deployed functions in Figure 10: Number of deployed functions in the dynamic scenario.



scenario.

6.4. Evaluation Results 1312

We analyse the system behaviour under different deployment modalities 1313 and network configurations, focusing on the defined KPIs. The experiments 1314 are structured to evaluate the impact of scheduling policies, knowledge dis-1315 semination strategies, and network scalability. 1316

Deployment Delay and Number of Deployments. To assess the efficiency of 1317 function execution, we compare the number of successfully deployed func-1318 tions across different configurations. The results in Figure 9 and Figure 10 1319 highlight the differences between the two main scheduling modalities, all and 1320 one, under various knowledge dissemination methods. In these figures, the 1321 blue bar represents the number of deployable nodes at function generation 1322

time, providing a reference to evaluate the appropriateness of each modalityw.r.t. the different scenarios.

In the static scenario, the all target modality achieves near-optimal re-1325 sults across all knowledge dissemination methods. This outcome is expected, 1326 as the static nature of the system ensures that all dissemination approaches 1327 effectively spread network information. The one modality also performs well, 1328 except in the progressive approach, where its distributed and probabilistic 1329 nature leads to function allocations on more than one node, increasing re-1330 source usage. Conversely, in the dynamic scenario, the hierarchical modality 1331 exhibits significantly lower performance. Both the Hierarchical and Hierar-1332 chical w/ Changes methods deploy less than half of the possible functions. 1333 This performance drop is primarily due to outdated network information at 1334 the Base Station, which prevents accurate function placement. Additionally, 1335 during function deployment, nodes move rapidly, meaning that some initially-1336 selected nodes may no longer be in the correct area, while others, suitable 1337 ones are not considered. This limitation is effectively mitigated by the pro-1338 gressive method, which achieves near-optimal deployment. Its hop-by-hop, 1339 locally adaptive decision-making allows the system to dynamically respond 1340 to network topology changes, significantly improving function placement. 1341

Figure 11 and Figure 12 illustrate the deployment delay, which follows the same trend observed in the function deployment analysis. The delay is inherently tied to the number of deployed functions—as more functions are successfully deployed, the average delay increases. Consequently, while the progressive method exhibits higher delay values, this is a direct result of its greater number of successful deployments, rather than an inefficiency in the scheduling mechanism.

Network Overhead Analysis. We evaluate the impact of the scheduling and
knowledge dissemination modalities on network overhead by measuring the
total number of control messages exchanged. The primary contributor to
network load is knowledge dissemination, while function deployment incurs
minimal overhead. We report the results in Figure 13 and Figure 14.

Full knowledge dissemination generates significant communication overhead, as each node periodically shares its entire network state with its neighbours. While this approach ensures high decision accuracy in static scenarios, it results in substantial bandwidth consumption. The Full with Change method improves performance by transmitting updates only when a node's state changes, providing a better trade-off between accuracy and efficiency.





Figure 13: Network overhead delay in the static scenario.



Figure 14: Network overhead delay in the dynamic scenario.



the static scenario varying the number of nodes.

Figure 15: Number of deployed functions in Figure 16: Number of deployed functions in the dynamic scenario varying the number of mobile nodes.

The Aggregated method, which shares only summary metrics, achieves the 1360 lowest overhead but at the cost of reduced scheduling precision. When com-1361 paring deployment modalities, the all and one approaches exhibit similar 1362 overhead levels due to the fixed periodic nature of knowledge dissemination. 1363 However, in the Full with Change method, the all targetting results in higher 1364 overhead. This increase is likely caused by network congestion, leading to 1365 packet losses and retransmissions, which ultimately inflate the total number 1366 of control messages exchanged. 1367

Scalability: Impact of Network Size on Deployment Efficiency. To evaluate 1368 the scalability of the system, we analyse how the number of deployed func-1369





Figure 17: Battery Capacity for energy and position prioritisation strategy as the number of nodes increases. Higher values indicate better energy-aware deployments.

Figure 18: Deployment Distance Ratio for energy and position prioritisation strategy as the number of nodes increases. Lower values indicate better distance-aware deployments.

tions evolves as the number of nodes increases. The results are shown in Figure 15 and Figure 16, where we present data exclusively for the all targetting.

In the static scenario, all methods achieve optimal results, successfully de-1373 ploying functions to all eligible nodes. However, in the dynamic scenario, the 1374 increasing number of nodes leads to a continuously changing network topol-1375 ogy, introducing challenges for hierarchical approaches. In this case, both Hi-1376 erarchical and Hierarchical with Changes exhibit limitations. The Hierarchi-1377 cal method struggles with outdated information and larger Heartbeat mes-1378 sages, increasing scheduling errors. The Hierarchical with Changes method, 1379 which benefits from selective updates in static scenarios, loses its advantage 1380 in dynamic networks due to the continuous need for updates, leading to fre-1381 quent state inconsistencies. 1382

In contrast, the Progressive approach scales effectively with the number of nodes. Its local and distributed decision-making allows it to adapt to mobility, ensuring a steady increase in the number of successfully deployed functions as more nodes become available. The results indicate that the Progressive method remains largely unaffected by network size and mobility, maintaining deployment efficiency across different scales.

¹³⁸⁹ *Prioritisation Strategy Analysis.* To assess how effectively the hierarchical ¹³⁹⁰ and progressive scheduling strategies follow the user-defined prioritisation strategy, we analyse two metrics: the Battery Capacity (Figure 17) and the Deployment Distance Ratio (Figure 18). The Battery Capacity measures how closely the selected nodes align with the goal of maximising residual energy, whereas the Deployment Distance Ratio evaluates position optimisation, indicating whether the selected nodes minimise the distance to a given target. Results show that both scheduling methods effectively adhere to the specified prioritisation strategies.

Specifically, in Figure 17, we report the Battery Capacity for both the en-1398 ergy and position prioritisation cases, using both hierarchical and progressive 1399 scheduling methods. When energy is prioritised, both strategies are able to 1400 select nodes with higher residual energy, achieving higher Battery Capacity 1401 values. However, the hierarchical scheduler shows some limitations as the 1402 number of nodes increases, mainly due to the difficulties in maintaining al-1403 ways up-to-date information in a dynamic scenario. Additionally, the Figure 1404 shows that when position is prioritised, the Battery Capacity tends to stabi-1405 lize around 0.5 — confirming that energy is not a selection criterion in this 1406 configuration and is therefore not significantly affected. 1407

Conversely, when distance is prioritised through the *position* strategy, both scheduling methods effectively select nodes closer to the target location. As shown in Figure 18, the Deployment Distance Ratio decreases with the increasing number of nodes, as more nodes become available near the target point. This feature allows the schedulers to better optimise the deployment with respect to distance, clearly improving that metric compared to the other strategy where position is not prioritised.

¹⁴¹⁵ Case Study: a realistic serverless workflow. This section presents the results ¹⁴¹⁶ for the fire detection and rescue scenario, the realistic running case study we ¹⁴¹⁷ introduced in Section 1 and detailed in Section 4. The experiment simulates a ¹⁴¹⁸ $3 \times 3 \text{ km}^2$ area where a swarm of drones is deployed. The simulation proceeds ¹⁴¹⁹ as follows:

- At intervals of 1 minute, the Base Station (BS) issues the execution of
 the ScanFireAlert function, which is deployed to all drones equipped
 with a camera.
- 1423
 1424
 1424
 1424
 1425
 1425
 1426
 1426
 2. 5 minutes from the beginning of the simulation, we simulate a fire outbreak at specific spots in the area, triggering the alert by the previous functions and the request, by the Base Station, to execute the ScanPeople function on all drones near the fires.



Figure 19: Successful deployment ratio in the fire detection and rescue scenario.

3. Similarly, we simulate the presence of survivors, some near the fire
areas, which triggers the request to execute individual (one) TrackSurvivor function instances to track the survivors.

For this experiment, all drones are assumed to be equipped with cameras and GPUs. As a performance metric, we calculate the *successful deployment ratio* across the three consecutive functions, i.e., the cases when all the three functions are successfully deployed.

Figure 19 presents the performance results as the number of drones varies. 1434 When the number of drones is low, only a small percentage of cases achieves 1435 successful deployment of the full function sequence, primarily due to limited 1436 area coverage. In this scenario, the Hierarchical method outperforms the 1437 Progressive method because the data exchange overhead is small in small-1438 scale networks. However, as the number of drones increases, the Progressive 1439 method surpasses the Hierarchical approach, reaching a near-optimal suc-1440 cess rate. Conversely, the Hierarchical method struggles with high network 1441 dynamism and an increasing number of nodes, causing its performance to 1442 stabilise around 0.25. 1443

Overall, the experiments confirm that hierarchical scheduling is better for small or static networks, whereas progressive scheduling scales better in large and dynamic environments.

1447 7. Related Work

Recent advancements in serverless computing have extended its capabilities to handle other environments than the cloud one, where the paradigm originated. The integration of serverless computing in mobile ad-hoc networks (MANETs) is a relatively new research direction, with most prior work focusing on edge, fog, and cloud computing methods for improving task offloading and resource allocation. In this section, we discuss related work from the literature.

One work close to ours is by Singh and Adhikari [33], who propose a hi-1455 erarchical federated learning framework designed for edge-fog-cloud environ-1456 ments. Their solution introduces a pop-up ad-hoc network for dynamic com-1457 munication between edge devices and fog servers, integrated with a serverless 1458 data pipeline to improve processing efficiency. To enhance participation and 1450 resource usage at the edge, the platform incorporates a game model that 1460 rewards edge devices based on their contributions. Patterson et al. [34] also 1461 present work close to ours, introducing a platform for swarm coordination 1462 using a centralised controller able to deploy serverless functions on both the 1463 cloud and edge devices, showing improved performance predictability, lower 1464 network traffic, and battery efficiency compared to existing alternatives. 1465

Improving function execution in serverless platforms remains a focal point 1466 of research for serverless for edge computing. Jindal et al. [35] introduce the 1467 Function Delivery Network (FDN), an innovative extension of FaaS that 1468 enables function execution across heterogeneous clusters. Their Function-1469 Delivery-as-a-Service model selects the most suitable execution platform based 1470 on platform characteristics, potential collaboration with other platforms, 1471 and data localisation. Jindal et al.'s evaluation shows that scheduling func-1472 tions on edge platforms can substantially reduce energy consumption while 1473 maintaining adherence to Service Level Objective requirements. Garbugli et 1474 al. [36] tackle Quality of Service (QoS) management in IoT applications lever-1475 aging edge cloud FaaS through TEMPOS, a middleware that ensures per-1476 formance reliability in environments with high stochastic variability. Their 1477 approach coordinates QoS monitoring, control, and management across the 1478 Cloud-to-Things continuum, tackling challenges such as bandwidth fluctua-1479 tions and shared virtualised resources to provide a more stable and efficient 1480 execution environment. 1481

Several studies have explored how edge, fog, and cloud computing can improve computation offloading in ad-hoc networks, particularly in UAV-

enabled systems. For instance, Grasso et al. [37] propose a reinforcement 1484 learning-based job offloading mechanism in a flying ad-hoc network (FANET), 1485 where UAVs with mobile edge computing (MEC) capabilities collaborate to 1486 balance computational loads and reduce delay. Similarly, You et al. [38] 1487 present a joint optimisation strategy for task scheduling, resource alloca-1488 tion, and UAV trajectory planning to enhance computing performance in 1489 Other approaches leverage multi-tier architectures to enhance FANETS. 1490 computation efficiency. For example, Luo et al. [39] introduce an adaptive 1491 scheduling model that dynamically distributes workloads across UAV-based 1492 edge servers, optimising energy consumption and execution latency. In the 1493 same direction, Niu et al. [40] focus on task scheduling in emergency scenar-1494 ios, where UAVs act as mobile computing nodes to process data and maintain 1495 network connectivity in disaster-affected regions, while Khan et al. [41] inte-1496 grate MEC with secure identity-based signcryption to ensure confidentiality 1497 and authentication in FANETs. 1498

The integration of serverless computing in mobile ad-hoc networks is still 1499 in its early stages, with most efforts focusing on making serverless frameworks 1500 adaptable to edge computing environments. One of the key challenges in this 1501 domain is function scheduling, as traditional serverless models rely on cen-1502 tralised cloud infrastructure, whereas ad-hoc networks require decentralised, 1503 self-organising strategies. Several studies have examined orchestration mech-1504 anisms for UAV-based ad-hoc networks. For example, Grigoropoulos and 1505 Lalis [42] propose a hierarchical orchestration model that spans edge, fog. 1506 and cloud layers to improve function deployment in UAV-based networks. 1507 Similarly, Keshavamurthy et al. [43] introduce a distributed control frame-1508 work for UAV relay networks, addressing key challenges in network topology 1509 management and function execution. Beyond terrestrial and aerial MEC sys-1510 tems, space-assisted computing has gained attention. Lin et al. [44] explore 1511 a game-theoretic framework for task offloading in tactical ad-hoc networks 1512 that integrate low-Earth orbit (LEO) satellites and UAV-based MEC nodes. 1513 Lin et al.'s work highlights the challenges of resource allocation and pricing 1514 strategies in heterogeneous environments where UAVs and satellites provide 1515 computation offloading capabilities. In terms of task scheduling, several stud-1516 ies propose methods based on heuristic and machine learning to improve exe-1517 cution efficiency. Qin et al. [45] introduce a time-sensitive task scheduling al-1518 gorithm for UAV-based MEC systems, prioritising reconnaissance operations 1519 based on dynamic environmental conditions. Similarly, Secinti et al. [46] ex-1520 plores fog computing techniques for UAV-based software-defined networks, 1521

¹⁵²² enabling adaptive task allocation and network-aware function placement.

In ad-hoc serverless computing, where network topology is highly dynamic 1523 and nodes frequently join or leave the network, service discovery is essential 1524 to ensure that function invocations can reliably locate and communicate with 1525 appropriate execution nodes. Ververidis Polyzos [47] provide a comprehensive 1526 survey of service discovery mechanisms for MANETs, categorising different 1527 strategies based on centralised, decentralised, and hybrid architectures. Be-1528 yond individual studies, Ferrer et al. [48] provide a comprehensive survey of 1529 mobile cloud, edge computing, and mobile ad-hoc computing, identifying op-1530 portunities and challenges in moving from centralised cloud models to fully 1531 distributed cloud paradigms. The study highlights ad-hoc cloud computing 1532 as a critical evolution in distributed computing, emphasizing the need for 1533 self-organising, fault-tolerant resource allocation mechanisms. 1534

Comparison with Our Work. Besides the works by Singh and Adhikari and 1535 Patterson et al., which are the closest to ours but concentrate on a centralised 1536 interpretation of serverless function distribution, the aforementioned studies 1537 focus primarily on computation offloading, task scheduling, and orchestration 1538 in MANET and FANET networks. Hence, none specifically address the de-1539 sign of a fully decentralised, topology-aware serverless execution model that 1540 integrates function scheduling with the constraints of a mobile ad-hoc envi-1541 ronment. Our approach introduces a serverless deployment-aware protocol 1542 that abstracts the complexity of MANETs while ensuring efficient function 1543 placement and execution through a combination of gossip-based knowledge 1544 dissemination and adaptive function allocation. 1545

Furthermore, our work introduces scheduling mechanisms that support 1546 both hierarchical and progressive deployment models, enabling flexible adap-1547 tation to different network conditions. By integrating these techniques, we 1548 present a first-of-its-kind solution for serverless computing in ad-hoc net-1549 works, addressing key challenges in network topology awareness, function 1550 invocation reliability, and execution efficiency. Unlike previous approaches, 1551 our work is the first to apply serverless computing to an ad-hoc context where 1552 nodes function both as execution hosts and as schedulers, enabling a fully 1553 decentralised and self-organising computing paradigm. 1554

To provide a broader overview of the positioning of our work, we present a comparative analysis with the main related studies discussed in the literature. We organise the comparison into two complementary tables.

¹⁵⁵⁸ The first table, reported in Table 4, focuses on the serverless comput-

Warles	Decentralised	Multi-node	Execution	Custom
WORKS	Scheduling	Execution	Offloading	Policies
Grasso et al. [37]	•	0	•	0
Grigoropoulos and Lalis [42]	0	0	0	•
Khan et al. $[41]$	0	0	Ð	0
Luo et al. [39]	0	0	0	0
Niu et al. $[40]$	0	0	0	0
Patterson et al. [34]	0	0	0	•
You et al. [38]	0	0	•	0
This Work	•	•	•	٠

Table 4: Comparison of our work (last row) with related one from the literature on serverless.

ing perspective, summarising the key contributions and limitations of the 1559 works targeting function scheduling and orchestration in distributed envi-1560 ronments. The comparison highlights the lack of solutions that specifically 1561 address the ad-hoc and decentralised execution context that characterises our 1562 proposal. Specifically, starting from the left-most column, we find "Decen-1563 tralised Scheduling", which is the main characteristic of our system, i.e., the 1564 fact that scheduling happens in a decentralised way, following the multi-hop 1565 protocol described in Section 3.4. Among these works, only one adopts a 1566 decentralised scheduling model, which is the solution presented by Grasso et 1567 al. [37]. Another attribute that distinguishes our work is "Multi-node Execu-1568 tion", which corresponds to a single user request triggering the execution of a 1569 function across multiple nodes. Ours is the only alternative that supports one 1570 such feature. However, Khan et al. [41] partially cover this feature by allow-1571 ing a task to be split into multiple subtasks processed in parallel by different 1572 nodes. "Execution Offloading" regards the possibility for the execution nodes 1573 (the drones) to delegate function executions to other systems. The propos-1574 als by You et al. [38] and Grasso et al. [37] also support this feature, while 1575 Khan et al. [41] implement a restricted version of it, limited to offloading to 1576 the edge. Finally, "Custom Policy" is the support for user-defined, custom 1577 scheduling policies (which our system offers through AHAPP). Patterson et 1578 al. [34] support this feature via HiveMind and Grigoropoulos and Lalis [42] 1579 enable users to specify custom policies through their Kubernetes-based setup 1580 and the custom Fractus scheduler (both centralised). In summary, our sys-1581 tem is the only one to support all five capabilities in a cohesive, decentralised 1582 architecture. 1583

Works	Ad-Hoc Network Support	Worker-Agnostic Deployment	Support for Dynamic Topology	Generic Resource Awareness
Grasso et al. [37]	O	0	Ð	Ð
Grigoropoulos and Lalis [42]	O	O	Ð	Ð
Keshavamurthy et al. [43]	Ð	0	O	Ð
Khan et al. [41]	Ð	0	O	0
Lin et al. [44]	Ð	0	O	Ð
Luo et al. [39]	O	0	Ð	Ð
Niu et al. [40]	O	0	Ð	Ð
Qin et al. [45]	Ð	0	Ð	•
Secinti et al. [46]	Ð	0	O	Ð
You et al. [38]	Ð	0	Ð	Ð
This Work	•	•	•	•

Table 5: Comparison of our work (last row) with related one from the literature on network and ad-hoc systems.

The second table, reported in Table 5, provides a comparative analysis 1584 from the network and ad-hoc system perspective. We consider key char-1585 acteristics that are crucial in the context of mobile ad-hoc networks and 1586 UAV-based systems: "Ad-hoc Network Support" indicates whether the sys-1587 tem is designed to operate over a decentralised, infrastructure-less network; 1588 "Worker-Agnostic Deployment" refers to the ability of the scheduling strategy 1589 to operate without requiring prior knowledge of the specific execution nodes 1590 at scheduling time — an essential property for progressive, fully decentralised 1591 deployments; "Support for Dynamic Topology" denotes the capability of the 1592 system to adapt to topology changes, node mobility, and unpredictable net-1593 work dynamics; finally, "Generic Resource Awareness" evaluates whether the 1594 system can take into account generic node characteristics like energy, po-1595 sition, and hardware capabilities during scheduling decisions, besides the 1596 traditional ones like load and memory occupancy. We deem all the consid-1597 ered related work only partially supporting ad-hoc networks (first column) 1598 because none of them supports multi-hop deployments (this aspect is comple-1599 mentary to the decentralised scheduling approach found in Table 4) — from 1600 the intersection between the first columns of the two tables, we note that, 1601 while Grasso et al. [37] support decentralised scheduling, they rely on 5G 1602 connectivity and they do not (need to) support multi-hop communications. 1603 Regarding "Worker-Agnostic Deployment", the only proposal that partially 1604 addresses this concern is by Grigoropoulos and Lalis [42], who, by parti-1605

tioning workers into homogeneous pools can abstract away from an exact 1606 knowledge of the workers, following a hierarchical scheduling logic that first 1607 selects the pool and then selects one member in that pool based on a compan-1608 ion worker-selection policy. None of the considered proposals have "Support 1609 for Dynamic Topology" because they take scheduling decisions based on a 1610 centralised view of the topology, which might substantially differ from the 1611 actual one — the topology might change drastically during a scheduling ses-1612 sion in highly dynamic scenarios. Lastly, only Qin et al. [45] give adequate 1613 support for "Generic Resource Awareness" since they handle task scheduling 1614 also considering energy and application-specific sensor values. 1615

1616 8. Conclusion

We present a novel proposal for serverless computing tailored for ad-1617 hoc drone networks, specifically designed to address the peculiar challenges 1618 posed by such dynamic and resource-constrained environments. Our ap-1619 proach leverages a two-layer network overlay architecture, combining a gos-1620 siping layer for efficient network management with a function scheduling layer 1621 that adapts to the varying conditions of drone swarm operations. Through 1622 the integration of hierarchical and progressive scheduling approaches, we en-1623 sure robust and flexible function execution, even in the face of intermittent 1624 connectivity and high latencies. 1625

Our experimental results demonstrate the effectiveness of our proposal. 1626 highlighting its ability to efficiently manage the execution of serverless func-1627 tions while reducing communication overhead and maintain high reliability. 1628 Specifically, our simulations show that progressive scheduling achieves nearly 1629 optimal deployment success (close to 100%) in highly dynamic scenarios with 1630 up to 50 drones, compared to approximately 25% for hierarchical schedul-1631 ing under similar conditions. Moreover, progressive scheduling significantly 1632 reduces network overhead, approximately halving the communication cost 1633 observed with hierarchical methods. These findings validate the potential 1634 of serverless computing to enhance the performance and adaptability of drone 1635 swarms, offering a promising avenue for future development in both civilian 1636 and commercial applications. 1637

Looking ahead, there are several exciting opportunities for further improving the system. For instance, we envision working on refining the AHAPP language to capture function-specific traits, e.g., linked to the performance of function execution on the disparate hardware of a given drone swarm. This idea comes from a dialect of APP, called cAPP [27, 23], which uses static
analysis to derive cost equations from the function's code and select the best
(cloud) node for executing that function given, e.g., recorded latencies to call
external services. In our settings, one can incorporate these cost estimations
w.r.t. the (recorded performance of the) different hardware found among the
components of the network and provide refined rankings for selecting the
most suitable allocation targets.

Beyond the refinement of the scheduling language, further research in 1649 ad-hoc networks can enhance function deployment. One avenue is adaptive 1650 knowledge dissemination, where the system dynamically selects the most 1651 suitable method—Full, Full with Changes, or Aggregated—based on topol-1652 ogy stability and network congestion. The current deployment-aware proto-1653 col considers only node capabilities, neglecting network conditions and link 1654 quality. These aspects are important in multi-hop networks, especially for 1655 functions that need to frequently push data outside the drone network (e.g., 1656 to report continuous updates to the user). Ignoring link reliability can lead 1657 to high-latency interactions and execution failures, highlighting the need for 1658 network-aware function placement. Another avenue is improving aggrega-1659 tion methods for Progressive deployment. The existing approach uses sum-1660 mary metrics, but further optimisation could involve probabilistic modelling, 1661 weighted aggregation based on node stability, or adaptive techniques that 1662 adjust dynamically to network variations. These refinements would enhance 1663 scheduling accuracy and efficiency in highly dynamic environments. Finally, 1664 improving resilience in dynamic topologies remains essential. While Progres-1665 sive scheduling adapts to mobility, predictive models leveraging reinforcement 1666 learning or distributed heuristics could anticipate topology changes, enabling 1667 pre-emptive function placement and reducing failures caused by connectivity 1668 disruptions. 1669

1670 Acknowledgements

The work has been partially supported by the PRIN 2022 project RAIN4C: "Reliable Aerial and satellIte Networks: joint Communication, Computation, Caching for Critical scenarios" (Id: 20227N3SPN, CUP: J53D23007020001), the research project FREEDA (CUP: I53D23003550006) funded by the framework PRIN 2022 (MUR, Italy), the French ANR project SmartCloud ANR-23-CE25-0012, and PNRR CN HPC - SPOKE 9 - Innovation Grant LEONARDO ¹⁶⁷⁷ - TASI - RTMER funded by the NextGenerationEU European initiative ¹⁶⁷⁸ through the MUR, Italy (CUP: J33C22001170001).

1679 **References**

1695

- L. Gupta, R. Jain, G. Vaszkun, Survey of important issues in UAV
 communication networks, IEEE Commun. Surv. Tutorials 18 (2) (2016)
 1123–1152. doi:10.1109/COMST.2015.2495297.
- ¹⁶⁸³ URL https://doi.org/10.1109/COMST.2015.2495297
- Y. Zeng, R. Zhang, T. J. Lim, Wireless communications with unmanned aerial vehicles: opportunities and challenges, IEEE Commun. Mag. 54 (5) (2016) 36–42. doi:10.1109/MCOM.2016.7470933.
- ¹⁶⁸⁷ URL https://doi.org/10.1109/MCOM.2016.7470933
- [3] P. Doherty, P. Rudol, A UAV search and rescue scenario with human body detection and geolocalization, in: M. A. Orgun, J. Thornton (Eds.), AI 2007: Advances in Artificial Intelligence, 20th Australian Joint Conference on Artificial Intelligence, Gold Coast, Australia, December 2-6, 2007, Proceedings, Vol. 4830 of Lecture Notes in Computer Science, Springer, 2007, pp. 1–13. doi:10.1007/978-3-540-76928-6\
 __1.
 - URL https://doi.org/10.1007/978-3-540-76928-6_1
- [4] S. A. F. Manssor, S. Sun, M. A. G. Al-Sadoon, S. Ali, Real-time human detection in thermal infrared imaging at night using enhanced tinyyolov3 network, J. Real Time Image Process. 19 (2) (2022) 261–274. doi:10.1007/S11554-021-01182-Z.
- ¹⁷⁰⁰ URL https://doi.org/10.1007/s11554-021-01182-z
- [5] E. Jonas, J. Schleier-Smith, V. Sreekanti, C.-C. Tsai, A. Khandelwal,
 Q. Pu, V. Shankar, J. Menezes Carreira, K. Krauth, N. Yadwadkar,
 J. Gonzalez, R. A. Popa, I. Stoica, D. A. Patterson, Cloud programming simplified: A berkeley view on serverless computing, Tech. Rep.
 UCB/EECS-2019-3, EECS Department, University of California, Berkeley (02 2019).
- [6] A. E. Yaacoub, L. Mottola, T. Voigt, P. Rümmer, Nerta: Enabling dynamic software updates in mobile robotics, in: A. Ferscha, M. C.
 Chan, S. S. Kanhere, R. R. V. Prasad (Eds.), Proceedings of the 2022

1710International Conference on Embedded Wireless Systems and Networks,1711EWSN 2022, Linz, Austria, October 3-5, 2022, Junction Publishing /1712ACM, 2022, pp. 120–125. doi:10.5555/3578948.3578959.

- URL https://dl.acm.org/doi/10.5555/3578948.3578959
- [7] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, G. Zavattaro, Webassembly at the edge: Benchmarking a serverless platform for private edge cloud systems, IEEE Internet Computing (01) (2024) 1–8.
 doi:10.1109/MIC.2024.3513035.
- [8] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, G. Zavattaro, Funless: Functions-as-a-service for private edge cloud systems, in: IEEE International Conference on Web Services, ICWS 2024, Shenzhen, China, July 7-13, 2024, IEEE, 2024, pp. 961–967. doi:10.1109/ICWS62655.
 2024.00114.
- [9] S. Hendrickson, S. Sturdevant, T. Harter, V. Venkataramani, A. C.
 Arpaci-Dusseau, R. H. Arpaci-Dusseau, Serverless computation with
 openlambda, in: 8th {USENIX} Workshop on Hot Topics in Cloud
 Computing (HotCloud 16), 2016.
- I1727 [10] J. Sampé, M. Sánchez-Artigas, P. García-López, G. París, Data-driven serverless functions for object storage, in: Middleware, Middleware '17, ACM, 2017, pp. 121–133. doi:10.1145/3135974.3135980.
 URL https://doi.org/10.1145/3135974.3135980
- [11] E. Oakes, L. Yang, D. Zhou, K. Houck, T. Harter, A. Arpaci-Dusseau,
 R. Arpaci-Dusseau, {SOCK}: Rapid task provisioning with {ServerlessOptimized} containers, in: 2018 USENIX Annual Technical Conference
 (USENIX ATC 18), 2018, pp. 57–70.
- [12] M. Shahrad, J. Balkind, D. Wentzlaff, Architectural implications of
 function-as-a-service computing, in: Proc. of MICRO, 2019, pp. 1063–
 1075.
- [13] K. Solaiman, M. A. Adnan, Wlec: A not so cold architecture to mitigate
 cold start problem in serverless computing, in: 2020 IEEE International
 Conference on Cloud Engineering (IC2E), 2020, pp. 144–153. doi:10.
 1109/IC2E48712.2020.00022.

- [14] A. Banaei, M. Sharifi, Etas: predictive scheduling of functions on worker
 nodes of apache openwhisk platform, The Journal of Supercomputing (9
 2021). doi:10.1007/s11227-021-04057-z.
- [15] C. P. Smith, A. Jindal, M. Chadha, M. Gerndt, S. Benedict, Fado:
 Faas functions and data orchestrator for multiple serverless edge-cloud
 clusters, in: 2022 IEEE 6th International Conference on Fog and Edge
 Computing (ICFEC), IEEE, 2022, pp. 17–25.
- [16] G. D. Palma, S. Giallorenzo, J. Mauro, G. Zavattaro, Allocation priority policies for serverless function-execution scheduling optimisation, in: E. Kafeza, B. Benatallah, F. Martinelli, H. Hacid, A. Bouguettaya, H. Motahari (Eds.), Service-Oriented Computing - 18th International Conference, ICSOC 2020, Dubai, United Arab Emirates, December 14-17, 2020, Proceedings, Vol. 12571 of Lecture Notes in Computer Science, Springer, 2020, pp. 416–430. doi:10.1007/978-3-030-65310-1_29.
- [17] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, G. Zavattaro, Affinity-aware serverless function scheduling, in: 22nd IEEE International Conference on Software Architecture, ICSA 2025, Odense, Denmark, March 31-April 4, 2025, IEEE, 2025.
- ¹⁷⁶⁰ [18] C. E. Perkins, Ad hoc networking, Pearson Education India, 2008.
- [19] F. Al-Turjman, M. Abujubbeh, A. Malekloo, L. Mostarda, Uavs assessment in software-defined iot networks: An overview, Comput. Commun.
 150 (2020) 519–536. doi:10.1016/J.COMCOM.2019.12.004.
- ¹⁷⁶⁴ URL https://doi.org/10.1016/j.comcom.2019.12.004
- [20] M. Mozaffari, W. Saad, M. Bennis, M. Debbah, Mobile unmanned aerial
 vehicles (uavs) for energy-efficient internet of things communications,
 IEEE Trans. Wirel. Commun. 16 (11) (2017) 7574–7589. doi:10.1109/
 TWC.2017.2751045.
- ¹⁷⁶⁹ URL https://doi.org/10.1109/TWC.2017.2751045
- [21] D. Shah, Network gossip algorithms, in: 2009 IEEE International Conference on Acoustics, Speech and Signal Processing, IEEE, 2009, pp.
 3673–3676.
- [22] X. Fang, B. Zhang, D. Yuan, Gossip-based asynchronous algorithms for
 distributed composite optimization, Neurocomputing 616 (2025) 128952.

[23] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, G. Zavattaro, 1775 Function-as-a-service allocation policies made formal, in: T. Margaria, 1776 B. Steffen (Eds.), Leveraging Applications of Formal Methods, Veri-1777 fication and Validation. REoCAS Colloquium in Honor of Rocco De 1778 Nicola - 12th International Symposium, ISoLA 2024, Crete, Greece, Oc-1779 tober 27-31, 2024, Proceedings, Part I, Vol. 15219 of Lecture Notes 1780 in Computer Science, Springer, 2024, pp. 306–321. doi:10.1007/ 1781 978-3-031-73709-1_19. 1782

[24] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, G. Zavattaro, 1783 An openwhisk extension for topology-aware allocation priority poli-1784 cies, in: I. Castellani, F. Tiezzi (Eds.), Coordination Models and Lan-1785 guages - 26th IFIP WG 6.1 International Conference, COORDINA-1786 TION 2024, Held as Part of the 19th International Federated Confer-1787 ence on Distributed Computing Techniques, DisCoTec 2024, Gronin-1788 gen, The Netherlands, June 17-21, 2024, Proceedings, Vol. 14676 of 1789 Lecture Notes in Computer Science, Springer, 2024, pp. 201–218. doi: 1790 10.1007/978-3-031-62697-5_11. 1791

- [25] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, G. Zavattaro, Formally verifying function scheduling properties in serverless applications,
 IT Prof. 25 (6) (2023) 94–99. doi:10.1109/MITP.2023.3333071.
- [26] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, G. Zavattaro, 1795 Custom serverless function scheduling policies: An APP tutorial, in: 1796 G. Dorai, M. Gabbrielli, G. Manzonetto, A. Osmani, M. Prandini, 1797 G. Zavattaro, O. Zimmermann (Eds.), Joint Post-proceedings of the 1798 Third and Fourth International Conference on Microservices, Microser-1799 vices 2020/2022, May 10-12, 2022, Paris, France, Vol. 111 of OASIcs, 1800 Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 5:1–5:16. 1801 doi:10.4230/OASICS.MICROSERVICES.2020-2022.5. 1802
- [27] G. De Palma, S. Giallorenzo, C. Laneve, J. Mauro, M. Trentin, G. Zavattaro, Serverless scheduling policies based on cost analysis, in: M. H. ter Beek, C. Dubslaff (Eds.), Proceedings of the First Workshop on Trends in Configurable Systems Analysis, TiCSA@ETAPS 2023, Paris, France, 23rd April 2023, Vol. 392 of EPTCS, 2023, pp. 40–52. doi: 10.4204/EPTCS.392.3.

- [28] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, G. Zavattaro, A declarative approach to topology-aware serverless function-execution scheduling, in: C. A. Ardagna, N. L. Atukorala, B. Benatallah, A. Bouguettaya, F. Casati, C. K. Chang, R. N. Chang, E. Damiani, C. G. Guegan, R. Ward, F. Xhafa, X. Xu, J. Zhang (Eds.), IEEE International Conference on Web Services, ICWS 2022, Barcelona, Spain, July 10-16,
- ¹⁸¹⁵ 2022, IEEE, 2022, pp. 337–342. doi:10.1109/ICWS55610.2022.00056.
- ¹⁸¹⁶ [29] YAML, YAML specification, https://yaml.org/spec/ (11 2022).
- [30] N. Kumar, S. Misra, M. S. Obaidat, Collaborative learning automatabased routing for rescue operations in dense urban regions using vehicular sensor networks, IEEE Syst. J. 9 (3) (2015) 1081–1090. doi: 10.1109/JSYST.2014.2335451.
- ¹⁸²¹ URL https://doi.org/10.1109/JSYST.2014.2335451
- [31] S. Liu, C. Dong, X. Zhu, L. Zhang, et al., Performance evaluation of batman-adv protocol on convergecast communication in uav networks, in: GLOBECOM 2022-2022 IEEE Global Communications Conference, IEEE, 2022, pp. 5105-5110.
- [32] B. Karp, H.-T. Kung, Gpsr: Greedy perimeter stateless routing for wireless networks, in: Proceedings of the 6th annual international conference
 on Mobile computing and networking, 2000, pp. 243–254.
- [33] N. Singh, M. Adhikari, Popfl: A scalable federated learning model in
 serverless edge computing integrating with dynamic pop-up network,
 Ad Hoc Networks 169 (2025) 103728.
- [34] L. Patterson, D. Pigorovsky, B. Dempsey, N. Lazarev, A. Shah, C. Steinhoff, A. Bruno, J. Hu, C. Delimitrou, Hivemind: a hardware-software system stack for serverless edge swarms, in: V. Salapura, M. Zahran, F. Chong, L. Tang (Eds.), ISCA '22: The 49th Annual International Symposium on Computer Architecture, New York, New York, USA, June 18 22, 2022, ACM, 2022, pp. 800–816. doi:10.1145/3470496.
 3527407.
- ¹⁸³⁹ URL https://doi.org/10.1145/3470496.3527407

[35] A. Jindal, M. Gerndt, M. Chadha, V. Podolskiy, P. Chen, Function delivery network: Extending serverless computing for heterogeneous platforms, Software: Practice and Experience 51 (9) (2021) 1936–1963.

- [36] A. Garbugli, A. Sabbioni, A. Corradi, P. Bellavista, Tempos: Qos management middleware for edge cloud computing faas in the internet of things, IEEE Access 10 (2022) 49114–49127. doi:10.1109/ACCESS.
 2022.3173434.
- [37] C. Grasso, R. Raftopoulos, G. Schembra, Deep q-learning for job offloading orchestration in a fleet of mec uavs in 5g environments, in: 2021
 IEEE 7th International Conference on Network Softwarization (Net-Soft), IEEE, 2021, pp. 186–190.
- [38] W. You, C. Dong, Q. Wu, Y. Qu, Y. Wu, R. He, Joint task scheduling,
 resource allocation, and uav trajectory under clustering for fanets, China
 Communications 19 (1) (2022) 104–118.
- [39] Y. Luo, W. Ding, B. Zhang, Optimization of task scheduling and dynamic service strategy for multi-uav-enabled mobile-edge computing system, IEEE Transactions on Cognitive Communications and Networking
 7 (3) (2021) 970–984.
- [40] Z. Niu, H. Liu, X. Lin, J. Du, Task scheduling with uav-assisted dispersed computing for disaster scenario, IEEE Systems Journal 16 (4) (2022) 6429–6440.
- [41] M. A. Khan, I. Ullah, S. Nisar, F. Noor, I. M. Qureshi, F. Khanzada,
 H. Khattak, M. A. Aziz, Multiaccess edge computing empowered flying
 ad hoc networks with secure deployment using identity-based generalized
 signcryption, Mobile Information Systems 2020 (1) (2020) 8861947.
- [42] N. Grigoropoulos, S. Lalis, Fractus: Orchestration of distributed applications in the drone-edge-cloud continuum, in: 2022 IEEE 46th Annual Computers, Software, and Applications Conference (COMPSAC),
 IEEE, 2022, pp. 838–848.
- [43] B. Keshavamurthy, M. A. Bliss, N. Michelusi, Maestro-x: Distributed
 orchestration of rotary-wing uav-relay swarms, IEEE Transactions on
 Cognitive Communications and Networking 9 (3) (2023) 794–810.
- [44] X. Lin, A. Liu, C. Han, X. Liang, K. Pan, Z. Gao, Leo satellite and uavs
 assisted mobile edge computing for tactical ad-hoc network: A game
 theory approach, IEEE Internet of Things Journal (2023).

- [45] Z. Qin, H. Wang, Z. Wei, Y. Qu, F. Xiong, H. Dai, T. Wu, Task selection
 and scheduling in uav-enabled mec for reconnaissance with time-varying
 priorities, IEEE Internet of Things Journal 8 (24) (2021) 17290–17307.
- [46] G. Secinti, A. Trotta, S. Mohanti, M. Di Felice, K. R. Chowdhury, Focus:
 Fog computing in uas software-defined mesh networks, IEEE Transactions on Intelligent Transportation Systems 21 (6) (2019) 2664–2674.
- [47] C. N. Ververidis, G. C. Polyzos, Service discovery for mobile ad hoc
 networks: a survey of issues and techniques, IEEE Communications
 Surveys & Tutorials 10 (3) (2008) 30–45.
- [48] A. J. Ferrer, J. M. Marquès, J. Jorba, Towards the decentralised cloud:
 Survey on approaches and challenges for mobile, ad hoc, and edge computing, ACM Computing Surveys (CSUR) 51 (6) (2019) 1–36.