

RainCloud: Decentralized Coordination and Communication in Heterogeneous IoT Swarms

Filip Loisel

DSG@TU Wien, AUXO GmbH

Email: filip.loisel@auxo-software.com

Geri Zeqo

DSG@TU Wien, AUXO GmbH

Email: name.surname@auxo-software.com

Andrea Morichetta*

Distributed Systems Group, TU Wien

Email: a.morichetta@dsg.tuwien.ac.at

Anna Lackinger

Distributed Systems Group, TU Wien

Email: a.lackinger@dsg.tuwien.ac.at

Schahram Dustdar

Distributed Systems Group, TU Wien

Email: dustdar@dsg.tuwien.ac.at

Abstract—The increasing volume and complexity of IoT systems demand a transition from the cloud-centric model to a decentralized IoT architecture in the so called Computing Continuum, with no or minimal reliance on central servers. This paradigm shift, however, raises novel research concerns for decentralized coordination, calling for accurate policies. However, building such strategies is not trivial. Our work aims to relieve the DevOps engineers from this concern and propose a solution for autonomous, decentralized task allocation at runtime for IoT systems. To this end, we present a semantic communication approach and an ad-hoc lightweight coordination strategy based on Ant Colony Optimization (ACO). We compare the ACO strategy with Random Search and Gossip protocol-based algorithms. We conduct accurate experiments with up to a hundred nodes in both a static and a dynamic environment, i.e., with device outages. We show that ACO finds a matching node with the smallest hops and messages sent. While the Gossip strategy can allocate the most tasks successfully, ACO scales better, thus being a promising candidate for decentralized task coordination in IoT clusters.

Index Terms—Decentralized IoT, Ant-Colony Optimization, Decentralized Coordination, Communication, IoT Computing

I. INTRODUCTION

Nowadays, IoT devices pervade several areas, from healthcare to urban scenarios, industrial settings. It is predicted that in 2025 41.6 billion devices will be in use.¹ This sheer number of IoT devices and the massive amount of data generated poses significant challenges to the current cloud-centric [1?] IoT model: a centralized architecture where requests and data flow from IoT devices to the cloud is becoming a bottleneck, making it challenging to handle latency, increased required bandwidth, vulnerability to failures, privacy concerns, and inefficient use of resources. Therefore, a paradigm shift must happen. As an alternative, Edge and Fog computing [2] can provide relief to the central server, especially if collaborating in a unified computing continuum [4?]. Furthermore, IoT and Edge devices can be more self-reliant, i.e., self-managing entities that make autonomous decisions and using local resources.

¹<https://blogs.idc.com/2021/01/06/future-of-industry-ecosystems-shared-data-and-insights/>




Fig. 1: High-level representation of RainCloud's scenario.

Still, failures can often occur; furthermore, in mobility, IoT nodes can arbitrarily leave and join the system. Therefore, IoT swarms [5] are a viable alternative to the cloud-centric paradigm. Creating swarms for device collaboration means solving an optimization problem aiming at finding the best paths and edges across nodes. Typically, finding an optimal solution is complex; this task becomes even more challenging in real-world IoT environments with real-time constraints due to the heterogeneous and often battery-powered devices, the diverse requirements of the applications, which might not be supported by some devices, and the varying conditions of the network environment [6].

This paper addresses this challenge by examining the feasibility of a lightweight and real-time decentralized task coordination and communication mechanism for IoT devices. First, we design a node discovery and communication system that emphasizes the adaptation to the dynamic nature of a decentralized IoT environment. We achieve this by developing a framework, Rain Cloud System (RCS) that continuously adapts the network through constant checks, following the self-actualization paradigm [7]. This framework leverages semantic metadata for message exchange and can adapt to different specific communication protocols. Second, we put forward a lightweight metaheuristic optimization algorithm based on Ant Colony Optimization (ACO) for task coordination, capable of finding effective paths and offering an efficient load distribu-

tion across the swarm by limiting the number of messages exchanged and the serving time. Furthermore, we focus the attention on providing a framework that can actually be used by infrastructure and system managers. Figure 1 illustrates these functionalities. The RainCloud System (RCS) facilitates the formation and updating of swarms in dynamic environments by continuously evaluating the quality of connections and optimal collaboration settings. Bolder and greener paths represent the most reliable edges, both in terms of connectivity and task offloading. Leveraging multiplatform technologies and languages, we provide an interoperable framework for decentralized coordination. Therefore, through this paper, we wish to contribute meaningfully to decentralized task coordination in heterogeneous and dynamic IoT swarms, offering insights into efficient and reliable task coordination strategies and effective communication mechanisms. Furthermore, we are focusing on communication dynamics. Combined, we provide a holistic view and tools, including monitoring, for letting IoT devices function collaboratively in decentralized swarms. Our contributions are as follows:

- We present a communication framework for decentralized coordination based on semantic information and self-actualization.
- We introduce a lightweight metaheuristic for swarm offloading based on Ant Colony concepts.
- The conduction and evaluation through a realistic simulation with up to 100 devices in a static environment and in a dynamic environment with device failures, offering a detailed comparison of our ACO-inspired coordination strategy with two other baselines: Random and Gossips.
- We offer a turnkey open-sourced Rain Cloud System (RCS)² framework for edge-based device connectivity.

II. RELATED WORK

In this Section, we explore the state of the art for decentralized task offloading in IoT scenarios with a focus on how other contributions handled communication in a decentralized system and we explore the use of variations of Ant Colony models for handling task offloading. We clarify how our approach goes beyond the state of the art by combining together these aspects in a unique, ready-to-deploy framework and draws away from previous research.

A. Communication in decentralized IoT

Communication plays an essential role in decentralized systems. Much effort has been put into exploring how to transmit information and what to consider as informative. The work of Bittman et al. [8] focuses on the technology, offering a new Remote Procedure Call (RPC) approach in decentralized environments. The authors propose shifting from location-centric abstractions like RPC to data-centric abstractions akin to distributed shared memory (DSM) for better

module composition. The work of Wang et al. [9] also goes in the direction of RPC. It focuses on extending RPC with a shared address space and first-class references to promote interoperability and reduce task duplication. However, this approach shares the limit of the classic RPC model. Kountouris et al. [10] propose a paradigm shift focused on the semantics of information, unifying information generation, transmission, and reconstruction. Semantic information spreading became the primary basis and paradigm for the RCS Framework Communication module with metadata information spreading. In particular, efforts toward semantic communication [11, 12, 13] seem to pave the path for richer communication patterns.

B. Ant-Colony Optimization in decentralized coordination

In the context of large and dynamic environments, such as IoT, Kumar et al. [14] propose a comparison between ACO and K-Means clustering algorithms for IoT job scheduling. These two approaches are studied to find the shortest route path, optimize QoS constraints, and reduce energy consumption. Wang et al. [15] present the Collection Path Ant Colony Optimization (CPACO) method, which improves upon traditional parallel computing task allocation methods and ACO algorithms. They update the strategy in the Ant-Cycle Model, creating a three-dimensional path pheromone storage space. Similarly, Zannou et al. [16] use ACO for task allocation, considering two parameters: the optimal path length and the nodes' capabilities. In distributed systems, ACO demonstrates the potential for improving QoS. Hussein and Mousa [17] focus on reducing latency for delay-sensitive applications in IoT-Fog systems. In a similar fashion, Kishor and Chakarbarty [18] propose a metaheuristic scheduler called Smart Ant Colony Optimization (SACO) to offload tasks in a fog environment. Latency is a key requirement for extremely dynamic applications, such as autonomous driving. Bui and Jung's study [19] explores the potential of ACO for dynamic decision-making in connected vehicles within an IoT environment through swarm intelligence. Tan et al. [20] employ a bi-objective optimization approach with a fixed number of ants and a primary focus on energy efficiency and task completion time. Similarly, Wang et al. [21] introduce a sophisticated bi-ACO framework, combining task offloading with UAV trajectory planning and using multiple heterogeneous colonies. Xu et al. [22] present the GA-ACO Fusion Algorithm, proposing an algorithm that combines GA and ACO. This approach focuses on delay and energy efficiency. Li et al. [23] explore another combination, employing quantum computing principles to enhance the traditional ACO; through quantum bits, it expands the search space without increasing the number of ants. A quantum-boosted model is also presented by Dong et al. [24] even if without explicit ACO-based techniques.

C. Summary

Unlike the other approaches, our ACO implementation offers a fast and lightweight solution for real-time optimization of multiple variables, as well as efficient queue size management. This approach is suitable for real-world and

²<https://github.com/auxo-gmbh/rcs-flooding-scenario>
<https://github.com/auxo-gmbh/rcs-documentation>
<https://github.com/auxo-gmbh/rcs-monitoring>
<https://github.com/auxo-gmbh/rcs-rain-cloud-system>

dynamic IoT and edge computing scenarios, thanks to its fast optimization through reducing the number of ants and carrying the workload during the exploration. Furthermore, we propose a communication approach that, through self-actualization methods, dynamically adapts to the state of the infrastructure. Finally, we offer all of this in an integrated and flexible framework.

III. METHODOLOGY

In the following, we present our framework for decentralized task offloading in IoT swarms. After presenting the structure of the system we address in our research, we delineate how to let the nodes connect and adjust over time through a semantic and self-actualization-based communication pattern. Furthermore, we present the metaheuristic for the optimization strategy, based on Ant Colony Systems.

A. Use Case Introduction

Device type	Task types	Queue	C_{comm}	C_{comp}	$C_{storage}$	$P(Fail)$	$P(Recover)$
s_1	T_1, T_2	10	1	1	1	0.1	0.4
s_2	T_3, T_4, T_5	15	1	2	1	0.1	0.4
m_1	T_1, T_4	10	1	2	2	0.2	0.3
m_2	T_5	5	2	2	2	0.2	0.3
m_3	T_2, T_3	10	2	2	3	0.2	0.3
w_1	T_1	5	3	3	3	0.3	0.2
w_2	T_2	5	3	2	3	0.3	0.2
w_3	T_3	5	3	3	3	0.3	0.2
w_4	T_4	5	3	3	3	0.3	0.2
w_5	T_5	5	2	3	3	0.3	0.2

TABLE I: Device types and their characteristics.

We expect our framework to work in the case of heterogeneous IoT devices equipped with different processing, storage, and communication capabilities, which serve different kinds of tasks. Each IoT node periodically produces tasks and, in a stable scenario, it takes care of executing them. However, a node may become overloaded and require assistance, rendering task delegation to other swarm nodes crucial. In our use case, the system has three resource groups (computation (C_{comp}), communication (C_{comm}), and storage ($C_{storage}$)), which can have values between 1 and 3. A *lower value* means that the *device is better* equipped. I.e., the strongest devices in terms of hardware are those where all three values are set to 1. The resources directly influence the processing time of the tasks in the system; Table I summarizes these properties.

All the devices in the swarm are connected and can communicate. Therefore, we leverage this scenario to implement autonomous coordination between devices to offload tasks. In the development of the framework and of the simulation, we consider the following assumptions: the first is **heterogeneity**, i.e., the members of an IoT swarm have different hardware specifications (*Device type*), queue sizes ($|Queue|$), and emitter types. The *Device types* include strong devices (s), mid-strong devices (m), and resource-constrained devices, which are the weakest with respect to resources in the swarm (w). Different IoT devices support different tasks (*Task types*), and the processing time of a task varies depending on the device's capabilities. Furthermore, some devices are more error-prone than others, which affects the reliability of successful task




Fig. 2: Main components and interaction for the RCS container responsible for a node.

processing. Secondly, we include **processing and communication overhead**. Most IoT devices are resource-constrained; therefore, task coordination must avoid negatively impacting the performance of the devices or overflowing the network with messages. Third, **scalability** is essential. As there is virtually no upper limit to the size of distributed IoT devices in a swarm, successful task coordination algorithms must perform at increasing scales. Finally, our approach is aware of dealing with a **dynamic environment**. Task coordination must assume that IoT devices might fail or join the system at any time. Hence, coordination must learn the topology and adapt it at runtime.

B. Rain Cloud System (RCS) Architecture

Figure 2 depicts the main components and flow for our RCS system, in particular for one node. In detail, we envision the stakeholder of our framework to be a DevOps engineer, responsible for configuring and managing the IoT devices and the running software. Therefore, the framework's responsibility is to ensure the machine's connectivity. The DevOps engineer manages the IoT swarms by adding new nodes, deleting old ones, or generally specifying connectivity rules, allowing for infrastructure flexibility. In this setting, our framework provides task coordination and communication tools. Tasks are generated by emitters, and each node decides, based on its task queue capacity, whether to process the task locally or forward it to other IoT devices selected according to a coordination strategy. An example of an IoT device in a smart city is a humidity sensor that provides data to an irrigation application. This application example motivated the name of our Rain Cloud System. After processing this data, the application can use it to determine whether the park's sprinkler system should be activated. At this point, the application leverages the framework to communicate with other IoT devices and offload tasks when needed. The application, in turn, transmits and exchanges monitoring data, to evaluate its own state. To offload a task, the nodes must first connect to other devices in a swarm to exchange their node profiles, which

summarize device and task properties. These node profiles are essential for obtaining information about neighboring devices and new members, which is crucial for task coordination as they reveal the available resources within the swarm. When a task is offloaded, the application appends a Time-To-Live (TTL) (expressed in seconds) to the request to indicate the maximum waiting time. Once the TTL has expired and no device has been found, the original node executes the task. The goal is to maintain a balanced task queue, avoiding both underutilization and overutilization and finding a node that can execute an offloaded task with a minimum number of hops.

C. Communication

The communication module facilitates interaction among devices within the swarm. We establish various protocols to enable efficient intra-swarm communication, employing TCP sockets for message exchange, and JSON as the message format. This interaction is not only essential for exchanging offloading information at runtime, but also for creating and updating the swarm. We handle these two aspects through the Node Discovery and Self-actualization phases.

1) *Node Discovery*: The Node Discovery (ND) Protocol plays a vital role in the adaptability of the RainCloud System Framework. Through the identification and connection of devices, the ND Protocol allows the implementation of a decentralized and interconnected swarm ecosystem. During the initial handshake phase, nodes exchange profiles. A profile carries a device's crucial information, such as its processing capabilities, memory capacity, available sensors, communication protocols, and supported services. Exchanging profiles during the handshake phase enables information sharing, essential for collaboration and task allocation. Devices within the swarm can manage the number of open connections they maintain with other members and specify the particular ports they use for establishing these connections. This level of customization ensures the optimization of nodes' performance and resource usage while balancing redundancy and scalability. This flexibility allows the swarm to adapt at runtime according to the environment dynamics. We design ND so that it relies on semantic information, whereby the protocol draws only the necessary metadata from the devices.³ This approach enables other protocols to adapt based on the information's meaning, allowing for a more efficient and dynamic data exchange between devices. By focusing on the semantic aspects of device information, the ND Protocol fosters a more intelligent and adaptable IoT ecosystem capable of responding to varying conditions and requirements. All protocols are designed to work asynchronously.

2) *Self-actualization process*: Self-actualization (SA) is a concept first introduced by the psychologist Abraham Maslow [7] in the 1950s. It refers to realizing and fulfilling one's potential, representing the highest level of psychological development and the ultimate goal of personal growth. We use

³Semantic information contains metadata as processing capabilities, memory capacity, available sensors, communication protocols, and supported services




Fig. 3: Concept of forward and backward ants in ACO. The forward ant is depicted on the left-hand side, going from the origin to the target, and the backward ant is on the right.

this concept to design a mechanism for dynamically adjusting the swarm topology and enhancing task coordination performance. The task coordination strategies store a specific value for each outgoing connection to other nodes, which indicates how attractive this connection is. For the ACO algorithm, this specific value is the pheromones that are collected over time on a connection variable related to the task type, and for the Gossips algorithm, this value is the number of messages based on the task type for the connection. SA takes this local knowledge about the swarm and calculates the weakest links. This computation occurs at each node and indicates which links should be disconnected. However, a task might be still offloaded to one of these nodes and, if the connection is terminated, the task might not return. For this reason, the SA process takes place in two phases. In the first stage, the nodes are informed that they should no longer contact specific nodes for offloading tasks because these nodes will soon disconnect. Then, each neighbor stores the detached devices in a block list. Here, connection requests from blocked devices are denied to prevent the same swarm topology from being constructed, forcing disconnected nodes to search for new neighbors. After some time, the marked devices are removed from the list.

D. ACO-inspired offloading mechanism

Here, we explore the proposed metaheuristic for decentralized task offloading. We base our approach on the Ant System (AS), a member of the ACO family of algorithms, bringing in some of its main characteristics for task offloading coordination. Specifically, we aim to minimize the interval between an offloading request and its completion.

1) *Novelty*: In our implementation, we first reduce the number of ants to only a pair per offload request, exploiting the concept of two different ant agents (see Figure 3), based on [25, 26], i.e., the *forward* and *backward* ants. More in detail, *forward ant* represents the offload request created by the origin device and it is responsible for finding the best path that leads to target nodes. Once it finds a suitable node for the task execution along its path, it offloads it there and terminates its role. The *backward ant* retraces the forward ant steps, updating pheromones based on a posteriori knowledge about the node's quality; in our ACO implementation, the *quality* is

a function of resource type and availability, plus queue length. We adopt this approach as it avoids the formation of loops [?] and keeps the process fast and lightweight. Furthermore, unlike traditional ACO and, to the best of our knowledge, any other approach, our forward ant *carries* the task and TTL. The goal is to minimize message exchange. Finally, we design the optimization process to last only one iteration, with the goal of improving scalability and reducing computation and communication overhead, accepting some performance loss in task distribution as a trade-off.

2) *Detailed Implementation*: Formally, in our implementation the transition rule is identical to the original AS's approach, i.e., based on a probabilistic decision-making procedure that considers both the pheromone levels and heuristic information associated with the possible transitions through local, a priori knowledge and the global, a posteriori knowledge, but adapted to our scenario. Like the traditional approach, we keep the balance between exploration and exploitation by adjusting the influence of these factors through parameters α and β , where α controls the influence of the a posteriori knowledge and β controls the influence of the a priori knowledge. In our case, we define the local knowledge η_{ij} as in Equation 1. It is exchanged during the Node Discovery (ND) protocol and indicates whether two neighboring nodes support the task type in the offload request. If so, a higher pheromone value defined by τ_{0+} is set. If not, the lower value τ_0 is used. The intention is to favor nodes that support the task type to reduce the offloading time and, thus, the latency and number of generated messages.

$$\eta_{ij} = \begin{cases} \tau_{0+}, & \text{if next node supports } T_{type} \\ \tau_0, & \text{otherwise} \end{cases} \quad (1)$$

In our approach we express the quality pheromones Q at the target node as the combination of the device's queue quality (Equation 3) and its available capacities (Equation 2). The device's queue quality Q_L is a function of the average queue occupation, L_{avg} , and the queue limit, L_{limit} , i.e., when the queue is overloaded. The queue quality is inversely proportional to the queue occupation. The second factor in the quality pheromones calculation is the device's capacities, i.e., communication (C_{comm}), computation (C_{comp}), and storage ($C_{storage}$).

$$Q = Q_L \cdot \frac{1}{\frac{1}{3}(C_{comm} + C_{comp} + C_{storage})} \quad (2)$$

$$Q_L = \begin{cases} 0, & \text{if } 1 - \frac{L_{avg}}{L_{limit}} \leq 0 \\ 1 - \frac{L_{avg}}{L_{limit}}, & \text{otherwise} \end{cases} \quad (3)$$

The node capable of processing the task then computes the quality pheromones (see Equation 2), which the backward ant then uses to update the pheromones in Equation 4. The evaporation of the pheromones does not happen as in AS when updating the pheromones but independently every n minutes. The backtrack ant updates the pheromones at each node using Equation 5 that combines the carried quality pheromones and

the distance from the current node to the source node, i.e., $Q_P = \frac{1}{hopsToOrigin}$.

$$\tau_{ij} \leftarrow \tau_{ij} + \Delta\tau_{ij} \quad (4)$$

$$\Delta\tau_{ij} = (Q \cdot Q_P) \cdot 100 \quad (5)$$

With this setting for the quality pheromones, we can capture the device's current availability and inherent capacities, leading to a more informed decision-making process in the ACO algorithm. The strategy aims to select target nodes with higher availability and better resources. Updating the pheromone tables at each node along the path is crucial for learning and guiding the search process. Evaporation prevents the algorithm from becoming overly biased towards previous solutions, enabling it to constantly explore new paths and encouraging a balance between exploration and exploitation. Each node in our system performs evaporation independently, removing a relative part of the pheromone values every n minutes. The amount of pheromone reduced is controlled by a parameter ρ .

We inspect the time complexity of the ACO-based task coordination strategy, considering different aspects of the algorithm, such as initialization of pheromones, quality pheromones calculation, applying state transition rules, updating of pheromones, and the evaporation process.

- **Pheromones Initialization.** When two devices are connected, the initialization involves updating the pheromone tables for the connected edges. This process has a time complexity of $O(n)$, where n is the number of connected edges.
- **Quality pheromones calculation.** The calculation of quality pheromones at the target node is a constant-time operation, resulting in a time complexity of $O(1)$.
- **State transition rule.** The selection of the next node depends on the state transition rule, which considers the connected edges and the supported task types of the selected node. The time complexity of this process is $O(n \cdot m)$, where n is the number of connected edges and m is the number of supported task types of the selected node.
- **Pheromones update.** Updating pheromones with the backtrack ant involves traversing the connected edges and updating the pheromone values accordingly. This results in a time complexity of $O(n)$, where n is the number of connected edges.
- **Evaporation.** The evaporation process involves updating the pheromone values for each connected edge, and each task type the node has heard of. This results in a time complexity of $O(n \cdot m)$, where n is the number of connected edges and m is the number of task types.

Given these adaptations of the AS for our system, we can present the pseudo-code of our ACO strategy. Algorithm 1 shows how the ants decide which node to visit next based on a *state transition rule*, which gives a score on the best

node. Once the node is selected, the task is computed, and the backward ant updates the pheromones. Finally, periodically, the pheromones evaporate to allow a strategy update.

Algorithm 1: Optimized Decentralized ACO Strategy

Data: Nodes $nodes$, Task t

```

// Concurrent activity at each node
1 foreach node  $n$  do
2   initializePheromoneTables()
3   while  $n$  is running do
4     if  $n$  is overloaded then
5       createForwardAnt( $t$ )
6       while  $t_{TTL} > 0$  do
7          $n_j = \text{applyStateTransitionRule}()$ 
8         goToNode( $n_j$ )
9         addNodeToPath( $n_j$ )
10        if  $n_j$  supports  $t_{type}$  and not overloaded
11          then
12            computeTask( $t$ )
13            createBackwardAnt( $\text{calculateQuality}()$ ,
14                                $path$ )
15            break
16          end
17        end
18      end
19      foreach backwardAnt do
20        while  $node \neq node_{origin}$  do
21           $node = \text{popNodeFromPath}(path)$ 
22          goToNode( $node$ )
23          applyPheromoneTrailUpdateRule()
24        end
25      end
26      if time for evaporation then
27        applyEvaporationRule()
28      end
29    end

```

IV. EVALUATION

This Section aims to evaluate the performance of the three task coordination algorithms Random, ACO, and Gossips, in the context of decentralized IoT systems. The focus of the evaluation is on considering how the ACO algorithm impacts task coordination’s performance compared to the remaining approaches. To this end, we run tests on servers on which VMs represented as IoT devices run the framework respectively. By running tests with different numbers of devices in each static and dynamic environment, we get a solid overview of how the strategies perform. In the static environment, the network topology remains unchanged throughout the experiment. On the other hand, the dynamic environment simulates decentralized IoT systems’ mobility and ever-changing nature. In this environment, devices can fail and rejoin the network. We compare the task coordination implementations through various

performance metrics with a focus on the load distribution of the queue, the service time, and the hops per hit.

A. Benchmarks

a) *Random search*: The Random strategy is a straightforward, and common [24, 27, 28] technique for decentralized task coordination in distributed IoT systems. Tasks are offloaded to randomly selected nodes without considering their specific characteristics or currently available resources. While its simplicity and ease of implementation guarantee a fast configuration for task offloading, this approach does not possess self-learning or self-adapting capabilities, limiting its ability to optimize task distribution and adapt to dynamic changes in the IoT environment. The *computational complexity* of selecting a random neighbor is $O(n)$, where n is the number of neighbors. Thus, the runtime of offloading or forwarding is also $O(n)$.

b) *Gossips*: Furthermore, we design and implement the Gossips approach, adapted from the Gossips protocol. It has two phases: the *Gossips Discovery Phase* and the *Gossips Offload Phase*. The goal of the first phase is to contact as many nodes as possible to find available targets. To this end, the origin node starts spreading messages, so-called Gossips discovery messages, to all its neighbors, which in turn propagate these messages to their neighbors. When a node receives a gossip discovery message and can process the task, it sends a discovery response along the path to the origin node. The first phase continues until the same gossip message is received by a node or the discovery TTL passed, determined by the origin node at the beginning. This deterministic approach aims to build extensive knowledge about the system and its available resources, at the same time quickly adapting to changes in the network. In the Gossips Offload Phase the origin node collects all received discovery response messages, which contain information about the current queue capacity and the on-board hardware resources of the target node. Then, this information plus the path length provide a metric for selecting the best target node. The complexity of the Gossips strategy can be broken down into two main components: spreading and forwarding gossip discovery messages and selecting the target node to offload the task.

- 1) **Spreading Gossips Discovery Messages.** The complexity of spreading and forwarding Gossips discovery messages can be represented as $O(n)$, where n is the number of connected edges at a node.
- 2) **Selecting the Target Node.** The selection process involves sorting the received response messages and choosing the highest-ranked node as the target for offloading the task. This sorting process has a complexity of $O(n \cdot \log(n))$, where n is the number of received gossip discovery response messages.

B. Scenarios

To thoroughly evaluate the performance of the task coordination strategies, we conducted experiments in two types of environments: *static* and *dynamic*. In the *static environment*,

the network topology remains unchanged throughout the experiment. This type of environment enables us to assess the performance of the strategies in a stable and controlled setting. On the other hand, in the *dynamic environment*, devices can fail and rejoin the network. Every 15 minutes, we calculate a probability to determine if a node will fail based on each node’s characteristics. If the probability is above a certain threshold, the device fails. For each “failed” device, every minute, we extract its probability to recover; if it is higher than a threshold, the device rejoins the network. These two environments ensure a thorough evaluation of how well the selected strategies can adapt to network changes and maintain performance. For each environment, we consider settings with different numbers of nodes. The experiments included 10, 25, 50, or 100 devices, sampling both robust nodes and resource-constrained devices (check Section III-A). This heterogeneity of devices types and cardinality helps us assess how well the strategies can handle diverse resource capacities in the network. For the processing time, we use previously collected application data. For the arrival rate of the tasks, we leverage the request generator from the tool Edge Runner⁴. Each task type is emitted at different times to increase the heterogeneity. T_1 and T_3 appear most often in the system and are short-lived, while T_4 and T_5 are the task types emitted least often but with a longer processing time. At any time, an emitter can send out precisely one task. Combined, however, the emitters can produce several tasks at once. Therefore, IoT devices may produce many tasks at once in a short period, resulting in a potential swarm overload that the offloading strategy should handle. We deploy a total of 100 VMs on the server machines, each of which is equipped with 3072 MB of RAM, 2 vCPU, and 20 GB of disk space.⁵

C. Performance metrics

To effectively compare the performance of the three task coordination strategies, we rely on a set of metrics emanated from related work [29, 30]. These metrics allow us to analyze various aspects of the strategies and are defined as follows:

- 1) **Load Distribution (LD)**. The average queue occupation measures load distribution. The metric reflects the system’s efficiency and aims balanced load, where the queue is neither empty nor too full.
- 2) **Service Time (ST)**. It measures when the origin node sends an offload request for a task until it receives a response back. I.e., how long does the origin have to wait until it can regard the offloaded task as finished? It is expressed in seconds.
- 3) **Hops per Hit (HPH)**. The metric measures the hops required to find the target node. It indicates how often another node must be contacted until the target node is found.

⁴<https://github.com/edgerun/request-generator>

⁵The virtual machines host a Rocky Linux 9.1 OS (minimal installation, <https://rockylinux.org/>) and the OpenJDK 11 Java Runtime Environment: OpenJDK 11 (<https://openjdk.org/>)

- 4) **Hit Miss Ratio (HMR)**. The ratio shows how often a target node was found divided by the number of times a target node was not found, but the offload request was lost (e.g., if TTL had already passed).
- 5) **Guarantee Ratio (GR)**. The ratio is calculated by dividing the number of offloaded tasks by the number of returned tasks (from the origin node to the target node and back to the origin node).
- 6) **Amount of Messages (AM)**. The metric counts the number of messages the active strategy produces while coordinating tasks during runtime.
- 7) **Messages per Request (MPR)**. The metric calculates the number of messages produced by the active strategy for a single offload request for a task.

D. Results

First, we evaluate the three strategies⁶ by analyzing Load Distribution (LD), Hops-per-Hit (HPH), and Messages per Request (MPR) in a 100-node environment. These selected metrics provide essential insights into the efficiency of the task relocation methods, the frequency of node contacts required to find the target node, and the number of messages generated by a strategy for a single task offload request. Subsequently, we present a comprehensive overview of the three strategies across all metrics described in the previous Subsection IV-C to offer a broader perspective on performance.

1) *Load Distribution (LD)*: The primary aim of LD is to assess the balancing of load across the system. A balanced LD ensures that tasks are offloaded to appropriate target nodes, preventing overload on specific nodes while others remain underutilized. We use the term queue capacity to describe how many tasks must be in the queue before subsequent tasks must be offloaded. Queue occupation describes the total number of tasks in the queue, including both local tasks and tasks that will be offloaded. If the queue occupation exceeds the capacity, we call the device overloaded. The average queue occupation in our system is the LD described above. We experiment with a varying number of nodes, in both static and dynamic settings. For the sake of brevity, we show in Figure 4 only the results for the case of 100 nodes. The rest of the values can be found in the summarizing tables in Section IV-D4. In Figure 4, we can observe how the static Random approach exhibits stronger deviations from the regression line and the queue occupation grows at a queue capacity of ten. With a queue capacity of 15, this trend is amplified: Random’s task coordination loses performance, and the prediction states that the devices are overloaded after one hour. Figure 5 emphasises this behavior. Especially for small queue sizes, the three strategies tend to have requests that overflow the queue size, calling for more accurate offloading mechanisms. This behavior is mitigated by larger queue sizes for ACO and Gossips, whereas Random Search still shows irregularities. We conclude that Random can only cope with smaller queue capacities for large swarm

⁶For ACO we have to specify some hyperparameter values. We use the following configuration: $\tau_0 = 50$, $\tau_{0+} = 100$, $\alpha = 2$, $\beta = 1$, $\rho = 0.3$.

sizes. Gossips shows the smallest LD values, while ACO has the highest values but never overloads the devices, which can be concluded as a solid task coordination mechanism.

In summary, ACO demonstrates the most effective utilization of queue resources across various scenarios and swarm sizes, likely due to its self-learning nature. Random shows that with large swarm sizes and larger queue capacities, it overloads the devices and thus degrades the overall system performance. The Gossips strategy shows the most stable values without the risk of overloading devices, but the queue utilization is in the lower range. In general, all three approaches have underutilization at the five and ten queue capacities, which can be improved in future works.

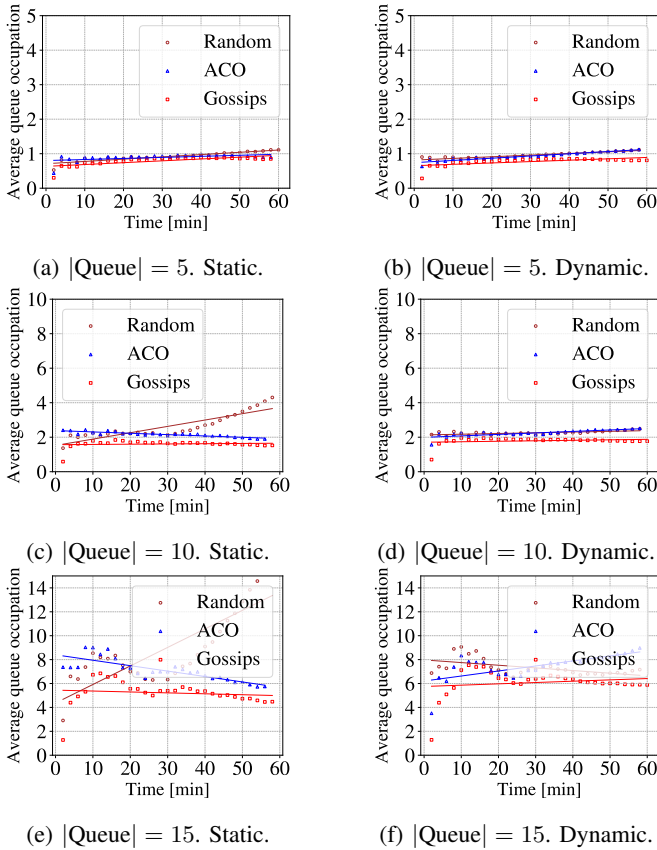


Fig. 4: LD in static (1st column) and dynamic (2nd column) environments with 100 devices of different queue capacities.

2) *Hops-per-Hit (HPH)*: Assessing the average HPH metric allows us to understand the efficiency and scalability of the Random, ACO, and Gossips strategies in the context of message propagation and task coordination. HPH measures the number of hops or nodes between the origin and target nodes during task offloading. A lower HPH value indicates fewer messages sent and a closer target node, leading to reduced latency, faster communication, and shorter service times. Figure 6 summarizes the test run with 100 nodes. Given a swarm size of 100, Random is the worst algorithm in terms of HPH, while ACO has the best mean value. Compared to the Gossips, Random and ACO have more outliers. In a static

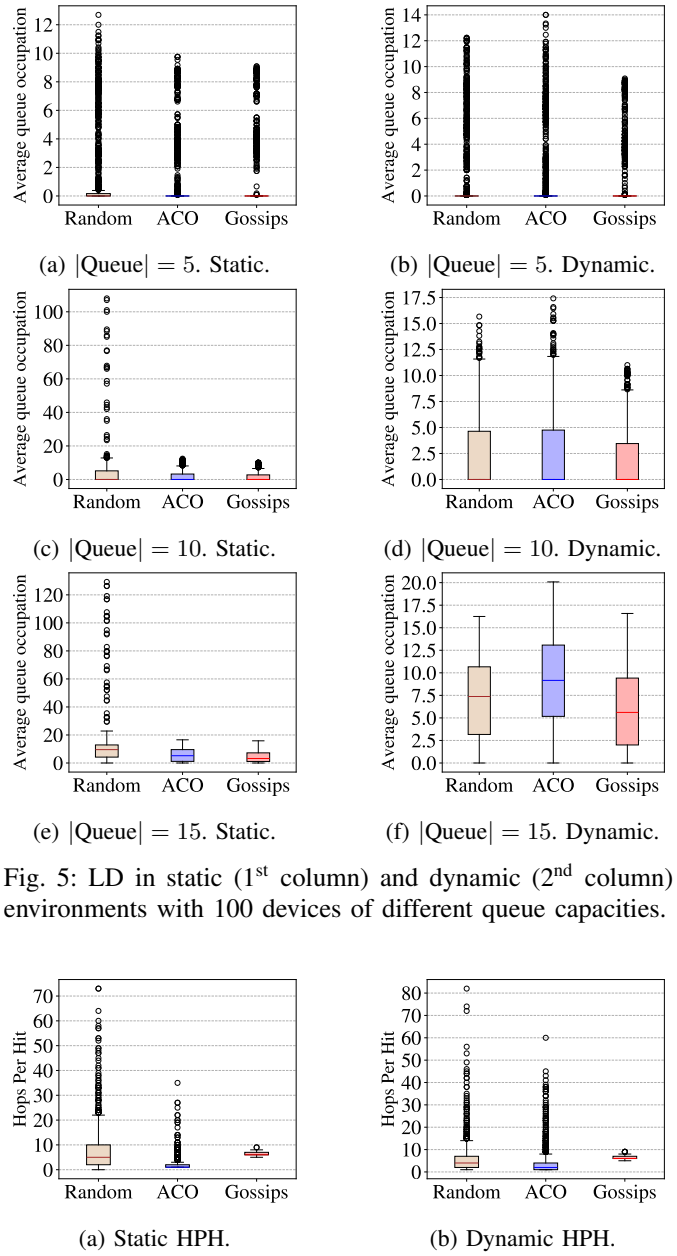


Fig. 5: LD in static (1st column) and dynamic (2nd column) environments with 100 devices of different queue capacities.

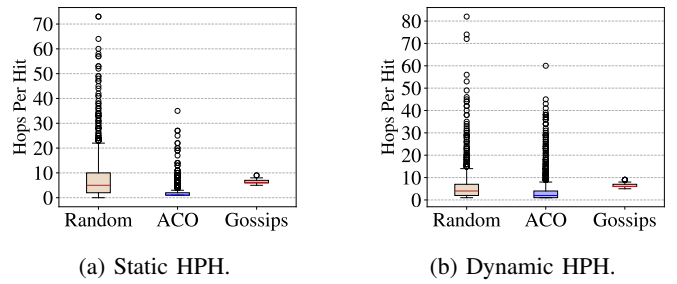


Fig. 6: HPH in static and dynamic environments with 100 devices.

environment, Random's highest value is even 75, meaning one-third of all nodes were traversed until the task offload request was accepted. Gossips shows stable and consistent values in both environments. In summary, the ACO strategy has the lowest and, thus, best HPH values in both environment setups, thanks to the deposition of pheromones. It would be interesting to know if the outliers are caused by the fact that the system needs more time to learn due to the larger number of nodes. I.e., the long paths (outliers) are formed at the beginning because not enough artificial ants have explored the network. Gossips scores well in this category and shows solid stability and consistency in static and dynamic environments, while

Random has the worst performance.




Fig. 7: Message per request over time in *static* (left) and *dynamic* (right) environments with 100 devices.

3) *Messages per requests (MPR)*: The MPR metric allows us to analyze the typical workload in RCS and to process its ability to deliver tasks to its destination. In Figure 7 we can analyze the results with a setup of 100 nodes. In this scenario, the ACO algorithm demonstrates its effectiveness in managing communication efficiency within distributed systems. In contrast, the Gossips algorithm exhibits instability in messages per request regard, and the Random plateaus at higher values.

4) *Summary*: An overview of the performance of the strategies in a static environment is provided in Table II, while the results for the dynamic environment are presented in Table III. They both report the average values for each metric, truncated after two comma digits. The winning strategy in the metric and swarm size category is highlighted in bold and green; the second place is colored orange, and the last place is red. For LD, weighting is based on whether the queue is neither too full nor too empty – strategies should target the occupation of 3/4 of the queue, given a capacity of 15. Lower values indicate better performance for ST, HPH, AM, and MPR. For the two metrics, HMR and GR, higher scores indicate better performance by the strategies.

Overall, the proposed ACO algorithm shows great potential for successfully coordinating nodes in a fully decentralized way. The proposed solution constantly shows the best performance in many of the used metrics. Random outperforms only slightly ACO on service time but has poor performance in terms of actual coordination, whereas Gossip shows better results for HMR and GR but requires much more Service Time and sends in the systems a much higher number of messages, with a serious risk of clogging the network. Furthermore, the test has only run for one hour; longer training times for ACO will intuitively improve its performance in HMR and GR, due to better exploration and refined pheromone trails, with the potential of making it solidly the preferred solution.

V. CONCLUSION

This paper introduced the Rain Cloud System (RCS) framework to facilitate decentralized swarm task offloading. In

Strategy	#Nodes	LD	ST	HPH	AM	MPR	HMR	GR
Random	10	4.46	64.06	4.28	713	5.40	45.45	45.45
	25	4.97	65.85	4.71	3215	7.29	48.29	48.29
	50	5.14	66.80	6.20	8865	11.40	46.07	46.07
	100	16.73	72.21	7.82	93818	31.50	35.75	35.65
ACO	10	5.03	62.79	1.34	300	2.37	57.26	57.26
	25	5.53	70.26	2.01	2309	4.75	50.84	50.84
	50	5.40	71.27	1.95	4664	5.75	47.15	47.15
	100	5.75	71.19	2.38	14243	8.93	48.69	48.69
Gossips	10	4.60	106.35	3.44	3900	33.42	69.36	69.36
	25	4.58	107.34	5.09	37400	114.01	58.07	58.07
	50	4.38	108.14	5.97	163549	296.12	53.11	53.11
	100	4.47	103.68	6.42	827354	757.37	58.02	58.02

TABLE II: Comparison of strategies in a static environment.

Strategy	#Nodes	LD	ST	HPH	AM	MPR	HMR	GR
Random	10	7.35	60.39	3.34	1082	5.07	43.19	43.19
	25	9.47	76.15	4.22	8220	8.40	42.33	41.30
	50	8.62	73.09	6.90	20100	13.74	45.55	44.39
	100	7.16	71.37	6.15	30211	13.70	48.79	47.02
ACO	10	4.97	63.78	1.11	295	2.36	52.94	52.94
	25	11.48	85.80	3.36	8354	7.99	40.22	38.79
	50	6.30	70.87	2.85	6693	7.00	47.06	46.75
	100	8.97	78.97	4.23	32160	12.03	47.64	46.44
Gossips	10	3.69	91.27	3.69	3259	32.47	78.16	74.71
	25	5.10	106.92	5.03	43816	121.22	58.60	58.60
	50	5.92	104.56	5.56	196590	304.28	51.20	51.04
	100	5.88	110.33	6.50	849374	720.20	52.32	52.23

TABLE III: Comparison of strategies in a dynamic environment.

particular, we developed and evaluated a semantic and self-actualizing communication strategy, coupled with a meta-heuristic based on the Ant Colony System and compared it with Random Search and Gossip, adapted for the use case. The Ant Colony Optimization (ACO) algorithm demonstrated promising results in network management and resource distribution, indicating its suitability for complex IoT challenges. ACO and the other algorithms embedded in the RCS framework allow the development of inventive strategies for resource distribution and task management within IoT networks.

In future work, we will incorporate path optimization for ACO, avoiding cycles. Furthermore, we aim to optimize the pheromones deposit during the initialization process, using more accurate rules. In general, the pheromone calculation could also take into account energy consumption or device geographical proximity. In addition, we want to explore a hybrid solution that combines the strengths of both ACO and Gossips strategies, as the latter performs better in dynamic settings. Lastly, we aim at performing additional experiments as, for example, an exhaustive parameter testing for ACO.

ACKNOWLEDGMENT

This work is funded by the HORIZON Research and Innovation Action 101135576 INTEND “Intent-based data operation in the computing continuum.”

REFERENCES

- [1] C. Tsiganos *et al.*, “Towards resilient internet of things: Vision, challenges, and research roadmap,” in *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2019, pp. 1754–1764.

- [2] M. Villari *et al.*, “Osmotic computing: A new paradigm for edge/cloud integration,” *IEEE Cloud Computing*, vol. 3, no. 6, pp. 76–83, 2016.
- [3] A. Morichetta *et al.*, “A roadmap on learning and reasoning for distributed computing continuum ecosystems,” in *2021 IEEE International Conference on Edge Computing (EDGE)*. IEEE, 2021, pp. 25–31.
- [4] —, “Cohabitation of intelligence and systems: Towards self-reference in digital anatomies,” in *2024 IEEE International Conference on Service-Oriented System Engineering (SOSE)*. IEEE, 2024.
- [5] A. J. Ferrer *et al.*, “Towards a cognitive compute continuum: An architecture for ad-hoc self-managed swarms,” in *2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. IEEE, 2021, pp. 634–641.
- [6] H. Xiao *et al.*, “Edge intelligence: A computational task offloading scheme for dependent iot application,” *IEEE Transactions on Wireless Communications*, vol. 21, no. 9, pp. 7222–7237, 2022.
- [7] A. Maslow, “Self-actualization and beyond.” *Non-Journal*, 1965.
- [8] D. Bittman *et al.*, “Don’t let rpcs constrain your api,” in *Proceedings of the Twentieth ACM Workshop on Hot Topics in Networks*, 2021, pp. 192–198.
- [9] S. Wang *et al.*, “In reference to rpc: it’s time to add distributed memory,” in *Proceedings of the Workshop on Hot Topics in Operating Systems*, 2021, pp. 191–198.
- [10] M. Kountouris and N. Pappas, “Semantics-empowered communication for networked intelligent systems,” *IEEE Communications Magazine*, vol. 59, no. 6, pp. 96–102, 2021.
- [11] H. Seo *et al.*, “Semantics-native communication with contextual reasoning,” *arXiv preprint arXiv:2108.05681*, 2021.
- [12] E. Uysal *et al.*, “Semantic communications in networked systems: A data significance perspective,” *IEEE Network*, vol. 36, no. 4, pp. 233–240, 2022.
- [13] H. Xie and Z. Qin, “A lite distributed semantic communication system for internet of things,” *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 1, pp. 142–153, 2020.
- [14] S. Kumar *et al.*, “Comparative study on ant colony optimization (aco) and k-means clustering approaches for jobs scheduling and energy optimization model in internet of things (iot),” *IJIMAI*, vol. 6, no. 1, pp. 107–116, 2020.
- [15] L. Wang *et al.*, “Ant colony optimization for task allocation in multi-agent systems,” *China Communications*, vol. 10, no. 3, pp. 125–132, 2013.
- [16] A. Zannou *et al.*, “A task allocation in iot using ant colony optimization,” in *2019 International Conference on Intelligent Systems and Advanced Computing Sciences (ISACS)*. IEEE, 2019, pp. 1–6.
- [17] M. K. Hussein and M. H. Mousa, “Efficient task offloading for iot-based applications in fog computing using ant colony optimization,” *IEEE Access*, vol. 8, pp. 37 191–37 201, 2020.
- [18] A. Kishor and C. Chakarbarty, “Task offloading in fog computing for using smart ant colony optimization,” *Wireless personal communications*, vol. 127, no. 2, pp. 1683–1704, 2022.
- [19] K.-H. N. Bui and J. J. Jung, “Aco-based dynamic decision making for connected vehicles in iot system,” *IEEE Transactions on Industrial Informatics*, vol. 15, no. 10, pp. 5648–5655, 2019.
- [20] L. Tan *et al.*, “Energy-efficient joint task offloading and resource allocation in ofdma-based collaborative edge computing,” *IEEE Transactions on Wireless Communications*, vol. 21, no. 3, pp. 1960–1972, 2021.
- [21] Y. Wang *et al.*, “Bi-objective ant colony optimization for trajectory planning and task offloading in uav-assisted mec systems,” *IEEE Transactions on Mobile Computing*, 2024.
- [22] F. Xu *et al.*, “Research on computing offloading strategy based on genetic ant colony fusion algorithm,” *Simulation Modelling Practice and Theory*, vol. 118, p. 102523, 2022.
- [23] C. Li *et al.*, “Dqn-enabled content caching and quantum ant colony-based computation offloading in mec,” *Applied Soft Computing*, vol. 133, p. 109900, 2023.
- [24] S. Dong *et al.*, “Quantum particle swarm optimization for task offloading in mobile edge computing,” *IEEE Transactions on Industrial Informatics*, vol. 19, no. 8, pp. 9113–9122, 2022.
- [25] G. Di Caro and M. Dorigo, “Antnet: Distributed stigmergetic control for communications networks,” *Journal of Artificial Intelligence Research*, vol. 9, pp. 317–365, 1998.
- [26] E. Michlmayr, “Ant algorithms for search in unstructured peer-to-peer networks,” in *22nd International Conference on Data Engineering Workshops (ICDEW’06)*. IEEE, 2006, pp. x142–x142.
- [27] J. Liu *et al.*, “Reliability-enhanced task offloading in mobile edge computing environments,” *IEEE Internet of Things Journal*, vol. 9, no. 13, pp. 10 382–10 396, 2021.
- [28] Q. Chen *et al.*, “Two-stage evolutionary search for efficient task offloading in edge computing power networks,” *IEEE Internet of Things Journal*, 2024.
- [29] M. Dorigo and L. M. Gambardella, “Ant colony system: a cooperative learning approach to the traveling salesman problem,” *IEEE Transactions on evolutionary computation*, vol. 1, no. 1, pp. 53–66, 1997.
- [30] J. Fan *et al.*, “Deadline-aware task scheduling in a tiered iot infrastructure,” in *IEEE GLOBECOM 2017*. IEEE, 2017, pp. 1–7.
- [31] T. C. Müller, “MI-based power consumption prediction models for edge devices,” Master’s thesis, TU Wien, 2023. [Online]. Available: <https://repositum.tuwien.at/handle/20.500.12708/175670>