

# A Decentralized Approach for Resource Discovery using Metadata Replication in Edge Networks

Ilir Murturi and Schahram Dustdar  
Distributed Systems Group  
TU Wien, Vienna, Austria

**Abstract**—Recent advancements in distributed systems have enabled deploying low-latency edge applications (i.e., IoT applications) in proximity to the end-users, respectively, in edge networks. The stringent requirements combined with heterogeneous, resource-constrained and dynamic edge networks make the deployment process a challenging task. Besides that, the lack of resource discovery features make it particularly difficult to fully exploit available resources (i.e., computational, storage, and IoT resources) provided by low-powered edge devices. To that end, this paper proposes a decentralized resource discovery mechanism that enables discovering resources in an automatic manner in edge networks. Through replicating resource descriptions (i.e., metadata), edge devices exchange information about available resources within their scope in a peer-to-peer manner. To handle the resource discovery complexity, we propose a solution to build edge networks as a flat model and enable edge devices to be organized in clusters. Our approach supports the system in coping with the dynamicity and uncertainty of edge networks. We discuss the architecture, processes of the approach, and the experiments we conducted on a testbed to validate its feasibility on resource-constrained edge networks.

**Index Terms**—Edge Computing, Internet of Things, Decentralized, Resource Discovery

## I. INTRODUCTION

Researchers from academia and industry stakeholders suggest adding more computational resources (i.e., perceived as *edge devices*) in proximity to the end-users to overcome high-latency issues between the cloud and the Internet of Things (IoT) domain [1]. Edge devices are low-powered computer entities featuring different capabilities; resources available may differ in terms of computational capabilities and IoT resources attached to them. A wide range of available resources at the edge has introduced new opportunities such as deploying low-latency, privacy-awareness, and resilient *edge applications* (e.g., IoT applications). In this regard, many studies have been carried to exploit edge networks for various purposes (i.e., from processing sensory data streams to EdgeAI applications) [2]. Notably, we consider edge networks as resource-constrained, heterogeneous, and dynamic environments where multiple low-powered edge devices in proximity are connected. In this sense, we may have various edge networks (e.g., smart building, smart home, drone network, etc.) where end-users may deploy different edge applications.

In the past few years, computer scientists have been mostly focused on proposing multiple techniques for resource allocation problems to minimize latency and maximize resource utilization at the edge. Notably, today's applications are not

monolithic; they are divided into multiple independent deployable tasks. Each task may have various resource requirements that need to be fulfilled by available edge devices upon deployment. Tasks are characterized by requirements such as computational (i.e., processing, memory, storage), energy, or bandwidth. However, resource allocation approaches often overlook the dependence between tasks and IoT resources (e.g., sensors and actuators) [3]. Additionally, despite the numerous advantages introduced by edge networks, communication between edge devices and the network organization has been neglected by many research papers [4], [5]. According to the paper [6], the communication and network organization type of a platform affects the functionality of the final applications deployed at the edge infrastructure.

Very few research works consider IoT resources as an application requirement that needs to be fulfilled when deploying them at the edge [7], [8]. For example, computing a local weather forecast in a smart agriculture setting may require various IoT resources such as temperature and humidity readings from available sensors across a crop field [5]. In such a scenario, application tasks dependent on particular IoT resources must be deployed on edge devices providing those resources. Notably, edge devices are not equipped with the same capabilities, and such a stringent constraint reduces the number of eligible deployments at the edge. For example, the allocation technique [7] tries to overcome the problem by enabling sharing IoT resource information within neighbor nodes. Similarly, the proposed solution [8] acknowledges the problem; however, it faces latency issues, and it considers a limited number of edge devices in the network topology. Nevertheless, both approaches do not address issues related to the communication between edge devices, network organization, and resource discovery.

To overcome these shortcomings, we discuss two major issues. First, edge networks should be designated to handle the complexity of discovering resources in a decentralized and automatic manner. Thus, we design edge networks in a flat model where edge devices in certain proximity are connected in a peer-to-peer (P2P) way. A set of edge devices form a *cluster*; while multiple connected clusters form an edge network, respectively *an edge neighborhood*. Besides that, we introduce system coordinators with their corresponding functionalities to organize edge devices and support the resource discovery process in an edge neighborhood. Second, resource managements' fundamental objective is to discover resources available at the edge [9]. Edge devices provide heterogeneous

resources and are equipped with a rich set of IoT resources. We refer to heterogeneous resources as computational, sensing, context data, or other domain-specific resources. Naturally, performing a resource discovery algorithm for each resource on the entire network is possible. However, such a process is computationally intensive, and resources are discovered by querying the entire network based on a keyword [10]. Thus, we advocate that exchanging information about available resources between edge devices in an automatic manner enables: i) sharing resources across the whole system and ii) edge devices to perform complex queries locally.

In this paper, we extend the framework of [11] with a methodology to build edge networks as a flat model and enable edge devices to be organized in clusters. Our proposed framework enables discovering heterogeneous resources and make them available at the runtime. A resource is described by providing certain core information about the functionality and its properties. This type of description, known as the resource's metadata, is replicated among edge devices and stored in a decentralized manner. Furthermore, we treat privacy aspects based on each edge device's resource preferences, ensuring that not all resources are advertised across the whole system. Specifically, our concrete contributions are as follows:

- We develop a prototype enabling edge devices to be connected in a P2P manner and form an edge network. Our approach is built on top of the Kademlia Protocol [12] used as the communication protocol between edges devices. To enable scaling of our approach, we propose a solution to break down edge networks into clusters as well as introduce new coordinators to handle the complexity to discover resources automatically.
- We advocate decentralization as the system can operate without a static entity for discovering available resources. Essentially, coordinators are placed dynamically and run on the most suitable edge devices providing various services. Our approach supports the system in coping with the environment's dynamicity and uncertainty and continuously re-evaluates coordinators' placements.
- To validate the approach's feasibility, we show that the prototype's footprint is limited to hardware resources and network bandwidth. We evaluate our prototype on a testbed composed of low-powered ARM-based edge devices.

The rest of the paper is structured as follows. After a motivating example used throughout this paper in Section II, related work is presented in Section III. Section IV presents our approach to organize edge devices in edge neighborhoods and resource modeling. Section V describes in detail the proposed algorithm in charge of determining system coordinators and the framework for automatic resource discovery in edge neighborhoods. Section VI provides the experiment results to evaluate the proposed solution. Finally, Section VII concludes the paper and outlines future work directions.

## II. MOTIVATION SCENARIO

To motivate our discussion, we consider emergencies such as natural disasters (e.g., earthquakes, fires, floods) in a city

[13]. Various city areas can be affected by natural disasters such as earthquakes, which can destroy infrastructure, cause injury or death, and trap people under buildings. In such situations, time is valuable, and drones could be used to analyze the situation and assist rescue teams in locating and communicating with victims trapped under a collapsed building. In such a scenario, multiple connected drones are essentially edge devices forming an *edge neighborhood*. Drones flying over the city's affected areas (i.e., neighborhoods) assist rescue teams in locating people trapped under a collapsed structure. Each drone is equipped with various computation capabilities and integrated sensors (e.g., radar sensors, infrared cameras, electronic noses, etc.). We consider drones as multipurpose devices where the rescue teams may deploy various services depending on the emergency. Furthermore, we consider three-tier Edge-Cloud infrastructure (i.e., cloud, fog, and edge) [14]. Fog devices (i.e., server-graded hosts) placed in base stations provide computational and storage capabilities to edge neighborhoods. In addition to that, base stations may provide docker charge stations for charging drones. Cloud hosts can be used to store data for long terms.




Fig. 1: IoT safety application [13].

We assume that drones are connected via a wireless connection provided by the ground users (i.e., rescue team) or by drones [15] covering a particular city area (e.g., neighborhood). Based on the situation seen in Figure 1, we assume that the rescue team deploys (1) a public safety IoT service that detects a dangerous zone in the affected area (i.e., discovering cracks, smoke, hazardous gases, etc.). Such a service aims at helping rescue teams (2) find a safe path, and avoid danger zones. The service is dependent on various resources such as multiple infrared cameras, radar sensors, and an electronic nose that are integrated into various drones. Since each drone is a potential candidate to run the service, it is then evident that each edge devices should be able to automatically discover resources in a decentralized manner and make them available at runtime. In such use case scenarios, we cannot depend on the service availability [16] offered from physically static entities (3-4). Additionally, edge-based systems with centralized architecture cannot run properly due to the network dynamicity (i.e., drones may join (5) and leave often).

### III. RELATED WORK

We divide related work into two categories. First, we review P2P approaches and communication types used at the edge. Afterwards, we review related techniques on resource management as they apply to IoT.

#### A. Communication in Edge Networks

Many studies have been carried out, leading to different approaches aiming at exploiting the edge and provide various allocation techniques to fulfill application requirements. In this context, many platforms use different communication types aiming to achieve particular system goals. The current literature recognizes and briefly discusses communication types at the edge [17], [18]. According to the paper [6], three main types of communications in Edge Computing are identified: hierarchical, P2P, and hybrid.

Notably, P2P approaches have shown great potential to handle edge infrastructures in a scalable manner [19]. Therefore, a lot of research has been conducted in this context, resulting in many approaches that aim at organizing edge devices using different communication types [20]. Due to their fully distributed nature, P2P architectures are both scalable, reliable [21], and fault-tolerant [22]. Many edge-based platforms use P2P to organize edge devices within the edge network [23]. Similarly, Cabrera et al. [24] propose a P2P-based fog platform that enables storing data generated from IoT devices. In the proposed approach, data is stored among fog devices in a distributed manner. A fog device can restore corrupted data by asking other fog device in the network. In the mentioned works, devices are organized in a P2P manner and are assumed to have partial view [25] or full view of the network (i.e.,  $\mathcal{O}(1)$  protocols). In contrast to the mentioned works, our approach provides a decentralized mechanism to organize edge devices in clusters. Our solution determines the most suitable edge devices to place coordinators in edge networks and adapts to the network changes that may occur. Nevertheless, the proposed solution makes edge neighborhoods autonomous environments and less dependent on centralized nodes.

#### B. Resource Management

Performing a resource discovery algorithm for each resource on the entire edge neighborhood is computationally demanding. For instance, queries like discovering all cameras in a certain area are becoming highly desirable for IoT applications. In the general sense, such system behavior in decentralized common DHT protocols is hardly achievable. Even though some DHT-based approaches support discovering data through multi-attribute queries [26]; however, such methods remain unsuitable in IoT systems and edge networks due to the high content lookup latency. In addition to that, resource discovery is a critical challenge for IoT application performance.

Service-based discovery has been widely studied [27], [28]. Paganelli et al. [10] introduce a DHT-based IoT service discovery that supports multi-attribute and range queries. Furthermore, the proposed solution enables real-time monitoring of resource positions since it updates resource location periodically. Santos et al. [22] propose a resource discovery

service for resource provisioning in fog environments. The proposed solution is based on DHTs and enables exchanging provisioning information about the available resources (i.e., performance metrics, workloads, etc.). However, in contrast to our approach, the proposed solution does not address privacy aspects, does not consider discovering IoT resources, and no actual provisioning mechanisms are discussed.

Resource allocation and management have been widely studied both in cloud and fog computing [17]. A taxonomy of resource management at the edge is presented in [9]. Up to now, many factors have been considered including time (e.g., computation [29]), data size [30], cost (e.g., networking and deployment [31], execution [32]), deploying self-adaptive IoT systems [33], user-application context [32], etc., which have been found to play important roles in resource and service provisioning. Jain et al. [34] propose a solution where the IoT application is divided into multiple tasks annotated with location information. The application is decomposed into fragments and deployed to the corresponding individual compute nodes based on the annotation. In contrast to the mentioned research papers, the resource discovery aspect has been mostly ignored. Furthermore, none of the papers have addressed privacy aspects when considering resource discovery.

Our approach's second novelty lies in a decentralized mechanism for automatic resource discovery in edge networks. Discovering resources at once represents a feasible and optimal solution for edge neighborhoods. Through replicating metadata between edge devices, we enable end-users or edge applications to perform locally various complex queries. Moreover, our resource discovery mechanism considers resource privacy preferences ensuring that not all resources are advertised across the whole system.

### IV. EDGE NETWORKS AND RESOURCE MODELING

This section introduces our approach to organize edge devices in an edge neighborhood. We describe basic definitions and then discuss communication protocol between edge devices. Subsequently, we discuss architecture modeling and resource modeling in Section IV-D.

#### A. Definitions

We refer to an *edge neighborhood* as a resource-constrained edge network, which is comprised of edge devices placed close to each other (see Figure 2). Edge neighborhoods are formed in various geographical areas and within different contexts (i.e., smart homes, drones network, etc.). Notably, they may vary in their sizes and settings; thus, our proposed system is configurable by the system designer. In our conception, edge devices are low-powered, heterogeneous, and resource-constrained computational entities in the system. Within the system context, edge devices may provide multiple functionalities (e.g., act as a client device, server device, and bootstrap device). Furthermore, edge devices are grouped in clusters to promote scaling, reduce bandwidth consumption, and manage the difficulty of discovering resources in an automatic manner in edge neighborhoods. Subsequently, edge devices within the same cluster are considered *neighbor devices*, as well as

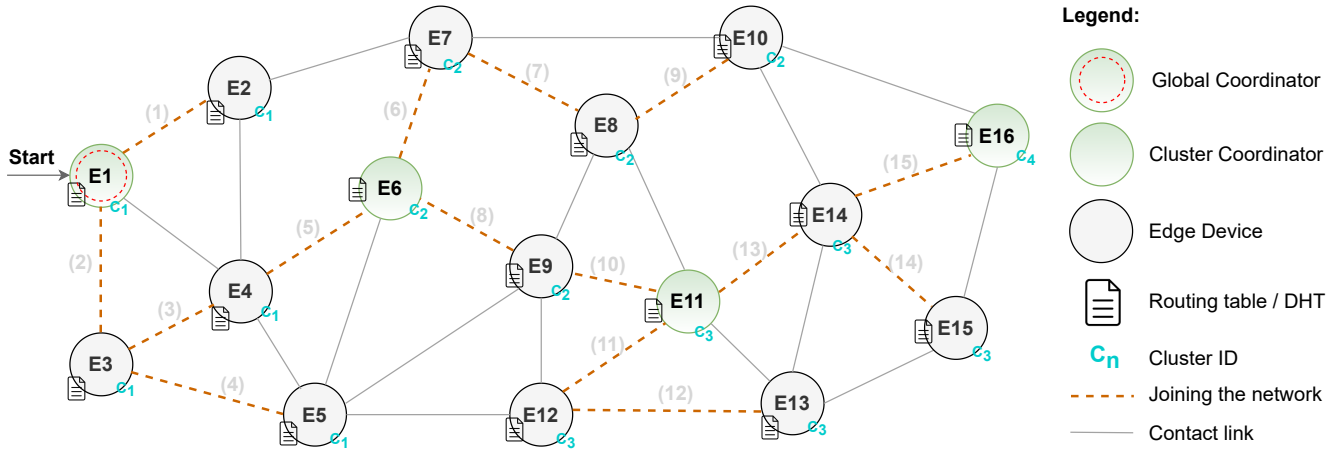


Fig. 2: An example of an edge neighborhood consisting of sixteen devices organized in four clusters.

clusters close to each other are considered as *neighbor clusters*. Each cluster consists of a finite number of edge devices, and each device belongs to one and only one cluster at a time.

In the system context, each edge device may serve specific roles such as *i) cluster coordinator* and *ii) global coordinator*. Such roles are dynamically assigned to the edge devices based on their resource capabilities (discussed in Section V-B). In Figure 2, each cluster has a coordinator device (i.e., with a green circle) and a single global coordinator device (i.e., with a red and green circles). We assume that cluster coordinators act as *superpeers* [35]. Each cluster coordinator keeps track of the other coordinators and devices within the same cluster (i.e., IP addresses, port). Similarly, edge devices in the same cluster store information for one another and are always aware of their cluster coordinator and global coordinator. A cluster coordinator may have various responsibilities for a subset of devices (i.e., monitoring, discovering resources, etc.). The global coordinator is responsible for monitoring coordinators, exchange resource descriptions between clusters, and orchestrate edge applications in Edge-Cloud infrastructure (i.e., controlling elasticity, migrating tasks, etc.). However, it remains the future work to provide a complete solution for the coordinators introduced in this paper. In this paper, we focus on three main aspects: *i) organizing edge devices in the edge neighborhood*, *ii) determining the most suitable devices to place the global and cluster coordinators*, and *iii) enabling automatic resource discovery on heterogeneous and dynamic edge neighborhoods*.

Each edge device may serve as the entrance door into the edge neighborhood. Essentially, edge devices provide core functionalities to assign newly added edge devices in the available clusters or create new clusters in the edge neighborhood. We introduce three functionalities to identify the maximum number of edge devices in a cluster. First, the system designer may define a system-wide parameter to bound the maximum size of clusters. Second, the system designer may configure cluster coordinators to generate random cluster size (i.e., not higher than the system-wide parameter). And third, the system designer may define a system-wide threshold value specifying the maximum CPU utilization for cluster coordinators. As a

result, depending on the functionality enabled, we may have edge neighborhoods with different cluster sizes.  $s$  in an edge neighborhood is not bounded (see Section III-C).

### B. Communication Protocol

In our proposed architecture, communications between edge devices is realized through implementing the Kademlia Protocol [12]. We have outlined [11] our main reasons to use Kademlia as the communication protocol between edge devices. Kademlia is a distributed hash table (DHT) for decentralized P2P computer networks. Essentially, DHT is a data storage that is kept consistent between all edge devices within the whole edge network. Essentially, when an edge device updates its local DHT, the changes are propagated to all other devices, allowing them to be queried and manipulated again. Likewise, information about current cluster coordinators and the global coordinator is stored in DHT.

Each edge device implements a distributed routing table and stores data in Kademlia buckets ordered by the local device's distance. The routing table size is bounded by  $\mathcal{O}(\log_2(l/k))$  where  $l$  is the number of edge devices in the edge neighborhood, and  $k$  is the bucket size. Once a bucket is full, it starts replacing unresponsive devices in favor of incoming devices. The routing table size in our proposed approach is configurable, and it depends on the expected edge neighborhood size. We consider edge devices as resource-constrained computers (i.e., in terms of computational and storage capabilities); thus, edge neighborhoods' size is not expected to be massive.

### C. Forming an Edge Neighborhood

Figure 2 shows an illustration of how our solution organizes edge devices in an edge neighborhood. Initially, the edge neighborhood is formed with an edge device E1. Since it is the only device in the edge neighborhood, it is automatically assigned to cluster C1 and determined as the cluster coordinator  $C1_{coord}$  and the global coordinator  $G_{coord}$ . At the same time, E1 serves as an entrance door into the edge neighborhood. We assume edge devices progressively join the edge neighborhood and each edge device also becomes another bootstrap device

$E_{n_b}$  (i.e.,  $n$  is a unique random ID). It is a common approach to keep a list of always-running edge devices to allow new edge devices to join an edge neighborhood. In our case,  $E_1$  is the first device contacted by  $E_2$  to join the edge neighborhood. The edge device  $E_2$  is assigned to cluster  $C_1$  by  $E_1$  in coordination with  $C_{1\text{coord}}$ . Furthermore,  $E_2$  is supplied with DHTs and the complete routing table from  $E_1$ . The process of adding new edge devices in an edge neighborhood is presented in Algorithm 1.

---

**Algorithm 1:** Process of adding a new edge device

---

**Input :**  $E_{\text{new}}$   
**Output:** Adding  $E_{\text{new}}$  to the edge neighborhood

```

1  $C_M = \text{this.MaxClusterSize}()$ 
2  $C_S = \text{this.CurrentClusterSize}()$ 
3  $C_N = \text{this.OtherClusters}()$ 
4 boolean flag = false
5 if  $C_M < C_S$  then
6    $E_{n_b}.\text{addDevice}(E_{\text{new}}, \text{this.cluster}())$ 
7 else
8   if  $C_N \neq \text{null}$  then
9      $C_{clo} = \text{this.findMostClosest}(C_N)$ 
10    foreach  $c \in C_{clo}$  do
11      if  $c.\text{size}() <$ 
12         $C_N.\text{cluster}(c).\text{maxSize}()$  then
13         $E_{n_b}.\text{addDevice}(E_{\text{new}}, c.\text{cluster}())$ 
14        flag = true
15        break
16    end
17  end
18 end
19 if flag = false then
20    $E_{n_b}.\text{addDevice}(E_{\text{new}}, \text{new Cluster}())$ 
21 end
22  $\text{this.updateRoutingTable}()$ 
23  $\text{this.updateDHT}()$ 

```

---

The process continues with adding a new edge device  $E_{\text{new}}$  by contacting  $E_{n_b}$ . Once  $E_{\text{new}}$  request to join the neighborhood, bootstrap device  $E_{n_b}$  initially provides a unique random ID (i.e., 160-bit). Afterward,  $E_{n_b}$  retrieves information from its cluster coordinator  $C_{n\text{coord}}$  where to assign  $E_{\text{new}}$ . A cluster coordinator  $C_{n\text{coord}}$  (i.e., referred as *this*) maintain information regarding: i) own cluster maximum size  $C_M$  (line 1), the current cluster size  $C_S$  (line 2), and other available clusters  $C_N$  (line 3).  $C_N$  provides clusters with available places where  $E_{\text{new}}$  can be assigned.

The  $\text{MaxClusterSize}()$  function is configurable and enables the system designer to define the maximum number of edge devices per cluster. Such value can be determined by i) the system-wide parameter (e.g., five edge devices per cluster or random value), implying that when the cluster exceeds the maximum allowed devices, more clusters with their corresponding coordinators should be designated to handle the resource discovering complexity, and ii) the CPU utilization threshold (e.g., CPU utilization set to 35%). When the system-

wide parameter is set and random cluster size is disabled, edge devices are grouped into clusters of the same size (as illustrated in Figure 2). Edge devices are grouped into clusters of different sizes when the CPU utilization threshold is set. More specifically, when  $C_M$  and  $C_S$  are equal,  $E_{\text{new}}$  is assigned to one of the existing clusters, or a new cluster is created. Nevertheless, both options can be used at the same time. However, the option that is violated decides whether or not the cluster can scale further. Furthermore, if the condition in (line 5) is not violated,  $E_{\text{new}}$  is assigned to the current cluster of the  $E_{n_b}$  (line 6).

Neighbor clusters are found through using a system call (i.e., *traceroute command*), which estimates proximity with other cluster coordinators (i.e., using hop count and latency). The function  $\text{findMostClosest}(C_N)$  uses *traceroute* command and returns the most suitable clusters that  $E_{\text{new}}$  can be assigned (line 9). This is especially useful in edge neighborhoods running on different networks. Finally,  $C_{n\text{coord}}$  provides information to  $E_{n_b}$  to assign the  $E_{\text{new}}$  to the particular cluster if and only if condition in (line 11) is not violated (lines 11-15). Otherwise, when there are no clusters with free places, then  $E_{\text{new}}$  is assigned to a newly created cluster by  $C_{n\text{coord}}$  (line 20). Note that each cluster in the edge neighborhood has a unique ID generated by  $E_{n_b}$  when the cluster is created. Notably, the bootstrap device's task is to cooperate with the cluster coordinator to assign a cluster ID to the newly joined devices. Finally, once  $E_{\text{new}}$  joins the edge neighborhood,  $C_{n\text{coord}}$  updates the routing table and other DHTs (lines 22-23).

#### D. Resource Modelling

We assume that an edge device may contain multiple resources (i.e., computational, sensing, or context data) represented as microservices [36]. When invoked remotely, such microservices yield resources; however, resource information needs to be shared among edge devices in the edge neighborhood beforehand. An essential step towards discovering these resources in the edge network is resource modeling at design time. To ensure automatic resource discovery in an edge setting, two types of resources must be modeled: i) *IoT resources* (i.e., sensors, actuators, etc.) and ii) *edge devices*.

To accomplish our goal in a pervasive environment, we have outlined resource representation structure in [11]. The resource structure provides seven main properties such as: i) resource identification, ii) resource connectivity, iii) resource capability, iv) resource accessibility, v) resource output, vi) resource location, and vii) resource administrative domain. Our resource structure is similar to the ontology-based structure proposed by Barnaghi et al. [37]. Unlike the ontology-based approach, we format resource descriptions in a JavaScript Object Notation (JSON). We advocate that exchanging metadata over JSON is a lightweight process, machine-readable, and provides rich information about resources. Besides that, size of metadata in JSON is very small. Thus, the replication process across many edge devices organized in clusters is feasible and does not degrade the overall network performance.

## V. A DECENTRALIZED EDGE-TO-EDGE RESOURCE DISCOVERY

In this section, we first discuss the overall design goals for the proposed approach in edge networks. Then, we discuss the process of determining the global coordinator and cluster coordinators. Next we explain in detail the framework for sharing edge device resources based on their privacy preferences. Finally, we discuss edge device failures, dynamicity and uncertainty of edge neighborhoods.

### A. Design Goals

Edge Computing introduces effective ways to overcome many of the limitations faced by the cloud [3]. Nevertheless, the paradigm alone also presents some limitations (i.e., computation capability and latency). To address such challenges, we identify three design goals that need to be established by any edge-based system:

- **Latency-aware.** The Edge Computing paradigm aims at providing low-latency services for the endpoint devices and the end-users. As a result, determining system coordinators at the edge neighborhood quickly is essential to fulfilling such a stringent requirement.
- **Dynamic Network.** An edge neighborhood changes with time (i.e., new edge devices can be added or excluded). Hence, the edge-based systems should be able to utilize newly added resources at the edge flexibly.
- **Adaptability.** Edge-based systems should be able to adapt to unexpected changes that might occur in the edge neighborhood. Thus, the coordinators should be dynamically placed among available edge devices and continuously re-evaluate placement decisions.

### B. The Process of Determining Coordinators

The process to determine coordinators in a distributed system should be carried out in system background, encapsulated from the end user's perspective, but indispensable for the correct and efficient execution of distributed tasks. System coordinators' role is versatile and can range between orchestrating applications, monitoring resources, or distributing data between devices. We define two leading roles, such as *i) cluster coordinator* and *ii) global coordinator*. Such roles are dynamically assigned to edge devices based on their resource capabilities. The process to determine new coordinators is triggered by an event when the global or a cluster coordinator experiences high utilization in specific hardware resources (i.e., CPU, memory, or storage) and requests to transfer leadership to other devices. The process to determine a new cluster coordinator occurs only between edge devices within the same cluster. The process to determine a new global coordinator occurs between cluster coordinators. The latter essentially consists of two phases: first, cluster coordinators are determined; second, new cluster coordinators determine the global coordinator. In Algorithm 2, we present the process to determine system coordinators.

The proposed algorithm runs on each edge device separately. The algorithm takes three inputs: i) an edge device

### Algorithm 2: Process of determining coordinators

---

**Input :**  $\phi_i, \theta, \sigma_i$   
**Output:**  $C_{coord}, G_{coord}$

```

1  $t = 0$ 
2  $\nu_{this} = \text{GetDeviceMetrics}(\phi_i)$ 
3  $\text{round} = \text{Random}()$ 
4  $\text{Initial\_Message}(\text{round}, \sigma_i)$ 
5  $\text{Parameter\_Message}(\nu_{this}, \text{round}, \sigma_i, \gamma)$ 
6  $I \leftarrow \text{Receive\_Parameter\_Messages}()$ 
7  $\text{Solution\_Found} \leftarrow \text{False}$ 
8 while  $t < \theta$  or  $!\text{Solution\_Found}$  do
9    $\nu = I.\text{getBest}()$ 
10  if  $\nu < \nu_{this}$  and  $\nu.\text{round} = \text{round}$  then
11     $\text{Solution\_Found} \leftarrow \text{True}$ 
12     $\text{Set}(E_{this}, \text{EdgeMode})$ 
13     $\text{Set}(En, \text{Coordinator}, \sigma_i)$ 
14  end
15 if  $!\text{Solution\_Found}$  then
16    $\text{Set}(E_{this}, \text{Coordinator}, \sigma_i)$ 
17 end
18  $E_{this}.\text{updateDHT}()$ 

```

---

hardware metrics denoted with  $\phi_i$ , ii) the deadline to find a solution denoted with  $\theta$ , and iii) the process type denoted with  $\sigma_i$ . The process to determine coordinators is designed by considering hardware metrics and bandwidth of edge devices. We consider hardware metrics both statically (e.g., CPU cores, storage capacity, etc.) and dynamically (e.g., current CPU load, current storage, etc.). Such hardware information can be monitored using *Hyperic Sigar* [38] while *Assolo* [39] enables collecting bandwidth probes. Notably, the algorithm gets only the current device metrics  $\nu_{this}$  specified at design time (line 2). In our case, we consider metrics  $\theta_i$  related to the CPU utilization degree. However, such a parameter is configurable based on system requirements. The deadline  $\theta$  is given at design time (e.g.,  $\theta = 50\text{ms}$ ). The third input  $\sigma_i$  specify the process type: i) determining the global coordinator ( $\sigma = \text{global}$ ) and ii) determining a cluster coordinator ( $\sigma = \text{cluster}$ ).

A unique random signature (i.e., SHA-1) called *round* is used to make each process unique (line 3). In lines (4-5), we define two types of messages exchanged between devices: *i) initial message* and *ii) parameter message*. First, the initiating coordinator (i.e., process initiator device) probes which edge devices are up and running in the edge neighborhood. Then, it sends an initial message only to the edge devices that responded on time. The initial message essentially contains a list of all participating devices (i.e., only those edge devices that replied) and a unique *round* value assigned to a process round. Second, once receiving an initial message and the list of participants, the process initiator device sends a parameter message containing its local metrics to all participating devices. The parameter message contains local hardware metrics, utilization values in percentage, a process round value, and timestamp when the process to determine the new coordinator is started  $\gamma$ . The timestamp  $\gamma$  is used to ensure that the metrics



are not older than the process initiation time. Each edge device responds to the initial message and sends its parameter message to other edge devices, including the initiating coordinator (line 6). Notice that the same proposed algorithm is used to determine the global and cluster coordinators.

After exchanging performance metrics, edge devices develop the same result independently. Each device compares received parameter metrics and determines which device is most suitable to serve as a cluster/global coordinator (lines 8-14). To find the most suitable edge device, we compare the number of CPU cores, clock speed, and current utilization (line 9).  $E_{this}$  changes the status to the EdgeMode (i.e., edge device only shares resources) only when an edge device  $E_n$  with better performance metrics is found. Line 16 is executed only when a solution is not found within the time  $\theta$ , and the coordinator role remains in the current edge device  $E_{this}$  since no better edge device is found. Furthermore, each edge device maintains a complete overview of all processes and saves the information in a DHT. The result is propagated to all edge devices within the edge neighborhood by using the DHT. The latter step is executed by all edge devices, which on the one side creates some redundancy but also improves the stability of propagating results (line 18).

### C. System Architecture and Resource Discovery

We propose a framework for enabling automatic discovery of heterogeneous resources in edge networks [11]. Figure 3 illustrates three main components of the framework: Edge-to-Edge Communication, the Metadata Container, and the Local Search Engine facility. The Edge-to-Edge Communication component implements communication between edge devices (i.e., Kademlia Protocol), organizes edge devices in clusters, determines coordinators, and exchanges resource descriptions in an edge neighborhood (1). The Metadata Container is responsible for analyzing resource descriptions based on their privacy preference defined by the system designer at design time. Besides that, the component is responsible for sharing resource descriptions system-wide (2), processing received resource descriptions from other edge devices (3), and store resource descriptions locally (4). The Local Search Engine component (4) is based on CouchDB document-oriented NoSQL database [40]. This component subsequently stores data locally and, through APIs, enables users to query stored content (6).

Referring to Figure 3, our resource discovery mechanism allows edge devices within the same cluster to exchange metadata through their cluster coordinator. Once coordinators are determined, edge devices are ordered to share their public metadata document with their cluster coordinator (1). Other edge devices contact their cluster coordinator to retrieve all public metadata documents contained within their cluster. This may happen once a new edge device is connected to the edge neighborhood or when an edge device wants to refresh its current storage. The frequency to refresh metadata documents is also configurable at design time. Besides that, the global coordinator regularly exchanges metadata documents with the cluster coordinators (5).

In Figure 4, we present the process to analyze resource descriptions and the process to share them in edge neigh-




Fig. 3: Resource discovery through metadata replication: High-level architecture.

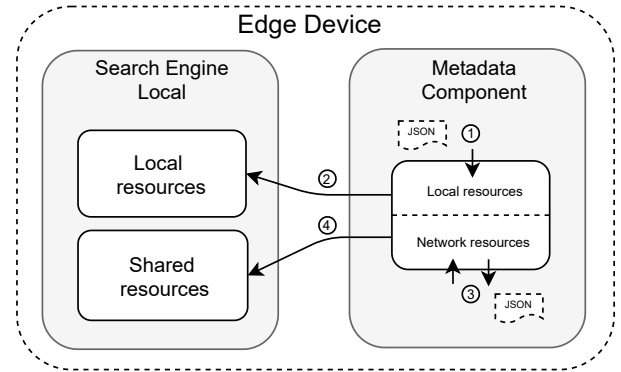


Fig. 4: Processing metadata on the edge.

borhood. As mentioned before, we consider edge devices as resource-constrained devices with a set of built-in sensors and actuators. Each resource is described by providing certain core information about the functionality and properties (see Section IV-D). We assume that resource descriptions (i.e., metadata) are provided by edge device manufacturers in JSON format. The Metadata Container is responsible for analyzing metadata in the edge device. The component stores resource descriptions based on privacy preference i) public resources (i.e., shared resources (4)) and ii) private resources (i.e., local resources (2)). Public resource descriptions are merged into a single metadata document and named with the edge device ID. The public metadata document is shared with the corresponding cluster coordinator system-wide (3). The local search engine component separates resource descriptions into private ones and those that are shared system-wide.

### D. Edge Device Failures

Consider a situation when a cluster coordinator triggers a new process to determine the new coordinator in the cluster. An edge device under a high workload may fail to participate in determining cluster coordinators or the global coordinator. Even though not participating in a determination process, edge devices continuously update their local DHTs with the closest devices. Besides that, those who have failed and re-joined the edge neighborhood may also contain some obsolete information (i.e., DHT is not up-to-date, coordinator information is outdated). However, edge devices must update their

local DHTs with the closest devices after joining the edge neighborhood.

Another situation may arise when a specific cluster coordinator fails. The global coordinator is responsible for detecting such failure and trigger a new process to determine the coordinator. As discussed previously, edge devices within the cluster are responsible for determining their new coordinator. When the global coordinator fails, one of the cluster coordinators triggers a new global coordinator determination process. Notably, each cluster coordinator verifies whether the global coordinator has failed and responds to the determination process. Furthermore, the resource discovery process is not repeated when edge devices frequently leave and join the edge neighborhood (i.e., due to the connectivity issues). Moreover, when an edge device went offline, the metadata documents remain stored in the other devices for some time. This is especially useful in unstable edge neighborhoods (e.g., wireless), prone to momentary loss of connection.

## VI. EVALUATION

In this section, we first discuss our evaluation setup environment, prototype details, and limitations. Next, we experimentally evaluate the approach's effectiveness by running multiple experiments and checking the proposed solution's behavior in different situations. We assess our proposed solution in terms of hardware and bandwidth consumption during runtime. Then, through a use case, we show how the proposed discovery mechanism increases the number of eligible deployments when deploying edge applications in an edge neighborhood. Finally, we conclude with a discussion in Section VI-D.

### A. Setup, Prototype Details, and Limitations

To assess the proposed approach's feasibility, we developed a prototype that implements the Kademlia Protocol and core functionalities to enable the automatic discovery of heterogeneous resources in edge networks. The prototype is written in Java, and it is tested on a testbed composed of (a) edge devices (i.e., Raspberry Pi 3 Model B V1.2) with 4×ARM Cortex-A53 CPU at 1.2 GHz, 1 GB of RAM, and 16 GB disk storage, and (b) virtual edge device instances. To evaluate our prototype, we exploited the testbed (i.e., edge neighborhood) composed of 60 edge devices placed close to each other. Edge devices contain multiple resource descriptions generated randomly at design time. Furthermore, edge devices in the testbed are connected through a wireless connection with a nominal speed of 10 Mbps and 5 Mbps in download and upload. We assume that every edge device trusts all other devices, and they all belong to the same local administrative domain.

In distributed systems, discovering devices is a significant challenge. In our current implementation, edge devices have an open designation port that listens for possible future connections. To join the edge neighborhood, edge devices require to know at least one device (i.e., IP, port) currently up and running. We acknowledge that the current implementation represents a limitation, and further investigations are required to develop an advanced approach to discover running edge devices. We acknowledge that IoT resources (e.g., sensors)

can also be connected on edge devices using various end-to-end communication protocols (e.g., Zigbee, etc.). However, in this paper, we treat communication and operational aspects of IoT resources as orthogonal to our approach; we assume edge devices are equipped with a set of IoT resources. Furthermore, we acknowledge that using the Java Virtual Machine (JVM) environment is resource-expensive. However, within this paper, we aim to show the approach's feasibility in resource-constrained edge neighborhoods.

### B. Experiments and Results

We evaluate our prototype on a testbed, whose size progressively increases to 60 edge devices as presented in Table I. The cluster coordinators can determine their cluster sizes based on i) the CPU utilization threshold (i.e., configured to 35%), ii) the system-wide parameter (i.e., configured to 30 devices per cluster), and iii) the random value (i.e., not bigger than the system-wide parameter). Furthermore, the routing table size is set to  $k = 20$ . We monitor edge devices through the *nmon* tool [41] and retrieve information regarding the hardware utilization and the data received and sent between edge devices.

TABLE I: Edge neighborhood

Neighborhood	Edge Devices (per cluster)	Total
C1	15	15
C2	20	35
C3	15	50
C4	10	60

The goal of the first experiment is to assess the prototype's footprint on hardware resources and bandwidth usage. Notably, we focus on resource consumption to determine the system coordinators and discover resources in the edge neighborhood. For each cluster created, we monitor the global coordinator for up to 60 seconds in the edge neighborhood (i.e., running on RPi3). We also monitor bootstrap devices (i.e., RPi3s) when a new request arrives to join the edge neighborhood. Notably, we observe that the time required to join the edge neighborhood and assign it to the existing cluster is between 0.07 – 0.2 seconds in all test cases. However, when a new cluster is required to be created, the average latency is slightly higher between 0.5 – 1.02 seconds, but reasonable for the resource-constrained edge neighborhood. Specifically, the overall CPU and memory utilization remains almost similar. Concretely, the average CPU consumption stayed around 2.5%, while the overall memory consumption stayed around 5 MB. In all test cases, edge devices have been successfully assigned to the corresponding clusters, or new clusters are created when required. Table II presents a detailed overview of the resource consumption and latency for edge devices to join the neighborhood.

Figure 5 plots the overall CPU utilization during the process to determine the global coordinator (i.e., the global and cluster coordinators are determined) and synchronization processes running in the background. Essentially, we simulate the situation when the global coordinator is overloaded, and the function to determine the new global coordinator is automatically



TABLE II: Average latency and resource consumption to join the edge neighborhood

Clusters	Latency (avg.)	CPU (avg.)	RAM (avg.)
C1	71.25 ms	2%	2 MB
C2	82 ms	2.5%	3 MB
C3	137.36 ms	2.5%	2 MB
C4	193.6 ms	3%	5 MB

triggered. We repeat the process more than 10 times for each test case (i.e., clusters). As shown in Figure 5, the overall CPU utilization is slightly increased when adding more edge devices/clusters in the edge neighborhood. Specifically, the average CPU consumption during the process to determine the system coordinators is around 10%, while the overall memory consumption is around 5 MB. The most important aspect is the accuracy of assigning edge devices to a particular cluster or forming new clusters when needed (i.e., when the cluster size is exceeded). In all test cases, creating new clusters was successful, including newly joined edge devices added to their adequate clusters. Since the overhead imposed by our approach is small, the experiments show that the proposed approach is feasible to operate on low-powered and battery-powered edge devices. Notably, the proposed approach has shown a very contained impact on hardware resources and bandwidth usage. To obtain consistent results, for each experiment, we calculated an 83% confidence interval of means.




Fig. 5: Analysis of the CPU utilization during the process to determine the system coordinators.

Figure 6 and Figure 7 plot the overall maximum and minimum data transfer/received by the global coordinator during the process to determine new system coordinators and the synchronization process. We consider analyzing the bandwidth due to the relationship between network traffic and the load on the memory system of an edge device [15]. Furthermore, in resource-constrained edge neighborhoods, the energy supply and bandwidth are among the main resource constraints of edge devices [42]. Therefore, it is necessary to know the total data size transferred/received (i.e., metadata sharing, processes to determine coordinators, synchronizing processes, etc.) between edge devices during the runtime.




Fig. 6: Analysis of the data received in kilobytes per second by the global coordinator.




Fig. 7: Analysis of the data transfer in kilobytes per second by the global coordinator.

Notably, the maximum and minimum values vary depending on the number of edge devices in the neighborhood. As shown in Figure 7, the data transfer is slightly increased by adding more devices in the edge neighborhood. The slight increase occurs due to the increased number of edge devices that impose the formation of new coordinators. Thus, newly formed coordinators synchronize their routing tables and their DHTs with the global coordinator.

The goal of the second experiment is to assess the time complexity to determine coordinators and place them on the edge neighborhood's most suited edge devices. Concretely, we show how the proposed mechanism discussed in Section IV and Section V performs to determine coordinators. The experiment results illustrated in Figure 8 show that the proposed approach in all test cases finds an edge device with the most suitable computation resources to assign the global coordinator. Besides that, the proposed approach successfully determines cluster coordinators for each cluster.

In Figure 9, we analyze the percentage of responsive edge devices during the process to determine the cluster coordinator on an edge neighborhood with a single cluster and different




Fig. 8: Time required to determine coordinators.




Fig. 9: Percentage of responsive edge devices during the process to determine coordinator on clusters with different sizes.

sizes (e.g., CS=15, etc.). We repeat the process more than five times for each test case. Notably, we observe that the time required to join the edge neighborhood and assign it to the existing cluster remains similar to the previous results (i.e., between 0.07 – 0.2 seconds in all test cases). The number of edge devices that respond to this process is critical since it enables us to find the most suitable edge device to place the cluster coordinator. In clusters with more than 40 edge devices, the percentage of responsiveness is decreased to 92%. In contrast, clusters up to 30 edge devices show a higher rate of responsiveness. To that end, we configure the system-wide parameter to bound the maximum number of edge devices per cluster (i.e., 30 edge devices).

### C. Use case: Deploying IoT safety application

The third experiment aims at demonstrating the proposed mechanism's function in the application deployment process. For the demonstration purpose, we adopt and extended parts of the FogTorchII simulator [43] to generate deployment plans in the edge infrastructure. FogTorchII was originally proposed

to support IoT designers in making decisions about where to deploy application components at the edge. The approach allows to specify IoT resources as an application requirement that must be met before applications can be deployed.

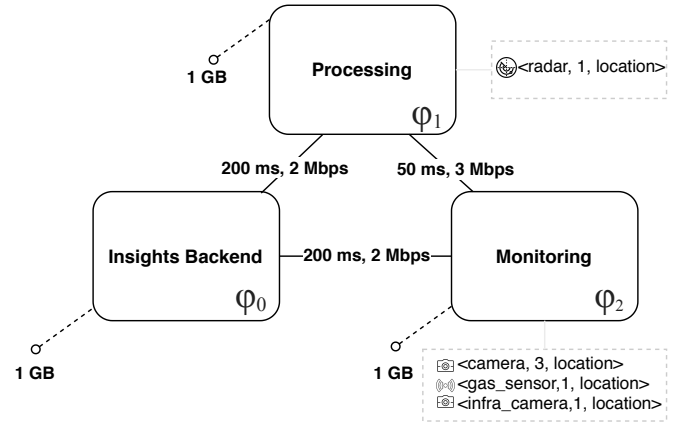


Fig. 10: IoT safety application.

We assume edge devices containing a set of built-in IoT resources (i.e., cameras, radars, gas sensors, etc.) that can be accessed remotely through APIs. Recalling our motivation scenario, consider the IoT safety application that provides a service to rescue teams (as discussed in Section II). As illustrated in Figure 10, the IoT safety application comprises three components: i) the insights backend component ( $\varphi_0$ ), ii) the monitoring component ( $\varphi_2$ ), and iii) the processing component ( $\varphi_1$ ). The insights backend component ( $\varphi_0$ ) enables users to interact with the service. The monitoring component ( $\varphi_2$ ) is responsible for monitoring resources specified at design time. The processing component ( $\varphi_1$ ) analyzes collected data into meaningful results and provides it to the end-users through the insight component. We assume deploying IoT safety application in the edge neighborhood as presented in Figure 2. Notice that only a single application component can be executed in parallel on edge devices. The other application requirements depicted in Figure 10 are assumed to be met by all edge devices.

In our given edge neighborhood, the edge device  $E_2$  provides the radar resource with privacy preference set to private. This means that the  $\varphi_1$  can be deployed only in  $E_2$ . The rest of IoT resources (e.g., cameras, gas sensors, etc.) are remotely accessible, and their privacy preference is set to public. Through the metadata replication process, edge devices exchange public resources in the edge neighborhood. To this end, each edge device is considered a potential candidate to execute one of the two other components. Meanwhile, each edge device shares private resources only with the global coordinator, responsible for generating deployment plans for the IoT safety application. Figure 11 shows the total number of eligible deployments plans generated through FogTorchII combined with our resource discovery mechanism. It is evident that  $\varphi_0$  and  $\varphi_2$  components dependent on particular resources can be executed on all other edge devices. More precisely, each edge device knows how to access resources on other edge devices due to metadata replication.




Fig. 11: Generating deployments plans for the IoT safety application in an edge neighborhood (see Figure 2).

Unlike our approach, FogTorchII does not support resource privacy preferences, does not provide a mechanism to discover resources (i.e., supports resource sharing with neighbor devices), and does not support the dynamic changes in their environments. Besides that, evaluating the deployment mechanism is out of this paper's scope, as we use it to demonstrate that the discovery mechanism allows us to completely exploit available resources in edge neighborhoods.

#### D. Discussion

We have demonstrated that by using our resource discovery mechanism, discovering heterogeneous resources in an edge setting increases the number of eligible deployments for applications dependent on IoT resources. We showed that organizing edge devices in clusters and discovering resources through the edge replication process is performant and feasible on a testbed with low-powered ARM-based devices. Our results showed that the overhead introduced by the proposed decentralized resource discovery mechanism is very low for realistic edge neighborhood sizes.

Our approach within the given testbed can operate with larger cluster sizes. However, edge devices may not react in time to participate in processes to determine coordinators in large cluster sizes. In such cases, the Kademlia Protocol requires much more processing capabilities to process faster incoming requests (i.e., updating routing tables, updating DHTs, resource metadata, etc.). Nonetheless, we plan to improve and optimize the current mechanism to support forming larger clusters in edge neighborhoods. We plan to investigate the maximum edge neighborhood size and the routing table size acceptable for low-powered edge devices.

Our approach does not provide a mechanism to detect which edge devices fail to send their parameter metrics and ask them to resend their messages. In addition to that, since multiple edge devices in an edge network come to the same result independently, it is necessary to introduce an additional mechanism to verify the result. In our approach, we overcome such an issue by distributing results using the DHT. This creates some redundancy - but also improves the stability of propagating results. However, using consensus algorithms where edge devices pick random participants to verify their final result is highly desirable. Thus, some assumptions underlying our methodology must be further explored.

## VII. CONCLUSION AND FUTURE WORK

Edge networks provide a seamless opportunity for deploying various edge applications providing multiple services to the end-users and the surrounding IoT devices. However, to deploy edge applications dependent on IoT resources, we require novel lightweight and decentralized mechanisms to automatically discover heterogeneous resources at the edge. To that end, we introduced a decentralized mechanism that enables edge devices to connect in a P2P manner, organize edge devices in clusters, and support automatic discovery of heterogeneous resources in edge networks. The proposed approach support resource discovery based on resource privacy preferences. Furthermore, we evaluated our approach in a testbed composed of a set of low-powered edge devices. Throughout the experiments, we showed the feasibility of the proposed approach to run on low-powered edge devices.

We believe that the proposed approach paves the way for utilizing available resources, leading to accomplish the promised high-quality and low-latency services deployed in edge networks (i.e., edge neighborhoods). As future work, we aim at providing a complete technical solution for the global and cluster coordinator; this includes both technical and architectural aspects. In addition to that, we plan to investigate techniques that will enable edge devices to discover nearby devices that allow them to join an edge neighborhood in an automatic manner. Finally, we plan to investigate several assumptions made in this paper.


## ACKNOWLEDGMENT

Research supported in part by the Research Cluster "Smart Communities and Technologies (Smart CT)" at TU Wien and it has received funding from the EU's Horizon 2020 Research and Innovation Programme under grant agreement No. 871525. EU web site for Fog Protect: <https://fogprotect.eu/>

## REFERENCES

- [1] Weisong Shi and Schahram Dustdar. The promise of edge computing. *Computer*, 49(5):78–81, 2016.
- [2] Schahram Dustdar and Ilir Murturi. Towards distributed edge-based systems. In *2020 IEEE Second International Conference on Cognitive Machine Intelligence (CogMI)*, pages 1–9. IEEE, 2020.
- [3] Cosmin Avasalcai, Ilir Murturi, and Schahram Dustdar. Edge and fog: A survey, use cases, and future challenges. *Fog Computing: Theory and Practice*, 2020.
- [4] Cosmin Avasalcai and Schahram Dustdar. Latency-aware decentralized resource management for iot applications. In *Proceedings of the 8th International Conference on the Internet of Things*, page 30. ACM, 2018.
- [5] Christos Tsigkanos, Ilir Murturi, and Schahram Dustdar. Dependable resource coordination on the edge at runtime. *Proceedings of the IEEE*, 2019.
- [6] Vasileios Karagiannis. Compute node communication in the fog: Survey and research challenges. In *Proceedings of the Workshop on Fog Computing and the IoT*, pages 36–40. ACM, 2019.
- [7] Antonio Brogi and Stefano Forti. Qos-aware deployment of iot applications through the fog. *IEEE Internet of Things Journal*, 4(5):1185–1192, 2017.
- [8] Cosmin Avasalcai, Christos Tsigkanos, and Schahram Dustdar. Decentralized resource auctioning for latency-sensitive edge computing. In *2019 IEEE International Conference on Edge Computing (EDGE)*, pages 72–76. IEEE, 2019.
- [9] Klervie Toczé and Simin Nadjm-Tehrani. A taxonomy for management and optimization of multiple resources in edge computing. *Wireless Communications and Mobile Computing*, 2018, 2018.

- [10] Federica Paganelli and David Parlanti. A dht-based discovery service for the internet of things. *Journal of Computer Networks and Communications*, 2012, 2012.
- [11] Ilir Murturi, Cosmin Avasalcai, Christos Tsigkanos, and Schahram Dustdar. Edge-to-edge resource discovery using metadata replication. In *2019 IEEE 3rd International Conference on Fog and Edge Computing (ICFEC)*, pages 1–6. IEEE, 2019.
- [12] Petar Maymounkov and David Mazieres. Kademia: A peer-to-peer information system based on the xor metric. In *International Workshop on Peer-to-Peer Systems*, pages 53–65. Springer, 2002.
- [13] Schahram Dustdar and Ilir Murturi. *Towards IoT Processes on the Edge*, pages 167–178. Springer International Publishing, Cham, 2021.
- [14] Schahram Dustdar and Ilir Murturi. Towards distributed edge-based systems. In *IEEE 2nd International Conference on Cognitive Machine Intelligence (CogMI 2020)*, pages 1–9. IEEE, 2020.
- [15] Mohammad Mozaffari, Walid Saad, Mehdi Bennis, and M  rouane Debbah. Communications and control for wireless drone-based antenna array. *IEEE Transactions on Communications*, 67(1):820–834, 2018.
- [16] Yoshitaka Shibata, Noriki Uchida, and Norio Shiratori. Analysis of and proposal for a disaster information network from experience of the great east japan earthquake. *IEEE Communications Magazine*, 52(3):44–50, 2014.
- [17] Redowan Mahmud, Ramamohanarao Kotagiri, and Rajkumar Buyya. Fog computing: A taxonomy, survey and future directions. In *Internet of everything*, pages 103–130. Springer, 2018.
- [18] Flavio Bonomi, Rodolfo Milito, Jiang Zhu, and Sateesh Addepalli. Fog computing and its role in the internet of things. In *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*, pages 13–16. ACM, 2012.
- [19] Adrien Lebre, Jonathan Pastor, Anthony Simonet, and Fr  d  ric Desprez. Revising openstack to operate fog/edge computing infrastructures. In *2017 IEEE international conference on cloud engineering (IC2E)*, pages 138–148. IEEE, 2017.
- [20] Federico Rizzo, Giovanni Luca Spoto, Paolo Brizzi, Dario Bonino, Giuseppe Di Bella, and Pino Castrogiovanni. Beekup: A distributed and safe p2p storage framework for ioe applications. In *2017 20th Conference on Innovations in Clouds, Internet and Networks (ICIN)*, pages 44–51. IEEE, 2017.
- [21] Stratis Ioannidis and Peter Marbach. Absence of evidence as evidence of absence: A simple mechanism for scalable p2p search. In *IEEE INFOCOM 2009*, pages 576–584. IEEE, 2009.
- [22] Jos   Santos, Tim Wauters, Bruno Volckaert, and Filip De Turck. Towards dynamic fog resource provisioning for smart city applications. In *2018 14th International Conference on Network and Service Management (CNSM)*, pages 290–294. IEEE, 2018.
- [23] Genc Tato, Marin Bertier, and C  dric Tedeschi. Koala: Towards lazy and locality-aware overlays for decentralized clouds. In *2018 IEEE 2nd International Conference on Fog and Edge Computing (ICFEC)*, pages 1–10. IEEE, 2018.
- [24] Frank HP Fitzek et al. On network coded distributed storage: How to repair in a fog of unreliable peers. In *2016 International Symposium on Wireless Communication Systems (ISWCS)*, pages 188–193. IEEE, 2016.
- [25] Genc Tato, Marin Bertier, and C  dric Tedeschi. Designing overlay networks for decentralized clouds. In *2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, pages 391–396. IEEE, 2017.
- [26] Yawei Zhao, Fei Cai, Junjie Xie, Lailong Luo, Xiaoqiang Teng, Honghui Chen, and Weijie Kong. A new dht supporting multi-attribute queries for grid information services. In *2013 IEEE 10th International Conference on High Performance Computing and Communications & 2013 IEEE International Conference on Embedded and Ubiquitous Computing*, pages 1675–1680. IEEE, 2013.
- [27] Jaeho Kim and Jang-Won Lee. Openiot: An open service framework for the internet of things. In *2014 IEEE world forum on internet of things (WF-IoT)*, pages 89–93. IEEE, 2014.
- [28] Edward Wang and Richard Chow. What can i do here? iot service discovery in smart cities. In *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*, pages 1–6. IEEE, 2016.
- [29] Deze Zeng, Lin Gu, Song Guo, Zixue Cheng, and Shui Yu. Joint optimization of task scheduling and image placement in fog computing supported software-defined embedded system. *IEEE Transactions on Computers*, 65(12):3702–3712, 2016.
- [30] Heng Shi, Nan Chen, and Ralph Deters. Combining mobile and fog computing: Using coap to link mobile device clouds with fog computing. In *2015 IEEE International Conference on Data Science and Data Intensive Systems*, pages 564–571. IEEE, 2015.
- [31] Lin Gu, Deze Zeng, Song Guo, Ahmed Barnawi, and Yong Xiang. Cost efficient resource management in fog computing supported medical cyber-physical system. *IEEE Transactions on Emerging Topics in Computing*, 5(1):108–119, 2015.
- [32] Mohammed A Hassan, Mengbai Xiao, Qi Wei, and Songqing Chen. Help your mobile applications with fog computing. In *2015 12th Annual IEEE International Conference on Sensing, Communication, and Networking-Workshops (SECON Workshops)*, pages 1–6. IEEE, 2015.
- [33] Fahed Alkhabbas, Ilir Murturi, Romina Spalazese, Paul Davidsson, and Schahram Dustdar. A goal-driven approach for deploying self-adaptive iot systems. In *2020 IEEE International Conference on Software Architecture (ICSA)*, pages 146–156. IEEE, 2020.
- [34] Rakesh Jain and Samir Tata. Cloud to edge: distributed deployment of process-aware iot applications. In *2017 IEEE International Conference on Edge Computing (EDGE)*, pages 182–189. IEEE, 2017.
- [35] Gian Paolo Jesi, Alberto Montesoro, and Ozalp Babaoglu. Proximity-aware superpeer overlay topologies. In *IEEE International Workshop on Self-Managed Networks, Systems, and Services*, pages 43–57. Springer, 2006.
- [36] Athman Bouguettaya, Munindar Singh, Michael Huhns, Quan Z Sheng, Hai Dong, Qi Yu, Azadeh Ghari Neiat, Sajib Mistry, Boualem Benatalah, Brahim Medjahed, et al. A service computing manifesto: the next 10 years. *Communications of the ACM*, 60(4):64–72, 2017.
- [37] Suparna De, Payam Barnaghi, Martin Bauer, and Stefan Meissner. Service modelling for the internet of things. In *2011 Federated Conference on Computer Science and Information Systems (FedCSIS)*, pages 949–955. IEEE, 2011.
- [38] Hyperic Sigar. "https://github.com/hyperic/sigar/wiki/overview". [Online accessed: January 2020].
- [39] Emanuele Golderi, Giuseppe Rossi, and Alberto Torelli. Assolo, a new method for available bandwidth estimation. In *2009 Fourth International Conference on Internet Monitoring and Protection*, pages 130–136. IEEE, 2009.
- [40] CouchDB. https://www.couchdb.com/. [Online accessed: January 2020].
- [41] IBM. Nmon. https://developer.ibm.com/technologies/systems/articles/aunmon\_analyzer/. [Online accessed: January 2020].
- [42] Changhun Bae and Wayne E Stark. A tradeoff between energy and bandwidth efficiency in wireless networks. In *MILCOM 2007-IEEE Military Communications Conference*, pages 1–7. IEEE, 2007.
- [43] Antonio Brogi, Stefano Forti, Carlos Guerrero, and Isaac Lera. How to Place Your Apps in the Fog-State of the Art and Open Challenges. *arXiv preprint arXiv:1901.05717*, 2019.



**Ilir Murturi** received the M.Sc. degree in computer engineering from the Faculty of Electrical and Computer Engineering, University of Prishtina, Prishtina, Kosovo. He is currently working toward the Ph.D. degree in Edge Computing under the supervision of Prof. S. Dustdar at the Distributed Systems Group, TU Wien, Vienna, Austria. His current research interests include the Internet of Things, Edge Computing, crowdsourcing, privacy, and smart cities.



**Schahram Dustdar (Fellow, IEEE)** is currently a Professor of computer science with the Distributed Systems Group, TU Wien, Vienna, Austria. He is an elected member of the Academia Europaea, where he is also the Chairman of the Informatics Section. He serves on the Editorial Board of IEEE Internet Computing and the IEEE Computer Magazine. He was a recipient of the ACM Distinguished Scientist Award in 2009, the IBM Faculty Award in 2012, and the IEEE TCSVC Outstanding Leadership Award for outstanding leadership in services computing in 2018. He is also the Co-Editor-in-Chief of the ACM Transactions on Internet of Things and the Editor-in-Chief of Computing (Springer). He is also an Associate Editor of the IEEE Transactions on Services Computing, the IEEE Transactions on Cloud Computing, the ACM Transactions on the Web, and the ACM Transactions on Internet Technology.