

Application-Aware Resource Allocation and Data Management for MEC-assisted IoT Service Providers

Simone Bolettieri^a, Raffaele Bruno^{a,*}, Enzo Mingozzi^b

^a*IIT-CNR, Via G. Moruzzi 1, 56124 Pisa, ITALY*

^b*Dept. of Information Engineering, University of Pisa, L.go L.Lazzarino, 1. I-56122 Pisa. ITALY.*

Abstract

To support the growing demand for data-intensive and low-latency IoT applications, Multi-Access Edge Computing (MEC) is emerging as an effective edge-computing approach enabling the execution of delay-sensitive processing tasks close to end-users. However, most of the existing works on resource allocation and service placement in MEC systems overlook the unique characteristics of new IoT use cases. For instance, many IoT applications require the periodic execution of computing tasks on real-time data streams that originate from devices dispersed over a wide area. Thus, users requesting IoT services are typically distant from the data producers. To fill this gap, the contribution of this work is two-fold. Firstly, we propose a MEC-compliant architectural solution to support the operation of multiple IoT service providers over a common MEC platform deployment, which enables the steering and shaping of IoT data transport within the platform. Secondly, we model the problem of service placement and data management in the proposed MEC-based solution taking into account the dependencies at the data level between IoT services and sensing resources. Our model also considers that caches can be deployed on MEC hosts, to allow the sharing of the same data between different IoT services with overlapping geographical scope, and provides support for IoT services with heterogeneous QoS requirements, such as different frequencies of periodic task execution. Due to the complexity of the optimisation problem, a heuristic algorithm is proposed using linear relaxation and rounding techniques. Extensive simulation results demonstrate the efficiency of the proposed approach, especially when traffic demands generated by the service requests are not uniform.

Keywords: IoT, Mobile Edge Computing (MEC), service placement, data management, traffic shaping, application-aware caching, optimisation.

1. Introduction

Recent years have witnessed the explosion of the Internet of Things (IoT) in terms not only of the number of deployed devices but also of the type of devices embedded with high-data-rate sensors (e.g., wearable devices, smartphones, vehicles) [1]. This new generation of IoT devices is enabling many new IoT applications, such as audio recognition, augmented reality, or cooperative driving, which involve delay-sensitive and computing-intensive processing tasks of a non-trivial amount of raw data [2]. Besides, the widespread deployment of IoT devices has fostered the development of IoT platforms and products to support a variety of IoT-based services, such as device and network management, application development and service provisioning [3].

In past years, the conventional approach followed by IoT platforms that offer solutions to build and deploy distributed IoT applications was data-centric and based on

cloud technologies [4]. Specifically, in such IoT frameworks, the streams of data that originate from IoT devices are transferred to a distant cloud infrastructure, where are stored and processed and, then, offered on-demand to multiple applications and end-users as a shared pool of virtualised sensing resources. Indeed, all top cloud-computing providers have expanded their offer with proprietary IoT solutions at all cloud levels (i.e. IaaS, PaaS and SaaS), being Azure IoT Suite¹ and AWS IoT² two of the most notable examples. However, the centralised management of IoT data in the cloud has proven problematic due to the overload in the access and core network and the communication latency between end-users and the cloud. To cope with these capacity and latency constraints, as well as to fulfil the QoS requirements of new IoT applications, *edge computing* has emerged as an alternative computing paradigm to support IoT applications [5, 6, 7]. According to this paradigm, computation resources are made available at the edge of the network, in the proximity to end-devices, significantly reducing bandwidth demands and

*Corresponding author

Email addresses: s.bolettieri@iit.cnr.it (Simone Bolettieri), r.bruno@iit.cnr.it (Raffaele Bruno), enzo.mingozzi@unipi.it (Enzo Mingozzi)

¹<https://azure.microsoft.com/en-us/overview/iot/>

²<https://aws.amazon.com/iot>

communication latency. Different architectures have been proposed to deploy edge computing platforms. Still, the Multi-access Edge Computing (MEC) technology, which allows deploying computing resources at various locations of the RAN part of 5G mobile networks [8] and, more generally, at different network edge devices such as BSs and WiFi access point, is rapidly gaining momentum to facilitate the deployment of delay-sensitive IoT applications [9]. Since MEC servers are much more resource-constrained than cloud data centres, it is a compelling challenge to design efficient resource management policies to maximise the number of IoT services that can be deployed while satisfying their QoS constraints.

Given the importance of MEC for future networks, there is already an extensive literature addressing resource-allocation problems in MEC systems when the service requests are received from mobile users. Specifically optimal solutions have been proposed for computation offloading [10, 11, 12], service placement [13], and routing of user requests to MEC servers [14]. However, existing solutions do not properly take into consideration some fundamental properties of emerging delay-sensitive and data-intensive IoT applications. First, many IoT applications require the periodic execution of computing tasks on a specific type of sensed data that is collected from *dispersed data sources*, and not directly uploaded by the user requesting the services that is offloaded on the edge platform. Thus, the placement of the requested IoT service in the MEC infrastructure not only impact the utilisation of available computing resources, but it also influences the routing of IoT data within the MEC network. This means that the service placement problem should be formulated considering the dependencies between data flows and data processing, as well as data redundancy. Second, existing resource allocation algorithms for MEC system mostly assume that service requested by end-users are executed independently, use separate input data, and are assigned a different set of resources. A few recent studies have introduced the concept of *shareable resources* to account for services that can share the same copy of data and code (e.g., libraries for video analytics) [15, 2]. In this work, we advocate the use of *data caches* to enable application information sharing and to reduce the bandwidth demands both in the access and core networks. Finally, in many cases the owners of the IoT infrastructure and the MEC infrastructure are separated, and they do not exchange their business-related data. For instance, the MEC system may not be aware of which data should be shared between the deployed IoT services, which complicates determining a suitable mapping between data sources and computing resources.

To overcome the above-mentioned shortcomings of a MEC-based IoT system, we propose an architectural solution to enable an IoT service provider to transparently interact with a MEC infrastructure integrated in a 5G network, which offers *edge IaaS capabilities*. It is important to point out that the edge IaaS model is the most popular service model for MEC platforms [16, 17]. In this

case, the edge resources are provided as virtual execution environments (VMs or containers) of variable sizes. Besides, virtualised and independent logical networks (called *network slices*) on the same physical network infrastructure are supported. Our approach enables the steering and shaping of IoT traffic within the MEC system, and the dynamic allocation of edge resources to IoT applications. The proposed framework is compliant to the ETSI MEC architecture as it relies on an entity, that we call *IoT Service Manager*, which is instantiated as a standard MEC application with specific support for data sharing between IoT services. In addition, our approach does not require the IoT service provider to disclose information to the 5G network operator about the traffic characteristics of the deployed IoT applications.

In the proposed MEC-enabled framework for IoT applications, a key problem consists in determining the MEC host that should execute the IoT services in order to fulfil the QoS requirements of the deployed services while minimising the consumption of network and edge resources. In this work, we propose a general (nonlinear) formulation of the optimal service placement, resource allocation and data management to maximise the number of admitted services while minimising the total cost of communication, computation and storage resources that the IoT service provider shall lease from the 5G network operator. Our formulation caters for the sharing of IoT data between IoT services with overlapping geographical scope, aiming to reduce the IoT traffic data that needs to be routed between MEC servers. Our approach allows the IoT service provider to keep full control of the cost of service deployment. Moreover, we consider the heterogeneous QoS requirements of IoT applications, such as the frequency of periodic task execution, to dynamically adjust the reporting rate of active data sources.

To address the computation efficiency issues of the nonlinear problem, we leverage linear relaxation and rounding techniques [18] to develop a heuristic algorithm that can determine a feasible solution of the original problem. Using simulation results, we show that in many practical scenarios, our algorithm finds close-to-optimal solutions with a significantly lower time complexity than the exact solver. To assess the heuristic performance, we carry out an extensive evaluation using three benchmarks, namely two greedy-based heuristics and a solution inspired by an operator placement algorithm for data stream processing applications that is proposed in [19]. Our results show that the proposed heuristic is more efficient than the considered benchmarks when traffic demands are not uniform. Furthermore, the proposed approach ensures a more uniform utilisation of the edge and network resources

The rest of this paper is organised as follows. Section 2 discusses the reference scenario and the use cases. Section 3 presents the proposed MEC architectural solution. The optimisation model is developed in Section 4. Section 5 explains the proposed heuristic algorithm. Section 6 presents the evaluation of our proposed algorithms, while

Section 7 reviews our contribution compared to related works.

2. Architecture and Use Cases

In this work, we envision a scenario in which the IoT platform offering IoT services to end-users and the 5G mobile network operator are two separate stakeholders in the IoT value chain. In the following, we describe the high-level architecture of the IoT system we consider, which is also illustrated in Figure 1. Then, we present a simple scenario and a few application use cases motivating the considered architecture. This discussion will provide the background to explain how the IoT platform and the 5G system can transparently interact through the use of MEC technologies without exchanging critical business-related information about their customers or infrastructure deployment.




Figure 1: Illustration of the high-level system architecture. Different colours represent different types of IoT services. Dashed coloured lines represent IoT data traffic. The figure shows that three services belonging to two service types are deployed in the MEC system.

2.1. High-level architecture

The *IoT Service Provider (IoTSP)* for brevity) is an organisation (e.g. a private company, a municipality, a utility provider, etc.) that owns and manages a sensing infrastructure, which comprises a potentially large number of different *IoT terminals*. We use the concept of IoT terminal as a generalisation of the classical model of IoT devices – that is to say, things that are equipped with micro-controllers, sensing/actuating capabilities, transceivers for digital communications, and protocol stacks to connect (typically wirelessly) to the Internet. For instance, an IoT terminal could be: *i)* a high-data-rate sensor (e.g. a surveillance camera deployed in a city), *ii)* a multi-sensing platform (e.g., a weather station), *iii)* a crowdsourced mobile device (e.g. a smartphone, a connected car), or *iv)* an *IoT gateway* that acts as a proxy of the sensing resources available in an in-situ WSN deployed in the area.

Then, the IoTSP may offer to end-users a *catalogue of different data-intensive and delay-sensitive IoT applications* to choose from. The IoTSP’s customers subscribe to a service from the pool of available ones selecting the demanded QoS parameters for the service, such as data granularity, frequency of service updates, geographical scope of the service execution [20]. The IoTSP aims at the maximisation of its revenues by satisfying the largest number of customers’ service requests. However, the IoTSP may not have the capabilities that are required to transfer, store, and process the huge amount of data which needs to be collected from geographically distributed IoT terminals to fulfil the QoS requirements of requested services. In the envisioned architecture, the IoTSP relies on the communication and edge-based computing resources that are provided by a *5G Mobile Network Operator (5G-MNO)* for brevity).

From a high-level perspective, the 5G-MNO implements a radio access network, which consists of dense deployment of heterogeneous base stations with different coverage ranges, supporting multiple communication standards (5G NR, LTE and WiFi). Then, the 5G network provides the connectivity platform enabling the IoTSP to collect the data generated from the IoT terminals under the coverage area of deployed base stations. In addition, a MEC system is integrated within the 5G network to offer edge-computing resources to its users. In this study, we assume that the *5G-MNO is a provider of edge IaaS capabilities*, i.e. it provides to its customers distributed compute and storage capabilities, usually in the form of a virtualised pool of resources [16]. Then, the IoTSP is one of the customers of the MEC system: *i)* it specifies the resources it needs across the MEC infrastructure to deploy the IoT applications requested by end-users (e.g. virtual machines or containers based on the virtualisation technology used by the 5G-MNO), and *ii)* it pays for resources according to use over time [17]. We recall that the core technical challenge we address in this study is to define a policy allowing the IoTSP to select the optimal set of network, computation and storage resources in the edge IaaS enabled by MEC platform. It is worth pointing out that we have considered a 5G-MNO only as an exemplification of a mobile communication technology with specific support for IoT environments. However, the same concept could be applied to other radio access technologies and network architectures by leveraging the ability of MEC technologies to enable heterogeneous networks (see Section 3.1 for a more detailed discussion on this aspect).

While the details of the MEC reference architecture will be discussed in Section 3.1, here we remind that one of the key components in a MEC system is the MEC host. Basically, a MEC host is a server that can be physically deployed in different locations between the base stations, a network aggregation point and the core network [21]. As shown in Figure 1, each MEC host can directly serve a given geographical area that comprises multiple base stations. Then, dedicated communication links are estab-

lished between the MEC host and the base stations in its serving area (fronthaul links), between MEC hosts (backhaul links), and between MEC hosts and the core network. The placement of the IoT application in the MEC infrastructure and the topology of the MEC network affect the routing of the traffic from the data sources to the targeted IoT applications.

2.2. Motivating scenario: IoT services for a Smart City

Inspired by previous studies [22, 9, 20], we show how a Smart City could benefit from the envisioned IoT system, and motivate the deployment of IoTSPs. Nowadays, it is a well-established concept that the IoT paradigm is a key enabler for realising the Smart City vision [23]. Indeed, an *urban IoT* system is required to collect a large variety of data on city resources and infrastructures, which can then be used for the optimisation of the public services that are offered to citizens. As observed before, the advancements on IoT technologies and devices is paving the way to a variety of new business models for smart cities that require the development of IoT applications with high QoS requirements. First, there is a plethora of IoT platforms³ that enable almost autonomous IoT device and network management. Thus, it is increasingly common that municipalities own their urban sensing infrastructures, acting as an IoTSP for different end-users such as police, public transportation companies, domestic energy providers, and business associations. In addition, besides classical sensors for environmental monitoring, car counting, on-street parking detection, etc., the use of high-resolution video cameras, microphones, and other high-data-rate sensors have been rapidly expanding in the urban environment, opening the way to new IoT applications that involve various types of perception-related tasks [1]. An important example of such computation-intensive tasks that require continuous processing of massive amount of data from high data-rate sensors is audio-visual event classification. Specifically, there is a vast literature on the design of various methodologies for scene analysis and event classification that leverages audio and visual information [24, 25]. Recent advancements in ML/AI techniques allow to develop complex real-time video analysis, which can be used to identify and classify persons, objects, events, define and use event-based rules (entering/exiting area, leaving/removing objects), and counting objects (vehicles) or people. The video analytics can be further enhanced by sound data. It is important to note that different applications may have different requirements in terms of resolution of the recorded images or quality of the audio recordings, frequency of data acquisition, and spatial granularity of the multimedia data. Clearly, the higher the data quality, the higher the granularity, and the higher the data rate of the sensory data streams to be processed. In addition, enhanced video/sound analytics

with real-time processing creates the need for executing the analytics functions close to the sources of these multimedia data streams for better performance and significant savings in communication latencies [26].

3. Solutions for MEC-Assisted IoT Service Providers

In this section, we first briefly overview the MEC specification, with a focus on the components and capabilities that are the most relevant to our solution. Then, we provide a detailed description of the proposed IoT Service Manager component, which is the MEC-based entity implementing the orchestration of IoT applications.

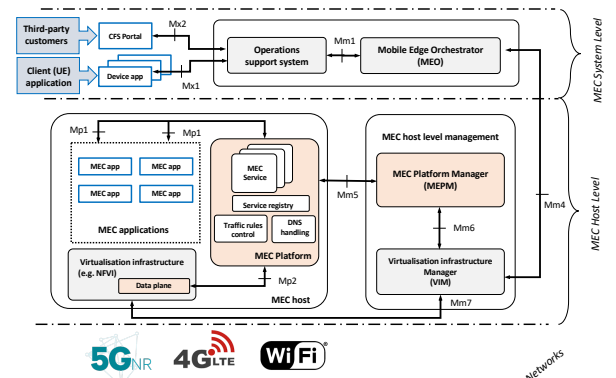


Figure 2: (Simplified) MEC reference architecture and functional components.

3.1. ETSI MEC reference architecture.

In 2014, ETSI established an Industry Specification Group (ISG) to specify a new standard and reference architecture, called Mobile Edge Computing (MEC), to complement the cloud radio access networks (C-RAN) of cellular operators with cloud-computing capabilities. However, over the years the scope of the MEC concept has been expanded to allow the deployment of MEC servers over multi-tier heterogeneous networks (for short, HetNets) and different elements of the network edge, such as BSs (aka eNB in 4G and gNB in 5G), optical network units, radio network controller sites, and WiFi access points [27]. To reflect this change of perspective in the MEC concept, the ETSI MEC initiative was renamed as Multi-access Edge Computing. Thus, not only MEC technologies will be important for 5G systems, but also for HetNets in general. Figure 2 shows a simplified functional view of this reference architecture, focusing on interfaces and components that are the most relevant to our proposal [8]. Three main logical entities constitute the core part of the MEC architecture: (i) the MEC host, (ii) the MEC system-level management, and (iii) the MEC host-level management. Basically, a MEC host is an edge node that provides physical computing, storage and communication resources, and

³<https://www.iot-survey.com/platform-provider>

includes a MEC platform and a virtualisation infrastructure. The former component provides the functionalities needed to run the MEC applications on top of the virtualisation environment offered by the MEC host (e.g. through VM or container-based technologies).

Furthermore, the MEC platform enables MEC applications to provide and consume MEC services. Specifically, a MEC service is exposed with a RESTful API, which MEC applications can subscribe to if authorised [28, 29]. A number of MEC services have been already defined in the MEC reference architecture, which are essential enablers for our IoT Service Manager. For instance, the “Location” service that, when available, provides authorised applications information about the radio nodes associated to a MEC host, as well as the location of all UEs currently served by these radio nodes. The “Bandwidth Manager” service, when available, allows allocation of bandwidth to specific traffic routed to and from MEC applications.

The virtualisation infrastructure of a MEC host includes a data plane that executes the traffic steering and filtering rules that are set by the MEC platform. Such traffic rules offer the capability to route and shape traffic between a specific UE and a MEC application, or among MEC applications. It is worth pointing out that the MEC platform can receive specific traffic rules from the MEC platform manager (MEPM) but also *authorised MEC applications can request the activation of MEC application traffic rules dynamically* [8]. As explained later, this capability is essential to allow the IoTSP to route IoT traffic between the IoT service Managers. In addition to managing traffic rules, the MEPM: (i) manages the full life cycle of the MEC applications running on MEC hosts; (ii) informs the MEC system management level of application-related events; and (iii) provides all the basic functionalities to run applications on a specific MEC host and to enable these applications to discover, produce or consume MEC services. The MEPM operates in cooperation with the Virtualisation infrastructure manager (VIM), which allocates the virtualised resources of the MEC hosts, maintain status information on the available resources, and monitors application faults and performance.

The core component of the control and management plane at the system level is the MEC orchestrator (MEO), which is responsible for selecting the MEC hosts to which MEC applications will be deployed based on constraints, such as latency, available resources, and available MEC services. To perform this task, the MEO maintains a broad set of information about the resource configuration and usage in the deployed MEC hosts and the topology of the entire MEC system. The service requests processed by the MEO are only the ones that are authorised by the Operating Support System (OSS) of the telco operator. The requests to run an application in the MEC system can originate from a device application, i.e. a MEC-aware component of the client application running on the user’s device (e.g. UE). Alternatively, *an operator’s third-party customer (e.g. an enterprise) can register with the MEC*

system to utilise and manage edge capabilities. In our scenario, an IoTSP can be seen as one of these third-parties that are allowed to instantiate applications on the MEC system, paying a fee based on resource usage [17].

3.2. A MEC-enabled IoT Service Manager

One of the goals pursued by an IoTSP is to maintain the control of the IoT data that is shared and routed to the IoT applications that are running in the MEC system. To achieve this objective, we envision that the IoTSP requests the 5G-MNO to host a software instance of a component, which we call *IoT Service Manager (IoT-SM for brevity)*, in each MEC host of the MEC system. The primary purpose of this component is to enable the dynamic loading/unloading of IoT services. Thus, it can be regarded as a sort of *middlebox* embedded into a MEC edge server specifically designed to host IoT services. However, this entity provides also several additional features: (i) it dynamically *scale* the usage of edge resources (memory, CPU, bandwidth) *within the constraints* set by the MEC system-level management, (ii) it supports the dynamic loading of IoT application services and their related data (e.g. libraries, databases, processing code); (iii) it manages traffic rules to steer IoT traffic among other instances of IoT-SMs, and between local IoT services; (iv) it supports a data cache and brokering service to enable reuse of IoT data between collocated IoT services; (v) it shapes the traffic that is generated from the activated IoT terminals under the serving area of the MEC host; and (vi) it maintains a map of the IoT terminals that are manageable from a MEC host.

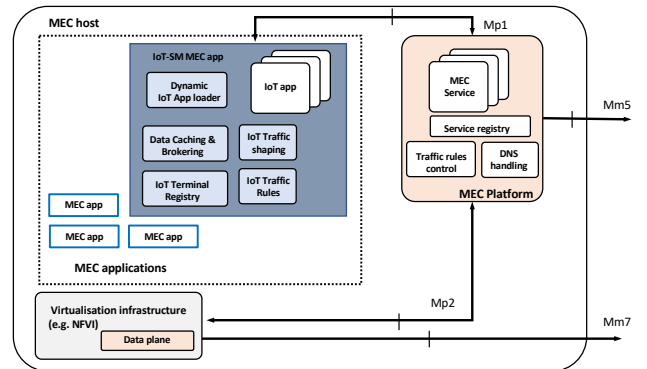


Figure 3: Functional architecture of the IoT Service Manager.

The functional blocks of the IoT-SM that support the above features are described in the following and illustrated in Figure 3. The *IoT Terminal Registry (RD)* functional block is in charge of maintaining a registry of the IoT terminals that are within the serving area of the MEC host on which the IoT-SM is running. To build such list of available data sources, the IoT-SM relies on the above-mentioned Location Service [30], which not only provides the list of UEs in a particular location area, but also the

location information of all UEs currently served by the radio node(s) associated with the MEC host. We remind that the 5G-MNO implements its specialised policy to associate the terminals to available BSs (e.g., to maximise spectral efficiency or improve load balancing [31]). Thus, the IoTSP has no control about which MEC host would serve individual IoT terminals. The *RD* also communicates with its counterparts running on the other MEC hosts to build a global map of IoT terminal–MEC host associations. The *Dynamic IoT App Loader (LOADER)* functional block is the core element of the proposed framework as it handles the management of the IoT applications and their associated resources. Specifically, it checks the current status of available resources and already deployed IoT services and, if a new IoT service can be loaded without violating the constraints set by the MEO, it downloads the library/code that is needed to run the IoT application, and requests to the MEC platform the necessary additional resources. It is worth pointing out that our architectural model is amenable to both centralised and distributed formulations of the resource allocation problem. In the former case, one IoT-SM is elected by the IoTSP as the leader manager that operates as a logically centralised orchestrator. In the latter case, an overlay exists between the IoT-SMs that take local decisions on resource allocation in a distributed manner [32]. The *LOADER* operates in collaboration with the *Traffic Shaping (TS)* and *Data Caching & Brokering (DCB)* functional blocks. The former is in charge of shaping the streams of IoT data that are needed to execute the IoT applications. Specifically, it executes the decisions received by the *LOADER* on which IoT terminals to activate, and which data streaming rate to assign to them. We remind that the optimal configuration of sensing resources that are assigned to an IoT application depends on the QoS parameters of the admitted applications (e.g. frequency of service execution, geographical scope of the service), and the bandwidth constraints. On the other hand, the *DCB* supports the cache of IoT data to enable different IoT applications instantiated on the same IoT-SM to share and reuse the same data. Furthermore, it is responsible for route traffic between IoT applications running on different MEC hosts, thus allowing the shared use of sensing resources that can be dispersed in a large deployment area. Finally, the *Traffic Rules (TR)* functional block is in charge of informing the MEC platform of the traffic rules that are needed to redirect the traffic between the IoT-SMs.

From the perspective of the MEC reference architecture and IoT-SM implementation, we observe that the proposed software component can be architected as a conventional MEC application, which is instantiated in the MEC system in response to a request of the IoTSP. Then, the Mp1 interface is used by the IoT-SM MEC applications to access the needed MEC services and to activate traffic rules in the MEC platform. Our architectural model is well aligned with the one specified in [20]. However, the model given in [20] focuses only on the virtualisation of IoT Gateways.

Finally, our system solution requires that the virtualisation technology used by the MEC system to implement the virtualisation infrastructure supports dynamic resizing of VMs at run-time. This is needed to allow the running IoT-SMs to dynamically scale up and down their allocated CPU and memory resources following the downloading of new IoT applications or the termination of previously admitted IoT applications, without restarting running IoT applications. It is important to point out that vertical scaling (i.e., changing CPU and RAM resources allocated to a running server) is a feature currently offered by many cloud PaaS platforms [33].

Implementation considerations. In past years, several experimental and open-source fog computing platforms were developed to allow IoT service providers to automatically program, manage and orchestrate dynamic data processing flows over cloud and edges, being FogFlow⁴ and Stacks4Things⁵ two relevant examples. Following a similar development path, there are various MEC platform under development at the moment. ETSI MEC ISG has recently established a new working group, called Deployment and Ecosystem Development working group (WG DECODE), to facilitate and promote the use of open-source components to build MEC systems by allowing applications that are virtualised in the edge, to access network and users' information from the local node using MEC-standardised open APIs⁶. In addition, non-profit technology consortia around the world, such as the Open Edge Computing Initiative and the Linux Foundation, are promoting open-source projects for edge computing platforms, which also provide supporting services that can be used to develop the IoT-SM concept.

4. System Model and Problem Formulation

The goal of this section is to formulate an optimisation problem for the joint service placement, allocation of resources and data management in the MEC system described above. First, we introduce the system model, the used notation (also shown in Table 1) and the modelling assumptions. Then, we formulate the problem in terms of constraints and optimisation objectives.

4.1. System model

Our model is aligned with the architecture illustrated in Figure 1. We consider an area A , where an IoTSP has deployed a set \mathcal{O} of IoT terminals, which could be part of a larger in-situ sensing platform deployed in the city. An IoT terminal is described by o_t , $t = 1, \dots, N$, where N is the total number of IoT terminals. Each IoT terminal o_t hosts a set \mathcal{R}_t of sensing resources, while $\mathcal{R} = \cup_{t \in [1, N]} \mathcal{R}_t$ is

⁴<https://fogflow.readthedocs.io/en/latest/>

⁵<http://stack4things.unime.it/>

⁶<https://forge.etsi.org/rep/mec>

Table 1: List of notation.

Parameter	Definition
Sets, vectors	
$\mathcal{W} = \{c_k\}$	set of cells; subscript k refers to cell c_k
$\mathcal{R} = \{r_i\}$	set of sensing resources; subscript i refers to resource r_i
$\langle l(i), c(i) \rangle$	data type measured by resource r_i , and cell where it is located
$\mathcal{O} = \{o_t\}$	set of IoT terminals; subscript t refers to terminal o_t
$\mathcal{BS} = \{bs_h\}$	set of base stations; subscript h refers to node bs_h
$\mathcal{H} = \{h_m\}$	set of MEC hosts; subscript m refers to host h_m
$\mathcal{BS}_m \subseteq \mathcal{BS}$	set of base stations that are served by MEC host h_m
$\mathcal{S} = \{s_j\}$	set of IoT services requested by end-users; subscript j refers to service s_j
$\langle l_j, \mathcal{W}^j, \lambda_j, \delta_j, \gamma_j^p \rangle$	Requirement vector of IoT service s_j : data type, geographical scope (i.e., set of cells that shall be covered to serve the s_j), frequency of service execution, processing load, persistent memory
$\mathcal{R}_t \subseteq \mathcal{R}$	set of sensing resources that are exposed by IoT terminal o_t
Decision variables	
x_{ij}	Binary variable indicating if the data collected by resource r_i is delivered to the IoT service s_j ($x_{ij} = 1$) or not
y_{mj}	Binary variable indicating if IoT service s_j is deployed on the MEC host h_m ($y_{mj} = 1$) or not
f_i	sampling rate assigned to sensing resource r_i
Auxiliary variables, constants	
ψ_{th}	binary variable indicating that terminal o_t is associated to base station bs_h
\hat{b}_l	payload size for a sample of data type d_l
$F_h, F_m, F_{m_1 m_2}$	total amount of IoT data transferred to base station bs_h , to MEC host h_m , and between MEC hosts h_{m_1} and h_{m_2} , respectively
$\langle \Delta_m, \Gamma_m \rangle$	maximum computing power and storage that can be allocated by MEC host h_m to running IoT services
$B_h^\uparrow, B_m, B_{m_1 m_2}$	uplink bandwidth capacity of base station bs_h , bandwidth capacity of fronthaul link to MEC host h_m , and bandwidth capacity of the link between MEC hosts h_{m_1} and h_{m_2} , respectively
$c_1^{bw}, c_2^{bw}, c^{cpu}, c^m$	unitary cost of wireless bandwidth, wired bandwidth, CPU cycles, storage

the whole set of sensing resources that are available in the reference area. Sensing resources are denoted by r_i . For simplicity, we assume that the area of interest is divided into a regular grid of K square cells, and the sensed data collected by an IoT terminal is associated to the cell where the IoT terminal is located. Cells are denoted by c_k , $k = 1, \dots, K$. Similarly to previous studies [34], we assume that a sensing resource can only collect data of a specific data type. Let L be the total number of data types of interest. Let us denote with $l(i) = 1, \dots, L$ the data type measured by resource r_i , and with $c(i) \in \mathcal{W}$ the cell where the resource r_i is located. Note that different attributes may be associated to a data type, such as the temporal and spatial resolution of the sensor reading, the energy dissipated by the sensor during the measurement, etc. For the sake of our analysis, we only consider the size (in bytes) of the message payload \hat{b}_l that is needed to upload the sensor reading for data type l .

The radio access network (RAN) of the 5G-MNO comprises H small base stations (SBSs) that are densely deployed to increase the network capacity and reduce latency by leveraging spatial reuse [35]. Let $\mathcal{BS} = \{bs_1, \dots, bs_H\}$ denote the set of available SBSs. An IoT terminal can be located under the (possibly) overlapping coverage regions

of multiple SBSs. It is out of the scope of this work to investigate cell association rules for base-station assignment (the interested reader is referred to [36] for a comprehensive survey). We assume that the 5G-MNO associates each IoT terminal to exactly one of the nearby SBSs, and a binary variable $\psi_{t,h} \in [0, 1]$ models the one-to-one mapping between IoT terminal $o_t \in \mathcal{O}$ and SBS $bs_h \in \mathcal{BS}$.

As described in Section 2, the area of interest is also served by a MEC system that consists of a set \mathcal{H} of MEC hosts. Each MEC host $h_m \in \mathcal{H}$ is responsible for managing a subset $\mathcal{BS}_m \subseteq \mathcal{BS}$ of the available SBSs. We assume that MEC hosts have limited storage and computing resources, and that are connected to the SBSs and among each other via links with limited capacity. Then, a portion of these resources (i.e. a *slice*) is pre-allocated by the 5G-MNO to the IoTSP tenant. Specifically, each MEC host h_m specifies the maximum storage capacity (i.e. hard disk) Γ_m and the maximum computation capacity (i.e. CPU cycles) Δ_m that can be assigned to the resident IoT-SM and the collocated IoT services. In addition, the 5G-MNO assigns to the IoTSP a maximum wireless uplink bandwidth capacity B_h^\uparrow to upload the IoT data from the terminals associated to SBS $bs_h \in \mathcal{BS}$. The fronthaul links connecting MEC host h_m to the radio nodes in \mathcal{BS}_m have a limited

total bandwidth equal to B_m . Finally, the backhaul link between MEC host h_{m_1} and MEC host h_{m_2} has fixed capacity B_{m_1, m_2} . In this study, we assume a *flat* network topology for the MEC system, in which all MEC hosts are interconnected. Investigating hierarchical edge architectures is left as future work. For the sake of simplicity, in this work, we assume that the capacity bounds of the IoT slice are constant during the life cycle of the IoT services. In other words, while the amount of network and edge resources that are leased by the IoTSP can be dynamically adjusted, the resource limits are set by the 5G-MNO and maintained fixed.

The IoTSP offers to its end-users a catalogue \mathcal{S} of IoT services. Let s_j denote the j -th IoT service requested by an end-user. A requirement vector $\langle l_j, \mathcal{W}^j, \lambda_j, \delta_j, \gamma_j^p \rangle$ is associated to a service s_j , which describes the specific QoS parameters of the IoT service, and the resources that an IoT-SM should request to the MEC host to deploy and execute the service dynamically. Specifically, this vector defines: (i) the type l_j of the sensor data that is needed for providing the service, (ii) the geographical scope of the service, expressed as the set \mathcal{W}^j of cells that constitutes the area of interest for the demanded service, (iii) the requested frequency λ_j of periodic service execution; (iv) the computation power δ_j that is needed to run the service, and (v) the memory γ_j^p that is occupied by the persistent state (i.e. code and data) of the service.

4.2. Problem formulation

For the sake of brevity, and when no ambiguity occurs, in the following we use the subscript index i to refer to sensing resource r_i , the subscript index t to refer to IoT terminal o_t , the subscript index m to refer to MEC server h_m , the subscript index k to refer to cell c_k , the subscript index h to refer to SBS bs_h , and the subscript index j to refer to service s_j . In addition, we introduce the indicator function $\mathbf{1}_a(X)$ of an element a of a set X as follows:

$$\mathbf{1}_a(X) = \begin{cases} 1 & a \in X \\ 0 & a \notin X \end{cases} \quad (1)$$

The IoTSP needs to decide: (i) the set of sensing resources and IoT terminals that should be activated to satisfy the QoS requirements of deployed IoT services; (ii) the IoT-SM in charge of loading the data and code that is necessary to execute the new IoT service; and (iii) how the data collected from the activated sensing resources should be distributed to the interested IoT services over the MEC network. As explained later, an IoTSP pursues a twofold objective: the maximisation of its revenue (in terms of successfully deployed IoT services), and the minimisation of the total cost of communication, computation and storage resources leased from to the 5G-MNO. To model these decisions we introduce three sets of optimisation variables: (i) $x_{ij} \in \{0, 1\}$ which indicates whether the data collected from sensing resource i is used as input data for the execution of service j ($x_{ij} = 1$) or not; (ii) $y_{mj} \in \{0, 1\}$

which indicates if the IoT-SM installed in MEC host m is in charge of executing service j ($y_{mj} = 1$) or not; and (iii) $f_i \in \mathbb{R}_{\geq 0}$ which specifies the frequency of data updates from resource i . We refer by *service placement* policy to the vector:

$$\mathbf{y} = (y_{mj} \in \{0, 1\} : m \in \mathcal{H}, j \in \mathcal{S}) , \quad (2)$$

and *data management* policy to the vectors:

$$\mathbf{x} = (x_{ij} \in \{0, 1\} : i \in \mathcal{R}, j \in \mathcal{S}) \quad (3)$$

$$\mathbf{f} = (f_i \in \mathbb{R}_{\geq 0} : i \in \mathcal{R}) . \quad (4)$$

The service placement and data management policies need to satisfy several constraints. First, an IoT service can be loaded by the IoT-SM installed in MEC host m only if this decision does not violate the capacity constraints set by the 5G-MNO. Specifically, the total computational load generated by deployed IoT services running on MEC host m must not exceed the computing capacity of that host,

$$\sum_{j \in \mathcal{A}} y_{mj} \times \delta_j \leq \Delta_m, \quad \forall m \in \mathcal{H} , \quad (5)$$

and the total amount of data stored in MEC host m for executing the IoT service must not exceed the memory capacity of that host:

$$\sum_{j \in \mathcal{A}} y_{mj} \times [\gamma_j^p + \hat{b}_{l_j} \times |\mathcal{W}^j|] \leq \Gamma_m, \quad \forall m \in \mathcal{H} . \quad (6)$$

It is worth noting that constraint (6) accounts for the fact that the *larger the geographical scope of a service, and the greater the amount of data that needs to be collected by the IoT-SM executing the service*. Second, each IoT service must run on exactly one MEC host:

$$\sum_{m \in \mathcal{H}} y_{mj} \leq 1, \quad \forall j \in \mathcal{S} . \quad (7)$$

Third, a sensing resource i is eligible for being activated only if it covers one of the cells that are requested by a deployed IoT service:

$$\sum_{\substack{i \in \mathcal{R} \\ l(i)=l_j}} \mathbf{1}_{c(i)}(\mathcal{W}^j) x_{ij} = \sum_{m \in \mathcal{H}} y_{mj} \quad \forall j \in \mathcal{S} \quad (8)$$

A core requirement of the formulated problem is that the frequency of data updates for an activated sensing resource must be able to satisfy the QoS requirements of all IoT services sharing the data produced by that sensing resource. This condition can be expressed as follows:

$$f_i \geq x_{ij} \times \lambda_j \quad \forall i \in \mathcal{R}, \forall j \in \mathcal{S} \quad (9)$$

$$f_i \leq \sum_{\substack{j \in \mathcal{S} \\ l(i)=l_j}} \mathbf{1}_{c(i)}(\mathcal{W}^j) x_{ij} \times \lambda_j \quad \forall i \in \mathcal{R} \quad (10)$$

Finally, the total bandwidth load that is generated by the IoT data traffic must not exceed the bandwidth capacity

of wireless and wired connections. To express this condition we introduce three auxiliary variables: (i) the total amount F_h of IoT data uploaded to SBS h , (ii) the total amount F_m of IoT data uploaded to MEC host m and (iii) the total amount F_{m_1, m_2} of IoT data transferred from MEC host h_{m_1} to MEC host h_{m_2} . It holds that

$$F_h = \sum_{t \in \mathcal{O}} \psi_{th} \left[\sum_{i \in \mathcal{R}_t} f_i \times \hat{b}_{l(i)} \right] \quad \forall h \in \mathcal{BS} \quad (11)$$

$$F_m = \sum_{h \in \mathcal{BS}_m} F_{bs_h} \quad \forall m \in \mathcal{H} \quad (12)$$

$$F_{m_1, m_2} = \sum_{j \in \mathcal{S}} y_{m_2 j} \sum_{h \in \mathcal{BS}_{m_1}} \sum_{\substack{t \in \mathcal{O} \\ \psi_{th}=1}} \sum_{\substack{i \in \mathcal{R}_t \\ l(i)=l_j}} x_{ij} \lambda_j \times \hat{b}_{l(i)} \quad (13)$$

It is worth pointing out that the set of equations (11)-(13) also define how the IoT traffic is routed within the MEC network. Owing to the model assumptions, it holds that the IoT traffic must not exceed the following bandwidth capacities.

$$F_h \leq B_h^\uparrow \quad \forall h \in \mathcal{BS} \quad (14)$$

$$F_m \leq B_m \quad \forall m \in \mathcal{H} \quad (15)$$

$$F_{m_1, m_2} \leq B_{m_1 m_2} \quad \forall m_1, m_2 \in \mathcal{H} \quad (16)$$

The goal of the IoTSP is to maximise its net utility when satisfying the service requests of the end-users. Intuitively, there is positive revenue for each IoT application that can be successfully served. Thus, the utility of running services is given by

$$\mathbf{J}_r = \sum_{j \in \mathcal{S}} \sum_{m \in \mathcal{H}} y_{mj} \cdot \quad (17)$$

On the other hand, the cost of an IoT service is generated by the price paid by the IoTSP to the 5G-MNO for the connectivity and edge resources. The total cost of the used resource is given by

$$\mathbf{J}_{edge} = \sum_{h \in \mathcal{BS}} c_1^{bw} F_h + \sum_{m \in \mathcal{H}} \left(c^{cpu} \delta_m + c^m \gamma_m + c_2^{bw} \left(F_m + \sum_{n \in \mathcal{H} \setminus \{m\}} F_{mn} \right) \right), \quad (18)$$

Using the above variables and functions, we can formulate our problem as follows:

$$\max_{\mathbf{x}, \mathbf{y}, \mathbf{f}} (\mathbf{J}_r - \gamma \mathbf{J}_{edge}) \quad (19)$$

$$\text{s.t. constraints (5) - (10), (14) - (16),}$$

where the weight parameter $\gamma \ll 1$ is used to assign a higher priority to the maximisation of the number of deployed services. Note that the problem (19) is a Mixed-Integer Non-linear Programming (MINLP) problem, which is NP-hard

in general. The non-linearity of the optimisation problem is due to constraint (16), which involves the cross-product of decision variables. However, we can apply standard linearisation techniques to convert the problem (19) into a Mixed Integer Linear Programming (MILP) problem [37]. Specifically, we introduce the following set of auxiliary binary variables:

$$\boldsymbol{\theta} = (\theta_{mij} \in \{0, 1\}, m \in \mathcal{H}, i \in \mathcal{R}, j \in \mathcal{S}). \quad (20)$$

and the following linear constraints

$$\theta_{mij} \leq x_{ij} \quad \forall m \in \mathcal{H}, i \in \mathcal{R}, j \in \mathcal{S} \quad (21)$$

$$\theta_{mij} \leq y_{mj} \quad \forall m \in \mathcal{H}, i \in \mathcal{R}, j \in \mathcal{S} \quad (22)$$

$$\theta_{mij} \geq (x_{ij} + y_{mj}) - 1 \quad \forall m \in \mathcal{H}, i \in \mathcal{R}, j \in \mathcal{S} \quad (23)$$

Using the above variables, it is possible to linearise constraint (16) by eliminating the cross-product terms as follows:

$$\sum_{j \in \mathcal{S}} \sum_{h \in \mathcal{BS}_{m_1}} \sum_{\substack{t \in \mathcal{O} \\ \psi_{th}=1}} \sum_{\substack{i \in \mathcal{R}_t \\ l(i)=l_j}} \theta_{m_2 i j} \lambda_j \times \hat{b}_{l(i)} \leq B_{m_1 m_2} \quad \forall m_1, m_2 \in \mathcal{H} \quad (24)$$

Owing to the new constraints, the equivalent MILP problem we can be written as:

$$\begin{aligned} \max_{\mathbf{x}, \mathbf{y}, \mathbf{f}, \boldsymbol{\theta}} & (\mathbf{J}_r - \gamma \mathbf{J}_{edge}) \\ \text{s.t.} & \text{constraints (5) - (10), (14) - (15), (21) - (24)}. \end{aligned} \quad (25)$$

It is important to point out that the exact solution of a MILP problem can be efficiently computed using the well-established Branch-and-Bound (B&B) algorithm [38]. Even though the worst-case complexity of such algorithms is still exponential, B&B methods can leverage structural properties of the problem and meta-heuristics to restrict the search space, thus reducing the time needed to compute an almost-exact solution, at least for small- to medium-scale problem instances.

5. Heuristic Algorithm

In this section, we describe a Linear Relaxation-guided heuristic algorithm (LR for brevity) for the solution of the problem (25). Our solution method is described in detail below and summarised in Algorithm 1 and Algorithm 2.

Our heuristic starts by solving a relaxed version of the original problem, called *P1* (line 1). Specifically, *P1* is obtained from (25) by removing the integrality constraint of each integer decision variable with the following collection of linear constraints:

$$0 \leq x_{ij} \leq 1 \quad \forall i \in \mathcal{R}, j \in \mathcal{S}, \quad (26)$$

$$0 \leq y_{mj} \leq 1 \quad \forall m \in \mathcal{H}, j \in \mathcal{S}, \quad (27)$$

$$0 \leq \theta_{mij} \leq 1 \quad \forall m \in \mathcal{H}, j \in \mathcal{S}, i \in \mathcal{R}. \quad (28)$$

Let us denote with $\{x_{ij}^\dagger\}$, $\{y_{mj}^\dagger\}$, $\{\theta_{mij}^\dagger\}$ the optimal relaxed solution values that are obtained in polynomial time using a linear program solver. If the optimal solutions of all relaxed variables are integer values, then a valid solution is found, and the algorithm is stopped (line 2). Otherwise, we round these values to obtain an integer solution, denoted by $\{\hat{x}_{ij}\}$, $\{\hat{y}_{mj}\}$ ⁷. To do so, we start from an empty solution set (line 7), and then tentatively try to allocate new services. It is important to point out that we leverage the fractional solution of $P1$ to express the probability that rounding a variable yields a resource allocation that is feasible and optimal. In addition, the algorithm keeps track of the utilisation of network and edge resources (line 8) to discard assignments that would violate capacity constraints.

The algorithm uses the fractional value y_{mj}^\dagger to define the order with which the services should be deployed. Specifically, let j' denote the service with the highest probability y_{mj}^\dagger to be deployed on a host m , for which an accept/reject decision is still pending (line 10). Then, the algorithm iteratively checks whether there is a host having a probability $y_{mj'}^\dagger > \mu$ (line 13) on which the service can be deployed without violating the problem constraints (line 16). The algorithm uses μ as a *cut-off* threshold to filter out unlikely assignments, thus reducing the number of iterations that are needed to find a sub-optimal solution. The order of host allocation follows the probability $y_{mj'}^\dagger$. If a feasible allocation on host m' exists, then $\hat{y}_{m'j'}$ is rounded to one (line 18). Afterwards, the algorithm checks if a feasible allocation of sensing resources that covers all the cells in $\mathcal{W}^{j'}$ exists. If yes, the algorithm updates accordingly the utilisation of edge resources (line 24). On the contrary, if there is not a feasible allocation of sensing resources, $\hat{y}_{m'j'}$ is rounded back to zero.

Algorithm 2 describes the procedure we implement to search for an allocation of sensing resources that can cover service j' if deployed on host m' without violating the bandwidth constraints. At the beginning, Algorithm 2 initialises auxiliary variables with the current state of bandwidth utilisation (line 2). Then, for each cell k in the set $\mathcal{W}^{j'}$ that is not yet covered (line 4), the procedure iteratively searches for the best sensing resource (i.e. the resource i' with the highest $x_{ij'}^\dagger$ value) to use for covering cell k . As in Algorithm 1, μ is a cut-off threshold used to filter out unlike assignments. Let us assume that the algorithm tentatively selects resource i' . First of all, a frequency $f_{i'}$ is selected to fulfil the requirement on frequency of data updates for service j' (line 11). Then, the algorithm checks if this sampling rate assignment violates one of the bandwidth constraints (line 12). If not, the sensing resource i' is allocated to service j' (i.e., $\hat{x}_{i'j'}$ is rounded to one), the cell is labelled as covered, the utilisation state of network links is temporarily updated, and

⁷We remind the θ_{mij} is an auxiliary decision variable used to remove the non-linear constraints with cross-product of variables. Thus, $\{\hat{x}_{ij}\}$ and $\{\hat{y}_{mj}\}$ do not depend on $\{\theta_{mij}^\dagger\}$.

Algorithm 1 - Heuristic Algorithm using linear relaxation and rounding techniques.

Input: $P1$ ▷ linear relaxation of original problem (25)

Output: $\{\hat{x}_{ij}\}$, $\{\hat{y}_{mj}\}$, $\{\hat{f}_i\}$ ▷ integer assignments and sampling rates

```

1: solve  $P1$  to obtain optimal solution  $\{x_{ij}^\dagger\}$ ,  $\{y_{mj}^\dagger\}$ ,  $\{\theta_{mij}^\dagger\}$ 
2: if  $(x_{ij}^\dagger, y_{mj}^\dagger, \theta_{mij}^\dagger \in \mathbb{Z})$  then
3:    $\hat{x}_{ij} \leftarrow x_{ij}^\dagger, \hat{y}_{mj} \leftarrow y_{mj}^\dagger$ ;
4:   return
5: end if

6:  $\Omega \leftarrow \{\emptyset\}$  ▷ Set of tested services
7:  $\hat{x}_{ij} \leftarrow 0, \hat{y}_{mj} \leftarrow 0, \hat{f}_i \leftarrow 0$  ▷ Initialise output variables
8:  $\hat{C}_m, \hat{S}_m, \hat{F}_h, \hat{F}_m, \hat{F}_{m_1 m_2}$  ▷ Auxiliary variables used to track resource utilisation

9: while  $(\mathcal{S} \setminus \Omega \neq \emptyset)$  do
10:   $j' \leftarrow \arg \max_{m \in \mathcal{H}, j \in \mathcal{S} \setminus \Omega} (y_{mj}^\dagger)$  ▷ higher priority is given to
    service  $j$  with the largest  $y_{mj}^\dagger$ 
11:   $\Sigma \leftarrow \{\emptyset\}$  ▷ Set of tested hosts
12:   $hostAllocated \leftarrow \text{False}$ 
13:  for all  $(m \in \mathcal{H}, y_{mj'}^\dagger > \mu)$  do
14:     $m' \leftarrow \arg \max_{m \in \mathcal{H} \setminus \Sigma} (y_{mj'}^\dagger)$  ▷ higher priority is given to
    host  $m$  with the largest  $y_{mj'}^\dagger$ 
15:     $\Sigma \leftarrow \Sigma \cup \{m'\}$ 
16:    if  $\text{CHECKEDGECONSTR}(m', j')$  then ▷ check
    if constraints (5) and (6) are not violated when deploying
    service  $j'$  on host  $m'$ 
17:       $hostAllocated \leftarrow \text{True}$ 
18:       $\hat{y}_{m'j'} \leftarrow 1$ 
19:      break ▷ exit for cycle
20:    end if
21:  end for

22:  if  $(hostAllocated = \text{True})$  then
23:    if  $(\text{ALLOCATESENSORS}(m', j') = \text{True})$  then
24:       $\hat{S}_m \leftarrow \hat{S}_m + \gamma_{j'}$ ;  $\hat{C}_m \leftarrow \hat{C}_m + \delta_{j'}$ 
25:    else
26:       $\hat{y}_{m'j'} \leftarrow 0$ 
27:    end if
28:  end if

29:   $\Omega \leftarrow \Omega \cup \{j'\}$ 
30: end while

```

Algorithm 2 - Function used to find a sensor allocation to cover service j without violating bandwidth constraints

```

1: function ALLOCATESENSORS( $m', j'$ )
2:    $F_h^\dagger \leftarrow \widehat{F}_h, F_m^\dagger \leftarrow \widehat{F}_m, F_{m_1 m_2}^\dagger \leftarrow \widehat{F}_{m_1 m_2}$ 
3:    $\widehat{\mathcal{W}} \leftarrow \{\emptyset\}$   $\triangleright$  set of tested cells
4:   while ( $\mathcal{W}^{j'} \setminus \widehat{\mathcal{W}} \neq \emptyset$ ) do
5:     pick  $k$  at random in  $\mathcal{W}^{j'} \setminus \widehat{\mathcal{W}}$ 
6:      $R_k \leftarrow \{\emptyset\}$ 
7:      $cellCovered \leftarrow \text{False}$ 
8:     for all ( $i \in \mathcal{R}, x_{ij'}^\dagger > \mu$ ) do
9:        $i' \leftarrow \arg \max_{i \in \mathcal{R} \setminus R_k, l(i)=l_{j'}, c(i)=c_k} (x_{ij'}^\dagger)$ 
10:       $R_k \leftarrow R_k \cup \{i'\}$ 
11:       $f_{i'}^\dagger \leftarrow (\lambda_{j'} \leq f_{i'}^\dagger ? \widehat{f}_{i'} : \lambda_{j'})$   $\triangleright$  set tentative
        sampling rate for resource  $i'$ 
12:      if CHECKNETWORKCONSTR( $f_{i'}^\dagger$ ) then
13:         $cellCovered \leftarrow \text{True}$ 
14:         $\widehat{x}_{ij'} \leftarrow 1$ 
15:         $[F_h^\dagger, F_m^\dagger, F_{m_1 m_2}^\dagger] \leftarrow \text{INCBAND}(\widehat{x}_{ij}, \widehat{y}_{mj}, f_{i'}^\dagger)$ 
16:        break  $\triangleright$  exit for cycle
17:      else
18:         $f_{i'}^\dagger \leftarrow \widehat{f}_{i'}$   $\triangleright$  revert to state before line 11
19:      end if
20:    end for
21:    if  $cellCovered = \text{False}$  then
22:      break  $\triangleright$  exit while cycle
23:    end if
24:     $\widehat{\mathcal{W}} \leftarrow \widehat{\mathcal{W}} \cup \{k\}$ 
25:  end while
26:  if  $cellCovered = \text{False}$  then
27:     $\widehat{x}_{ij'} \leftarrow 0 \quad \forall i \in \mathcal{R}$   $\triangleright$  Undo allocation
28:    return False
29:  else
30:     $\widehat{F}_h \leftarrow F_h^\dagger, \widehat{F}_m \leftarrow F_m^\dagger, \widehat{F}_{m_1 m_2} \leftarrow F_{m_1 m_2}^\dagger$ 
31:     $\widehat{f}_i \leftarrow f_i^\dagger$   $\triangleright$  Commit updates
32:    return True
33:  end if
34: end function

```

the algorithm directly passes to process the next cell in set $\mathcal{W}^{j'} \setminus \widehat{\mathcal{W}}$. On the contrary, if cell k cannot be covered by anyone of the available sensing resources we exit the main loop (line 22), all $\widehat{x}_{ij'}$ are rounded to zero (line 27), and the function notifies the failure of sensor allocation to Algorithm 1 (line 28). Finally, if all cells in $\mathcal{W}^{j'}$ are successfully covered, the function commits the updates of state variables (line 31), and it notifies Algorithm 1 that the allocation of sensing resources to service j' is successful (line 32).

The worst-case time complexity, say $T(A1)$, of algorithm 1 can be expressed as follows

$$T(A1) = T(P1) + \mathcal{O}(|\mathcal{S}| \times (|\mathcal{H}| + T(A2))), \quad (29)$$

where $T(A2) = (|\mathcal{W}| \times |\mathcal{R}|)$ is the worst-case execution time of Algorithm 2. Specifically, the first term $T(P1)$ is the time required to solve the linear relation of the original

problem (25), called $P1$, which is polynomial in the number of variables and constraints. The second (polynomial) term accounts for the time needed to build the integral solution from the relaxed one. In the worst case, we have to check that the constraints are not violated on every MEC host ($|\mathcal{H}|$) and from all the sensing resources in every cell ($|\mathcal{W}| \times |\mathcal{R}|$ in Algorithm 2). Then, this test should be repeated for each service, i.e. $|\mathcal{S}|$ times. It is worth pointing out that, in practise, the average time complexity of the proposed heuristic, can be much lower as the *cut-off* threshold allows us to reduce the number of service-host and cell-sensor pairs that have to be tested, trading-off accuracy with speed. Furthermore, each service may require to cover only a subset of the available cells, which also reduces the average execution time of the heuristic.

6. Results and Discussion

In this section, we evaluate the performance of the proposed solution. In this study, we resort to simulation as it offers a controllable and repeatable environment. The custom-based simulation environment of the system described in Section 4 is developed using Matlab. The model in 4.2 is also solved using Matlab. We adopt an evaluation setup similar to [14], which is also depicted in Figure 4. More precisely, we consider a reference area of size $400m \times 400m$, which is divided into a regular grid of $K = 400$ cells with side 20 meters. Within this area, $H = 12$ SBSs are also regularly deployed in a grid-like topology. Each SBS has a coverage radius of 120 meters. $N = 1200$ IoT terminals are distributed uniformly at random over the reference area. IoT terminals are associated with the nearest SBS. Without loss of generality, we assume that a MEC host is collocated with each SBS, i.e., $|\mathcal{BS}_m| = 1$. We define two base configurations for the *resource slice* assigned by the 5G-MNO to the IoTSP tenant, which are summarised in Table 2. Specifically, *Scen A* represents a use case with abundant resources. For each MEC host m , we set the memory capacity to $\Gamma_m = 20\text{GB}$, the computation capacity to $\Delta_m = 4\text{GHz}$, the uplink wireless capacity $B_h^\dagger = 60\text{Mbps}$, and the backhaul link capacity $B_{m_1, m_2} = 10\text{Mbps}$. On the other hand, *Scen B* represents a use case with more constrained resources. Basically, we halve the bandwidth and computation resources, and we reduce the memory capacity to $\Gamma_m = 15\text{GB}$. We point out that a backhaul link can be implemented as a virtual link/tunnel in the metro core network of the 5G-MNO. As described in Section 4.1, in this study, we consider a flat network topology for the MEC system. Thus, the virtual mesh network interconnecting the MEC hosts would require a total bandwidth of $H(H-1) \times B_{m_1, m_2}$. Note that in the following evaluation, the values of the base resource configurations are also varied to assess the impact of each constraint on the system performance.

As reference scenario, we focus on data-intensive and latency-sensitive IoT applications that require the sensing, delivery and processing of *large volumes of video and*

	<i>Scen A</i>	<i>Scen B</i>
Γ_m	20 GB	15 GB
Δ_m	4 GHz	2 GHz
B_h^\uparrow	60 Mbps	30 Mbps
B_{m_1, m_2}	10 Mbps	5 Mbps

Table 2: Base configurations for resource slices assigned by the 5G-MNO to the IoTSP tenant.




Figure 4: Evaluation setup

sound data [14, 39]. To this end, we assume that each IoT terminal is equipped with high-quality microphones and cameras. Then, we consider two types of IoT services, namely visual scene recognition (VR), and audio event classification (AC). The first one requires the acquisition of high-resolution video frames that are compared with a database of possibly thousands of images. We assume that the size of a video frame is 2Mb and that end-users request video frames to be processed with a frequency from 0.125 to 1 fps. This design choice corresponds to bandwidth for video frame uploading from 0.25 to 2 Mbps. The computation is set in the range [0.001, 0.08] GHz per cell, assuming 40 CPU cycles per bit, and up to 4 GB of storage are consumed. The audio event classification (AC) service is similar to VR. In this case, acoustic models are compared to sound data collected from the on-board microphones [40]. AC services are typically less resource-demanding than VR services. Specifically, we assume that the size of a sound recording is 1 Mb (half of a high-quality video frame), and that end-users request sound data to be processed with a frequency from 0.125 to 1 Hz. The computation of AC services consumes 30 CPU cycles per bit and require up to 2 GB of storage. Finally, the geographical scope of both AC and VR services is fixed to 10 cells.

6.1. Benchmarks

In the following, we present in detail the three benchmarks, namely, Greedy First-Fit, Greedy Best-Fit, and DSP, we use for performance comparison with our algorithm (called *Opt*) and the approximation algorithm described in Section 5 (called *LR*).

6.1.1. Greedy First-Fit (G-FF)

G-FF iteratively selects an IoT service j at random in the set \mathcal{S} , and for each $k \in \mathcal{W}^j$ greedily searches for a sensing resource that can cover the cell k with the requested frequency of data updates. If all cells in \mathcal{W}^j can be covered without violating the constraints on the wireless bandwidth of SBS, this algorithm tries to place the service on the MEC host that collects data from the largest fraction of cells requested by that service. If this assignment is not feasible (due to a lack of computing, storage or network backhaul resources), service j is simply discarded.

6.1.2. Greedy Best-Fit (DSP)

G-BF algorithm uses the same approach of the G-FF policy to allocate sensing resources to an IoT service j selected at random in the set \mathcal{S} . Differently from G-FF, this algorithm selects the MEC host where to place the service from a sorted list L in a best-fit manner. Specifically, the heuristic adds to a list L all the available MEC hosts and sorts them in descending order of the fraction of cells in \mathcal{W}^j that are covered by each MEC host. Then, the heuristic selects the MEC host from L based on the host rank until computing, storage and wired bandwidth resources are available. If there is not a feasible allocation in the set L , service j is simply discarded.

6.1.3. DSP

The third benchmark is inspired by the optimal operator placement algorithm for data stream processing applications proposed in [19]. Specifically, the problem formulated in [19] aims to place processing elements that carry out a specific operation on a fixed set of data streams on a set of computing nodes in order to minimise overall expected response time, availability, and network usage. There is a clear analogy between a data stream and a sensing resource and between an operator and an IoT service. Specifically, the DSP algorithm sorts all the IoT services in ascending order of the overall data traffic that is requested by the service. Then, sensing resources are statically associated to IoT services using the same rules of the G-BF and G-FF schemes. For each IoT service from the sorted list the DSP algorithm solves the optimisation problem defined in [19] considering only the network usage in the objective function. If there is not a feasible allocation for the IoT service, that service is discarded and the DSP algorithm moves forward to allocate the next service in the ordered list.

6.2. Homogeneous scenario

The first set of results focuses on investigating the efficiency of resource utilisation in a homogeneous scenario where each IoT service selects the cells of interest uniformly at random in the reference area, and end-users select randomly IoT services only from the VR class. Then, we conduct a stress test by evaluating the percentage of deployed IoT services versus the number of service requests,

Service j	Data \hat{b}_l (Mb)	Frequency λ_j (Hz)	Computation (cycles per bit)	Storage γ_j^p (GB)	$ \mathcal{W}^j $
<i>visual scene recognition (VR)</i>	2	[0.125,1]	40	[1,4]	10
<i>audio event classification (AC)</i>	1	[0.125,1]	30	[1,2]	10

Table 3: Resource requirements of IoT services

using the resource slice of *Scen A* in Table 2. The goal of these experiments is to highlight the trends in the utilisation of networking and edge resources when each MEC host collects data from a similar number of activated sensing resources.




Figure 5: Percentage of deployed IoT services versus the number of service requests for uniform cell popularity (*Scen A*).

Figure 5 shows that all policies are able to admit 60 IoT services. However, the performance of G-FF rapidly degrades, and with 100 service requests, the deployment ratio is below 70%. With 100 requests Opt performs slightly better than G-BF, LR and DSP, with a gain up to 4.3%. To explain these results, we take a closer look into the utilisation of resources of individual MEC servers. Figure 6 show the resource utilisation of each resource type using a boxplot for 60 and 100 service requests. Boxplots are a standard method to represent a statistical distribution of values in a compact way. In the graphs, we plot the uplink wireless bandwidth, the total incoming and outgoing traffic load of a MEC host, and the whole storage and computation load (as a percentage). We observe that Opt and the heuristic algorithm are able to ensure a uniform resource utilisation between MEC hosts. At the same time, the greedy-based policies and DSP experience much higher variability in resource utilisation, especially as far as storage and computing power are concerned. Furthermore, for most of the MEC hosts, the benchmarks utilise more storage and computation resources than Opt and LR. Finally, with 100 service requests all the available storage resources are utilised, and this resource bottleneck causes services requests to be discarded. The superior performance of Opt in terms of admitted applications can be explained by the more efficient service placement found by Opt. Interestingly, in the tested configurations networking resources are largely underutilised, and there are no significant differences between the different policies. Clearly, the more service requests and the larger the volume of IoT data. However, this increase is not linear with the number

of service requests thanks to data caching for services with overlapping geographical scopes.



Figure 6: Boxplots of resource usage in percentage. Triangle points are the mean value. The lower and upper whiskers are the 5% and 95% percentile, respectively. Boxes are the outliers.

6.3. Impact of cell popularity

The main lesson we learned from the above experiments is that when the traffic demands are uniformly distributed at random over the network, and the application workloads are homogeneous, a greedy allocation strategy can still achieve near-optimal performance. However, in a real-world urban setting the states to classify or the events to detect are not uniformly distributed in the area of interest, but there are locations more popular than others (a.k.a. point of interests - PoIs). Without loss of generality, we assume that in the reference area A there are two POIs and that the cell popularity around each PoI follows a Zipf distribution with shape parameter α ⁸. The results in Figure 7 are obtained using the resource slice of *Scen B* in Table 2, which ensures to have enough resources to satisfy at least 60 service requests in the homogeneous case (i.e. $\alpha = 0$) using the optimal allocation policy. As expected, the higher the scaling value of the zipdf distribution, the smaller the number of deployed IoT applications, with Opt that goes from a service satisfaction ratio of 100% with $\alpha = 0$ (homogeneous scenario) to 72% with $\alpha = 1.5$. This is due to the fact an increasing number of applications requests data from the most popular cells, and there may not be enough bandwidth capacity to route the IoT

⁸When multiple PoIs exist, the cell popularity distributions generated by each POIs are simply superimposed and then re-normalised.




Figure 7: Percentage of deployed IoT services for different values of zipf α parameter. The results are obtained with two PoIs and 60 IoT requests (*Scen B*).

traffic in the MEC network towards IoT applications that are deployed on farther MEC hosts, where computation and storage resources are still available. However, Opt performs significantly better than the greedy-based solutions and DSP, with gains up to 400% over G-FF, up to 33% over G-BF and up to 20% over DSP. We can observe that the performance gap between Opt and the greedy-based solutions increases with α , while the gap between Opt and DSP remains quite stable. It is important to point out that the optimal policy explores the full solutions space and determine the largest set of applications that can be admitted. On the contrary, the other schemes tries to allocate the IoT services one by one following a fixed order. Specifically, DSP sorts the IoT services based on the amount of data requested by each service, while LR leverages the optimal relaxed solutions of the original problem. The results confirm the effectiveness of our approach and the quality of the solutions computed by the proposed approximation algorithm being the performance gap between Opt and LR lower than 7%. Interestingly, we can observe that DSP performs better than G-BF when increasing the α value, i.e, increasing the popularity of cells around the PoIs. Two reasons can be identified to explain this behaviour. First, when α is large, a service requires data from sources that are spatially concentrated and the backhaul links to/from the MEC hosts that are close to the PoIs rapidly saturate, becoming a resource bottleneck. Thus, DSP ensures a more efficient utilisation of network resources as it formulates the operator placement problem in such a way to minimise the network usage. Second, DSP gives priority to services with low data demands, which is beneficial to admit more services.

Figure 8 shows the utilisation of each resource type for $\alpha = 0.6$ and $\alpha = 1.5$ using a boxplot. Differently from the homogeneous scenario, we can observe that an increase in the popularity concentration, the average usage of computation, storage and bandwidth resources decreases. This is due to the fact that a lower number of IoT services is admitted and fewer resources are needed to execute them. Clearly, this decrease is remarkable with G-FF, as the service satisfaction ratio is very small at $\alpha = 1.5$. Inter-

estingly, although Opt admits the largest number of IoT services, the utilisation of backhaul and radio bandwidth for most MEC hosts is similar to the one observed with G-BF. Finally, DSP shows a higher variability in the usage of storage resources for $\alpha = 0.6$, since this policy gives priority to services that process a low amount of data. To explain the observed behaviours we can note that the optimal data management policy of Opt allows activating the sensing resources that can maximise the caching opportunities for admitted applications. Furthermore, the optimal service placement allows reducing the volume of traffic that needs to be delivered between MEC servers. An important difference with respect to the homogeneous case is that there are MEC hosts that are significantly more utilised than others. These overloaded MEC hosts are typically the ones closer to the PoIs, where most of the sensory streams originate. The benchmarks aim to locate the services as close as possible to the PoIs. On the contrary, Opt allows a more uniform use of the available resources, which not only delays the occurrence of resource bottlenecks but also allows to take full advantage of the available resources.




Figure 8: Boxplots of resource usage in percentage. Square points are the mean value. The lower and upper whiskers are the 5% and 95% percentile, respectively. Boxes are the outliers.

To investigate more in-depth the efficiency of the proposed solutions, Figure 9 depicts the distribution of the leading causes of service reject for different values of the zipf α parameter. More precisely, we consider each rejected service request, and we conduct the following tests: (i) if there is not any MEC host with enough compute (memory) resources to accommodate the new service, we assume that CPU (storage) constraints are the main cause of service reject, respectively; (ii) if the service can be deployed in at least one MEC host, but not all the requested cells can be covered (due to lack of sensing resources or not enough wireless bandwidth to collect the data updates from activates sensing resources) we assume that the fronthaul constraint is the main cause of service reject; finally (iii) if the service can be deployed in at least one MEC host, and all the requested data can be collected from the

activated sensing resources, but there is not enough wired bandwidth to deliver it to the MEC host where the service is deployed, we assume that the backhaul constraint is the main cause of service reject. As shown in Figure 9a and Figure 9b, G-FF is the only allocation policy that rejects service requests due to unavailability of CPU and storage constraints. On the contrary, when cell popularity is highly concentrated around the PoIs, the backhaul limitation is the only cause of service reject for all the other policies. These results confirm our claim that considering the dependencies at the data level when placing services in the MEC system is essential to minimise the flow of data exchanged effectively.




Figure 9: Distribution of the main cause of service reject for different cell popularity values.

We finally highlight the running times of the presented algorithms, summarised in Table II. These running times are based on a Matlab implementation running on an HP ProDesk desktop with 3.6 GHz Intel Core i7700 processor and 16GB RAM. The exact solution of the optimisation problem is computationally demanding as a branch-and-bound method is used. In addition, the time complexity rapidly increases with the α parameter as more backhaul links need to be tested. For instance, the running time when α goes from zero (homogeneous scenario) to 0.6 increases by 480%. When $\alpha = 1.5$, all the iterations of the exact solver reach the imposed time limit (12 hours). As expected, the greedy-based allocation algorithms require significantly lower time to return a solution. However, the proposed heuristic requires only slightly more time than a greedy-based approach to return a close-to-optimal solution and is only up to 4.3 times slower than the G-BF policy in the worst case. The running times of DSP are comparable with those of LR, although the computational efficiency of DSP improves as the α value increases. It is worth pointing that the running times of LR are compatible with an online execution of the allocation policy, while Opt is suitable only for offline allocation.

6.4. Impact of resource constraints

In this section, we explore the impact of each resource on the efficiency of the allocation policies in terms of percentage of deployed services. Due to space limitations, we focus on the case $\alpha = 0.6$ and 60 service requests. In Figure 10a, We first explore in the impact of computation capacity

Algorithm	Running time (seconds)		
	$\alpha = 0$	$\alpha = 0.6$	$\alpha = 1.5$
OPT	722.52 (± 345)	3470 (± 1300)	43200 (± 0)
G-FF	12.92 (± 0.89)	8.15 (± 0.39)	3.64 (± 0.38)
G-BF	13.09 (± 1.9)	8 (± 0.63)	3.3 (± 0.3)
LR	12.77 (± 0.69)	13.3 (± 0.78)	14.15 (± 1.8)
DSP	13.81 (± 2.2)	12.53 (± 0.6)	8.6 (± 0.3)

Table 4: Running times for different algorithms. Table reports mean and variance (in parenthesis) values measured over the 25 runs.

Δ_m considering a range of values starting from 0.5GHz to 2.5GHz. As expected, increasing computation capacity increases the number of deployed service slightly. In addition, the performance gap between the different policies remains stable in all combinations. Next, we show the impact of storage capacity Γ_m in Figure 10b. The observed trends are similar to those depicted in Figure 10a, and again the effect of increasing storage capacity flattens out rapidly. More interesting results are shown in Figure 10c and Figure 10d, which depict the percentage of deployed services for different values of fronthaul (B_h^\uparrow) and backhaul ($B_{m_1 m_2}$) bandwidth capacities, respectively. We can observe that an increase of fronthaul bandwidth capacity from 5Mps to 10 Mbps results in doubling the percentage of deployed services. This is due to the fact that many sensing resources should be activated around the PoIs, and a limited uplink bandwidth can prevent the activation of the needed sensors. However, a further increase of fronthaul bandwidth capacity has only a marginal impact on the efficiency of the proposed algorithms. On the contrary, the percentage of deployed service increases almost linearly with the backhaul bandwidth capacity. These results show that the amount of data exchanged between MEC hosts is non-negligible, and effective management of traffic flows within the MEC network eventually leads to superior performance. Finally, we conclude by noting that the gap between Opt and LR is small in all combinations, while DSP is the second best performing scheme, thus confirming the efficiency of the proposed heuristic.

6.5. Impact of service variability

To conclude our evaluation, it is worth investigating which types of service requests are more efficiently handled by the proposed allocation policies. To this end, Figure 11 depicts the distribution of fulfilled service requests across the two types of services that we have previously defined, i.e. VR and AC, when there are 50 service requests per class and α value is varied from 0 to 1.5. Similarly to what observed in Figure 7, the total number of deployed applications decreases when α increases. However, AC services are less affected by the α value. For instance, when $\alpha = 1.5$ 46% of the VR services is discarded while only 18% of AC services are discarded with Opt. In addition, Opt performs significantly better than the greedy-based solutions, with gains up to 500% over G-FF and up to 54% over G-BF (for $\alpha = 1.5$). These performance gains are higher than




Figure 10: Percentage of deployed applications for different (a) computation, (b) memory, (c) fronthaul bandwidth, and (d) backhaul bandwidth capacities for $\alpha = 1.5$ and 60 service requests.

those observed when only VR services are requested by end-users (see Figure 7). Finally the performance gap between Opt and LR is small, which confirms the efficiency of the heuristic also with heterogeneous service requests.

To explain these results we can observe that AC services consume less resources than VR services. Thus, the objective function (25) assigns a higher utility to AC services over VR services. Clearly, this depends on the adopted revenue model, which assign the same value to each type of service, independently of the resource usage (see equation (17)). In addition, since AC services have a smaller resource footprint it is easier to distribute AC services in the MEC system to fully utilise the spare resources of individual MEC hosts. Opt and LR are more efficient than the greedy-based solutions in exploiting this multiplexing effect. On the other hand, DSP only aims to minimise the network usage and it treats the two services equally.




Figure 11: Distribution of fulfilled service requests across the two types of services for different values of zipf α parameter. The results are obtained with two PoIs and 50 service requests per class (*Scen B*).

7. Related Work

There is a vast literature on resource management in edge computing environments, and MEC systems in particular, and the interested reader is referred to [41, 42] for a

comprehensive survey. Most of the earlier work on this research area focuses on the problem of optimal placement of computation tasks on suitable edge nodes based on the availability of edge resources and the demand patterns under static (e.g. in [10]) or dynamic conditions (e.g. in [11]). Recently, the problem of service placement and request routing in MEC systems where storage resource (i.e. caches) may be shared by requests of the same service from multiple end-users has received significant attention [43]. For instance, the authors in [44] study the problem of computing task caching and formulate a MILP to minimise the task duration of all tasks under the caching capacity constraints. In [45] the joint content caching, computing and bandwidth resources allocation problem is formulated as a MINLP problem to minimise the total energy consumption and network usage, and an ADMM-based distributed algorithm is proposed to solve it. The work in [15] considered a similar problem as [45] under the assumption that the coverage regions of the base stations are non-overlapping. A recent work extended these problem formulations to also account for services that need to pre-store a non-trivial amount of data, or consume significant downlink bandwidth to send the outcome of the processing task to the requesting user [14]. However, these works did not consider that multiple data streams generated by dispersed devices may need to be shared among many services with heterogeneous requirements through data caches. In addition, these studies overlooked the data dependencies when placing services, e.g. in terms of minimising the flow of data exchanged.

Another thread of research related to this work are the studies that extend MEC capabilities to efficiently support IoT deployments and IoT applications [9]. One possible approach is to embed MEC-server functionalities within the conventional IoT gateways to allow local execution of IoT applications [46]. An alternative architecture is proposed in [20], where an additional MEC component, called MEC IoT platform, is introduced to virtualise IoT-gateway functionalities and to enable other MEC applications to interact with deployed IoT devices, hiding the complexity and heterogeneity of underlying IoT physical infrastructures. A network slicing framework for IoT services in hybrid architectures with MEC and centralised cloud systems is proposed in [47], which is based on a series of micro-services encapsulated as Virtual Network Functions (VNFs). SDN and VNF technologies are also

used in [48] to develop a virtualised MEC platform for IoT applications. The authors in [22] propose a new architecture comprising a central IoT Broker, which interacts with IoT applications and shapes IoT traffic to meet QoS requirements at application level, and a 5G Network Slice Broker in charge of negotiating with each IoT Broker network slice resources. A framework that allows uncoordinated access of IoT clients to serverless micro-services in MEC systems is developed in [32]. While our system is inspired by these works, our solution allows IoT tenants to have a better control of resource allocation during IoT service deployment.

There are also many similarities between our scenario and the body of work that addresses the problem of resource allocation for data stream processing (DSP) applications, namely applications that can process unbounded streams of data in a near real-time fashion [49]. Typically, DSP applications can be decomposed into components that execute different processing tasks. Thus, they can be represented as a network of operators that connect data sources with data sinks (i.e., receivers of results). In this context, a classical problem is deciding on how to deploy these components on computing nodes in order to fulfil application requirements and considering computing and communication constraints. The reader is referred to [49, 50] for a comprehensive survey on DSP frameworks, while in this overview we are concerned with solutions that focus on edge/fog environments. For instance, the authors in [51] formulated an *assignment problem* as a MILP to map DSP operators onto available resources in a mixed cloud-edge infrastructure, while jointly minimising task response time and deployment costs. The proposed model also allows to parallelise the execution of operators onto multiple computing resources. To address the scalability issue of the model in [51], in a follow-up papers the authors propose a resource selection technique to reduce the state space of the problem [52]. Two greedy strategies are described in [53] to solve the operator assignment problem, which prioritise the allocation of resources close to the data sources. A MILP formulation of the operator assignment problem is also proposed in [19] aiming to jointly minimise the expected application response and unavailability, as well as network usage. Then, a set of model-free and model-based heuristics are proposed to find a solution. Recently, the authors of [54] have formulated the operator placement problem for DSP applications in IoT environments. Our work differs from these solutions because we consider the possibility of sharing data flows among processing applications using edge caches. Furthermore, most of the problem formulations assumes that the mapping between data sources and operators is fixed, while we allow a dynamic assignment to optimise the routing of data flows within the network of edge devices.

8. Conclusions

In this paper, we have studied the problem of joint service placement and data management in mobile edge comput-

ing systems. We have targeted a scenario where an IoT service provider leverages network, computing and storage resources of a MEC-enabled system to fulfil the heterogeneous requirements of delay-sensitive and data-intensive IoT applications. The proposed approach aims at the maximisation of the revenue due to the application deployment process (in terms of number of deployed applications) while minimising the cost of the required physical resources. Furthermore, our model accounts for data dependencies when placing services, e.g. in terms of minimising the flow of data exchanged between edge servers. We have developed both an optimisation framework and a heuristic algorithm. Finally, we have proposed a MEC architectural solution to enable the IoT service provider to transparently interact with the MEC system. The results obtained with realistic scenarios show the efficiency of the proposed approach, especially when traffic demands generated by the service requests are not uniform. Interesting directions for future work include exploring the cooperation and competition between multiple IoTSPs using the same MEC systems, as well as investigating general MEC topology and different data models.

References

- [1] R. Du, P. Santi, M. Xiao, A. V. Vasilakos, C. Fischione, The Sensable City: A Survey on the Deployment and Management for Smart City Monitoring, *IEEE Communications Surveys Tutorials* 21 (2) (2019) 1533–1560. doi:10.1109/COMST.2018.2881008.
- [2] V. Farhadi, F. Mehmeti, T. He, T. L. Porta, H. Khamfroush, S. Wang, K. S. Chan, Service placement and request scheduling for data-intensive applications in edge clouds, in: *Proc. of IEEE INFOCOM 2019*, 2019, pp. 1279–1287.
- [3] C. M. De Farias, W. Li, F. C. Delicato, L. Pirmez, A. Y. Zomaya, P. F. Pires, J. N. De Souza, A Systematic Review of Shared Sensor Networks, *ACM Computing Surveys* 48 (4) (2016) 51:1–51:45.
- [4] P. P. Ray, A survey of IoT cloud platforms, *Future Computing and Informatics Journal* 1 (1) (2016) 35–46.
- [5] P. Hu, S. Dhelim, H. Ning, T. Qiu, Survey on fog computing: architecture, key technologies, applications and open issues, *Journal of Network and Computer Applications* 98 (2017) 27–42.
- [6] Y. Ai, M. Peng, K. Zhang, Edge computing technologies for Internet of Things: a primer, *Digital Communications and Networks* 4 (2) (2018) 77–86.
- [7] N. Hassan, S. Gillani, E. Ahmed, I. Yaqoob, M. Imran, The role of edge computing in internet of things, *IEEE Communications Magazine* 56 (11) (2018) 110–115.
- [8] ETSI ISG MEC, Multi-access Edge Computing (MEC); Framework and Reference Architecture, ETSI GS MEC 003 V2.1.1 (January 2019).
- [9] P. Porambage, J. Okwuibe, M. Liyanage, M. Ylianttila, T. Taleb, Survey on Multi-Access Edge Computing for Internet of Things Realization, *IEEE Communications Surveys & Tutorials* 20 (4) (2018) 2961–2991. doi:10.1109/COMST.2018.2849509.
- [10] X. Chen, L. Jiao, W. Li, X. Fu, Efficient Multi-User Computation Offloading for Mobile-Edge Cloud Computing, *IEEE/ACM Transactions on Networking* 24 (5) (2016) 2795–2808.
- [11] H. Guo, J. Liu, J. Zhang, W. Sun, N. Kato, Mobile-Edge Computation Offloading for Ultradense IoT Networks, *IEEE Internet of Things Journal* 5 (6) (2018) 4977–4988.

- [12] M. Alkhalaileh, R. N. Calheiros, Q. V. Nguyen, B. Javadi, Data-intensive application scheduling on Mobile Edge Cloud Computing, *Journal of Network and Computer Applications* 167 (2020) 102735.
- [13] F. A. Salaht, F. Desprez, A. Lebre, An overview of service placement problem in fog and edge computing, *ACM Computing Surveys* 53 (3). doi:10.1145/3391196.
- [14] K. Poularakis, J. Llorca, A. M. Tulino, I. Taylor, L. Tassiulas, Service placement and request routing in mec networks with storage, computation, and communication constraints, *IEEE/ACM Transactions on Networking* 28 (3) (2020) 1047–1060.
- [15] T. He, H. Khamfroush, S. Wang, T. La Porta, S. Stein, It's Hard to Share: Joint Service Placement and Request Scheduling in Edge Clouds with Sharable and Non-Sharable Resources, in: *Proc. of IEEE ICDCS'18*, 2018, pp. 365–375.
- [16] G. Castellano, A. Manzalini, F. Risso, A Disaggregated MEC Architecture Enabling Open Services and Novel Business Models, in: *Proc. of IEEE NetSoft'19*, 2019, pp. 178–182.
- [17] Jonas Strobel, Edge computing: Five viable telco business models, *STL Partners – Executive Briefing* (November 2017).
- [18] L. C. Lau, R. Ravi, M. Singh, *Iterative Methods in Combinatorial Optimization*, Cambridge Texts in Applied Mathematics, Cambridge University Press, 2011.
- [19] M. Nardelli, V. Cardellini, V. Grassi, F. L. Presti, Efficient Operator Placement for Distributed Data Stream Processing Applications, *IEEE Transactions on Parallel and Distributed Systems* 30 (8) (2019) 1753–1767. doi:10.1109/TPDS.2019.2896115.
- [20] L. Zanzi, F. Cirillo, V. Sciancalepore, F. Giust, X. Costa-Perez, S. Mangiante, G. Klas, Evolving Multi-Access Edge Computing to Support Enhanced IoT Deployments, *IEEE Communications Standards Magazine* 3 (2) (2019) 26–34.
- [21] Sami Kekki et al., MEC in 5G networks, *ETSI White Paper N. 28* (June 2018).
- [22] V. Sciancalepore and F. Cirillo and X. Costa-Perez, Slice as a Service (SlaaS) Optimal IoT Slice Resources Orchestration, in: *Proc. of IEEE GLOBECOM'17*, 2017, pp. 1–7. doi:10.1109/GLOCOM.2017.8254529.
- [23] A. Zanella, N. Bui, A. Castellani, L. Vangelista, M. Zorzi, Internet of things for smart cities, *IEEE Internet of Things Journal* 1 (1) (2014) 22–32.
- [24] M. Cristani, M. Bicego, V. Murino, Audio-Visual Event Recognition in Surveillance Video Sequences, *IEEE Transactions on Multimedia* 9 (2) (2007) 257–267.
- [25] K. Ahmad, N. Conci, How Deep Features Have Improved Event Recognition in Multimedia: A Survey, *ACM Transactions on Multimedia Computing, Communications, and Applications* 15 (2).
- [26] D. Sabella, A. Vaillant, P. Kuure, U. Rauschenbach, F. Giust, Mobile-Edge Computing Architecture: The role of MEC in the Internet of Things, *IEEE Consumer Electronics Magazine* 5 (4) (2016) 84–91.
- [27] Q. Pham, F. Fang, V. N. Ha, M. J. Piran, M. Le, L. B. Le, W. Hwang, Z. Ding, A survey of multi-access edge computing in 5g and beyond: Fundamentals, technology integration, and state-of-the-art, *IEEE Access* 8 (2020) 116974–117017. doi:10.1109/ACCESS.2020.3001277.
- [28] ETSI ISG MEC, Mobile Edge Computing (MEC); Mobile Edge Management; Part 1: System, host and platform management, *ETSI GS MEC 010-1 V1.1.1* (October 2017).
- [29] ETSI ISG MEC, Mobile Edge Computing (MEC); Mobile Edge Management; Part 2: Application lifecycle, rules and requirements management, *ETSI GS MEC 010-2 V1.1.1* (July 2017).
- [30] ETSI ISG MEC, Mobile Edge Computing (MEC); Location API, *ETSI GS MEC 013 V1.1.1* (July 2017).
- [31] Y. Sun, G. Feng, S. Qin, S. Sun, Cell Association With User Behavior Awareness in Heterogeneous Cellular Networks, *IEEE Transactions on Vehicular Technology* 67 (5) (2018) 4589–4601.
- [32] C. Cicconetti, M. Conti, A. Passarella, Uncoordinated access to serverless computing in MEC systems for IoT, *Computer Networks* 172 (2020) 107184.
- [33] Amazon web services: Auto scaling., <https://aws.amazon.com/autoscaling/>.
- [34] Q. Zhu, M. Y. Sarwar Uddin, N. Venkatasubramanian, C. Hsu, Spatiotemporal Scheduling for Crowd Augmented Urban Sensing, in: *Proc. of IEEE INFOCOM'18*, 2018, pp. 1997–2005. doi:10.1109/INFOCOM.2018.8485869.
- [35] X. Ge, S. Tu, G. Mao, C. Wang, T. Han, 5g ultra-dense cellular networks, *IEEE Wireless Communications* 23 (1) (2016) 72–79.
- [36] M. Kamel, W. Hamouda, A. Youssef, Ultra-Dense Networks: A Survey, *IEEE Communications Surveys & Tutorials* 18 (4) (2016) 2522–2545.
- [37] F. Glover, E. Woolsey, Converting the 0-1 Polynomial Programming Problem to a 0-1 Linear Program, *Operations Research* 22 (1) (1974) 180–182.
- [38] N. V. T. R. Horst, Panos M. Pardalos, *Introduction to Global Optimization*, 2nd Edition Edition, Springer, 2000.
- [39] C. Delgado, M. Canales, J. Ortín, J. R. Gállego, A. Redondi, S. Bousnina, M. Cesana, Joint Application Admission Control and Network Slicing in Virtual Sensor Networks, *IEEE Internet of Things Journal* 5 (1) (2018) 28–43. doi:10.1109/JIOT.2017.2769446.
- [40] J. Park, S. Kim, Sound learning-based event detection for acoustic surveillance sensors, *Multimedia Tools and Applications* 79 (2020) 16127–16139.
- [41] Y. Mao, C. You, J. Zhang, K. B. Letaief, A Survey on Mobile Edge Computing: The Communication Perspective, *IEEE Communications Surveys & Tutorials* 19 (4) (2017) 2322–2358.
- [42] C.-H. Hong, B. Varghese, Resource Management in Fog/Edge Computing: A Survey on Architectures, Infrastructure, and Algorithms, *ACM Computing Surveys* 52 (5).
- [43] S. Wang, X. Zhang, Y. Zhang, L. Wang, J. Yang, W. Wang, A Survey on Mobile Edge Networks: Convergence of Computing, Caching and Communications, *IEEE Access* 5 (2017) 6757–6779.
- [44] M. Chen, Y. Hao, L. Hu, M. S. Hossain, A. Ghoneim, Edge-CoCaCo: Toward Joint Optimization of Computation, Caching, and Communication on Edge Cloud, *IEEE Wireless Communications* 25 (3) (2018) 21–27.
- [45] Q. Chen, F. R. Yu, T. Huang, R. Xie, J. Liu, Y. Liu, Joint Resource Allocation for Software-Defined Networking, Caching, and Computing, *IEEE/ACM Transactions on Networking* 26 (1) (2018) 274–287.
- [46] C. Vallati, A. Viridis, E. Mingozzi, G. Stea, Mobile-Edge Computing Come Home Connecting things in future smart homes using LTE device-to-device communications, *IEEE Consumer Electronics Magazine* 5 (4) (2016) 77–83.
- [47] R. Sanchez-Iborra, S. Covaci, J. Santa, J. Sanchez-Gomez, J. Gallego-Madrid, A. F. Skarmeta, MEC-Assisted End-to-End 5G-Slicing for IoT, in: *Proc. of IEEE GLOBECOM'19*, 2019, pp. 1–6.
- [48] H.-C. Hsieh, J.-L. Chen, A. Benslimane, 5G Virtualized Multi-access Edge Computing Platform for IoT Applications, *Journal of Network and Computer Applications* 115 (2018) 94–102.
- [49] H. Isah, T. Abughofa, S. Mahfuz, D. Ajerla, F. Zulkernine, S. Khan, A Survey of Distributed Data Stream Processing Frameworks, *IEEE Access* 7 (2019) 154300–154316.
- [50] M. Dias de Assunção, A. da Silva Veith, R. Buyya, Distributed data stream processing and edge computing: A survey on resource elasticity and future directions, *Journal of Network and Computer Applications* 103 (2018) 1–17.
- [51] F. Rodrigo De Souza, M. Dias de Assuncao, E. Caron, A. da Silva Veith, An Optimal Model for Optimizing the Placement and Parallelism of Data Stream Processing Applications on Cloud-Edge Computing, in: *Proc. of IEEE SBAC-PAD'20*, Porto, Portugal, 2020, pp. 59–66.
- [52] F. Rodrigo De Souza, A. Da Silva Veith, M. Dias de Assuncao, E. Caron, Scalable Joint Optimization of Placement and Parallelism of Data Stream Processing Applications on Cloud-Edge Infrastructure, in: *Proc. of IEEE ICSOC'20*, Vol. 12571

of Lecture Notes in Computer Science, Springer, Cham, 2020, pp. 149–164.

- [53] A. C. Franco da Silva, P. Hirmer, B. Mitschang, Model-Based Operator Placement for Data Processing in IoT Environments, in: Proc. of IEEE SMARTCOMP'19, 2019, pp. 439–443.
- [54] Y. Zhang, J. Yan, L. Pu, S. Chen, Dynamic component placement and request scheduling for iot big data streaming, IEEE Internet of Things Journal 7 (8) (2020) 7156–7170. doi:10.1109/JIOT.2020.2982458.