**QUEEN'S
UNIVERSITY
BELFAST**

DYVERSE: DYnamic VERtical scaling in multi-tenant Edge environments

Wang, N., Matthaiou, M., Nikolopoulos, D. S., & Varghese, B. (2020). DYVERSE: DYnamic VERtical scaling in multi-tenant Edge environments. *Future Generation Computing Systems*, 108, 598-612.
<https://doi.org/10.1016/j.future.2020.02.043>

Published in:
Future Generation Computing Systems

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

Copyright 2020 Elsevier.

This manuscript is distributed under a Creative Commons Attribution-NonCommercial-NoDerivs License (<https://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits distribution and reproduction for non-commercial purposes, provided the author and source are cited.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

DYVERSE: DYnamic VERTical Scaling in Multi-tenant Edge Environments

Nan Wang¹, Michail Matthaiou¹, Dimitrios S. Nikolopoulos², and Blesson Varghese^{1,*}

¹*Queen's University Belfast, UK*

²*Virginia Polytechnic Institute and State University, USA*

Abstract

Multi-tenancy in resource-constrained environments is a key challenge in Edge computing. In this paper, we develop 'DYVERSE: DYnamic VERTical Scaling in Edge' environments, which is the first light-weight and dynamic vertical scaling mechanism for managing resources allocated to applications for facilitating multi-tenancy in Edge environments. To enable dynamic vertical scaling, one static and three dynamic priority management approaches that are workload-aware, community-aware and system-aware, respectively are proposed. This research advocates that dynamic vertical scaling and priority management approaches reduce Service Level Objective (SLO) violation rates. An online-game and a face detection workload in a Cloud-Edge test-bed are used to validate the research. The merit of DYVERSE is that there is only a sub-second overhead per Edge server when 32 Edge servers are deployed on a single Edge node. When compared to executing applications on the Edge servers without dynamic vertical scaling, static priorities and dynamic priorities reduce SLO violation rates of requests by up to 4% and 12% for the online game, respectively, and in both cases 6% for the face detection workload. Moreover, for both workloads, the system-aware dynamic vertical scaling method effectively reduces the latency of non-violated requests, when compared to other methods.

© 2019 Published by Elsevier Ltd.

Keywords: Edge computing, multi-tenancy, vertical scaling, dynamic priority, dynamic scaling

1. Introduction

The vision of next-generation distributed computing is to harness the network edge for computing [1, 2]. In contrast to servicing all user requests from the Cloud, a workload may be distributed across the Cloud and nodes, such as routers and switches or micro data centres, that are located at the edge of the network [3, 4, 5].

Figure 1 shows a three-tier Edge computing architecture. The Cloud tier is represented by data centres that provide compute resources for workloads. The Edge tier uses nodes that are closer to users. These include: (i) traffic routing nodes – existing nodes that route Internet traffic,

*Corresponding author; E-mail: b.varghese@qub.ac.uk

Figure 1. Example of a three-tier Edge computing architecture.

for example, Wi-Fi routers, which may be augmented with additional compute resources, and (ii) dedicated nodes – additional micro data centres, for example, cloudlets. Workloads hosted on the Edge could either be the same or a subset of functionalities of those hosted on the Cloud based on the availability of resources at the Edge. The end device tier represents user devices and sensors; 29 billion of these are estimated to be connected to the Internet by 2022¹.

In Edge computing, end devices are connected to Edge nodes instead of directly to servers in the Cloud. The benefits of distributing a workload across the Cloud and the Edge have already been established in the literature. They include reduced communication latencies and reduced traffic to the Cloud, which in turn improves response times and Quality-of-Service (QoS) [6, 7].

There are challenges in achieving the vision of using Edge computing for distributing Cloud workloads, especially when Edge nodes are resource constrained as in traffic routing nodes. This paper aims to investigate the problem of supporting *multi-tenancy* in Edge computing nodes with limited hardware resources when compared to the Cloud. In this case, the servers of multiple workloads hosted on the same Edge node are anticipated to compete for insufficient resources [8]. Similar to Cloud computing, multiple workloads are expected to share hardware resources in Edge computing [9] since it would not be cost-effective to customise individual Edge computing nodes for a specific application [10]. Although when multiple workloads are asking for computing resources from a resource-constrained Edge node, a few of the workloads may continue using Cloud servers. This would minimise the opportunity to leverage the edge of the network for reducing the distance

¹<https://www.ericsson.com/en/mobility-report/internet-of-things-forecast>

of data transfer and subsequently improving the QoS of applications [11, 12]. Therefore, it is important to efficiently support multi-tenancy on the Edge.

Multi-tenancy on the Cloud is well researched to host multiple Virtual Machines (VMs) on the same underlying hardware [13, 14]. Nonetheless, it needs to be revisited in the context of Edge computing since: (i) Edge resources have limited processing capabilities, due to small form factor and low-power processors when compared to data centre resources [15]. (ii) The Edge is a more transient environment (availability of resources change over time and may be available only for short periods) compared to the Cloud [16].

Multi-tenancy causes resources contention. Mechanisms employed on the Cloud to mitigate resource contention include *vertical scaling*, which is a process of allocating/deallocating resources to/from workloads so that multiple workloads can coexist [17]. These solutions are not suited for Edge environments since: (i) The mechanisms to monitor and optimise the allocation/deallocation of resources to meet user-defined objectives, specified as *Service Level Objectives (SLOs)* are typically computationally intensive [18, 19]. (ii) Predictive models used for estimating resource demands will have insufficient data for training [20, 21]. Edge services are expected to have short life cycles and may result in insufficient data to feed into an accurate Machine-Learning (ML) model; (iii) Edge environments are expected to have more transient system states compared to the Cloud and Edge workloads may run in a come-and-go style. Solutions designed for the Cloud platforms do not consider this. Therefore, a light-weight vertical scaling mechanism to facilitate multi-tenancy on the Edge is required, which is proposed and presented in this article.

As the execution of a workload progresses in a multi-tenant environment, there may be workloads that require more or fewer resources to meet their SLOs. If required resources are not available for Edge workloads, then SLOs will be violated. A static resource provisioning method is unsuitable given the frequent changes in an Edge environment and resource scaling is therefore required. The research in this paper proposes a dynamic vertical scaling mechanism that can be employed in a multi-tenant and resource-limited Edge environment.

The proposed mechanism is underpinned by a model that accounts for *static priorities* (set before the execution of a workload) and *dynamic priorities* (that changes during execution) of workloads on the Edge. While priorities have been exploited in Cloud computing [22, 23], we investigate it in the context of Edge computing in this article. The Edge is expected to be a premium service for Cloud workloads, and therefore, selecting Edge service users becomes important. We propose three

dynamic priorities that are workload-aware, community-aware and system-aware to this end. We hypothesise that dynamic vertical scaling along with priority management approaches will improve the QoS of multi-tenant workloads and reduce SLO violation rates.

This paper makes the following novel contributions:

- i. The development of a framework for supporting multi-tenancy in a resource-constrained Edge computing node. In this framework, the Edge service QoS maximisation problem in a three-tier environment is formulated by considering the SLO violation rate, server co-location, dynamic vertical scaling, and priorities of workloads.
- ii. The design of the dynamic priority approaches for managing Edge applications in a multi-tenant environment, which accounts for Edge-specific characteristics and different economic models. Currently, research using priority management (for example the decision for offloading or task queuing) focuses on the pre-deployment phase. DYVERSE, on the other hand, applies priority management after deployment.
- iii. The development of the lightweight dynamic vertical scaling mechanism that adjusts resource allocations for prioritised Edge applications after deployment. Existing resource management techniques are heavyweight (require significant processing) since they are computationally intensive [24]. DYVERSE offers resource management suited for resource-constrained Edge nodes.
- iv. The evaluation of DYVERSE on two different use-cases in a realistic experimental environment. Much of existing research on resource management for Edge computing is evaluated using simulators [25, 26, 27]. The benefits of applying DYVERSE in an Edge node is on the contrary presented in a test-bed with a location-based mobile game and a real-time face detection application.

The feasibility of the proposed priority management approaches and dynamic vertical scaling mechanism is validated using an online-game and a face detection workload in an Edge environment. These workloads are a natural fit for using the Edge since they are latency critical - the response time is affected by the distance between the user device and the server. The merit of DYVERSE is observed in that they only have a sub-second overhead per Edge server when 32 servers are deployed on an Edge node. Additionally, we observe that scaling using static and dynamic priorities reduces

the SLO violation rates of user requests by up to 4% and 12% for the online game (6% for the face detection workload) respectively, when compared to executing workloads on the Edge node without dynamic vertical scaling. Moreover, the proposed dynamic vertical scaling with the system-aware dynamic priority approach improves the latencies of requests that are not violated. The key result is that our initial hypothesis is confirmed.

The remainder of this paper is organised as follows. Section 2 presents the background and develops the problem model. Section 3 proposes one static and three dynamic priority management approaches for a multi-tenant Edge environment. Section 4 presents a dynamic vertical scaling mechanism for Edge nodes. Section 5 experimentally evaluates DYVERSE against a catalogue of metrics - system overhead, SLO violation rate and latency. Section 6 highlights the related work. Section 7 concludes this paper.

2. Background and Problem Model

The architecture considered in this paper is based on a three-tier model (as shown in Figure 1) in which compute resources are located at the edge of the network closer to end devices [28]. In the Cloud tier, workloads are hosted on servers in a data centre. To enable the use of the Cloud in conjunction with the Edge nodes, we deploy a Cloud Manager along with each server. This manager is responsible for offloading a workload onto an Edge node. The Edge tier comprises an Edge node with an Edge Manager communicating with the Cloud Manager and managing the Edge resources. On-demand Edge servers are deployed to service requests from end devices. In the end device tier, multiple devices are connected to the Cloud servers. When an Edge server is deployed, the end devices connect to the Edge servers.

In this paper, we consider an Edge node to be a multi-tenant environment hosting multiple workloads. It would not be feasible to build bespoke Edge systems for individual workloads to improve their QoS. Therefore, workloads will need to share an Edge node. Unlike Cloud computing, where additional resources can be purchased when needed, Edge environments are resource limited and hence multi-tenancy is challenging.

To support multi-tenancy on an Edge node and given that a suitable economic model is available for utilising the Edge, workloads need to be prioritised for using resources. If they are not prioritised, they have to equally share resources; this is inefficient since resource demands may vary and will be affected by the workloads executed. Prioritising workloads makes it possible to assign

Table 1. Notation used in the Cloud-Edge model

Notation	Description	Source
S	A set of Edge servers s , each hosted in a container on an Edge node.	Edge Manager
PS_s	Priority score of Edge server s ; it is calculated statically SPS_s or dynamically $(w/c/s)DPS_s$	
uR	One unit of CPU and memory	
U_s	A set of users u to connect to s	
L_s	The service level objective (in terms of latency) of Edge server s	Cloud Manager
$donation_s$	flag for the willingness of donating resources from Edge server s	
$dThr_s$	Threshold of L_s to allow the scaling down of Edge server s	
aL_s	Average latency of Edge server s	
R_s	CPU and memory used by Edge server s on the Edge node	
VR_s	Service level objective violation rate of Edge server s in the previous round of dynamic vertical scaling	Monitor
aR_s	Resources to be added to Edge server s for scaling up	
FR	Free CPU and memory on the Edge node available for S	
$decision$	Flag for 'scaleup' or 'scaledown' containers on an Edge node	Auto-scaler

resources based on the demand of each workload and reduces any under-utilisation of resources. In this research, a priority management mechanism for an Edge node that (i) maintains a Priority Score (PS) for each running workload, and (ii) determines which workload should be first scaled is proposed.

Resource provisioning has been considered in Edge environments [9]. However, existing research focuses on managing resources during the scheduling and deployment of workloads, referred to as pre-deployment management. In a multi-tenant Edge environment, where resources are limited and the availability of free resources keeps changing, it is important to efficiently manage resources periodically to not only avoid overloading the Edge node but also ensuring that the overall QoS of Edge workloads is satisfactory. Post-deployment resource management refers to managing resources after workloads are deployed [29]. In this paper, we consider post-deployment management along with dynamic priorities to address the resource allocation problem for the multi-tenant Edge

130 environment.

131 2.1. Problem Model

The notation to represent an Edge environment is shown in Table 1. Let $S = \{s_1, s_2, \dots, s_n\}$ be a set of latency-sensitive servers hosted on an Edge node e . Each server $s \in S$ is a subset of functionalities of the Cloud server that is offloaded onto the Edge node. When deploying s onto the Edge node, its Cloud Manager provides U_s , a set of users that are to be serviced by s , a latency objective L_s , its willingness to donate resources $donation_s$, and the threshold for permitting a scale-down action $dThr_s$. Then the multi-tenant Edge node is modelled as

$$s_n = \langle U_{s_n}, L_{s_n}, donation_{s_n}, dThr_{s_n} \rangle, \forall s_n \in S \quad (1)$$

132 Our research employs four components on an Edge node to facilitate multi-tenancy: (i) *Edge*
133 *Manager* maintains the Edge server registry, makes decisions on starting and terminates Edge
134 servers; (ii) *Monitor* periodically monitors a number of metrics related to each Edge server s .
135 These metrics include: (a) CPU and memory usage R_s , (b) average latency aL_s for all users
136 $u \in U$, (c) workload intensity, for example, number of requests $Request_s$, and (d) scaling frequency.
137 (iii) *Auto-scaler* dynamically makes decisions to scale up or down an Edge server s based on its
138 performance; and (iv) *Edge Server* s interacts with end devices.

The research objective is to: *develop the mechanisms required for sharing computing resources of an Edge node among multiple workloads while minimising the latency of each workload and maximising its performance. The overall performance of workloads on the Edge node e is measured by the SLO violation rate of all servers S :*

$$VR_e = \frac{\sum_{s \in S} Request_s[aL_s > L_s]}{\sum_{s \in S} Request_s} \quad (2)$$

139 The equation above may be affected by varying the resources allocated to an Edge server and
140 the sequence of allocating/deallocating resources on Edge servers. We hypothesise that dynamic
141 resource allocation along with priorities will reduce the SLO violation rate.

142 3. Priority Management

143 Uniform allocation of resources to multiple tenants on an Edge node can occur at the same time.
144 However, this is a static allocation technique and does not consider the specific needs of individual

tenants. Customised allocations cannot always proceed concurrently since Edge nodes are resource constrained. Therefore, the sequence of allocating resources for running servers on an Edge node needs to be considered. In this paper, it is assumed that Edge and Cloud service providers are different. Hence, priority management, scaling schemes, and resource management in the Cloud and Edge are decoupled. Pricing models also affect the priority of an Edge server. For instance, it may be unfair to assign the same priority to two tenants with similar computational requirements – when one tenant pays for a fixed period and the other pays for the resource. In this case, a different priority needs to be assigned to the tenants.

In this section, we propose Dynamic Priority Management (DPM) approaches when using different pricing models. DPM is compared against a Static Priority Management (SPM) approach, which is used as a baseline. We envision Edge computing to be a utility-based service relying on the *pay-for-X* (pay for what you use) principle. The following three pricing models are considered: (i) Pay-For-Resources (PFR) – the resource used is paid for, (ii) Pay-For-Period (PFP) – the time for utilising resources is paid for, and (iii) Hybrid [30] – combines both PFR and PFP models.

3.1. Static Priority Management

We define SPM as a baseline in the priority management of multiple tenants. PS is the importance of an Edge server (high value means high priority and the server is provided resources before others), which is calculated when the server is launched. The PS remains the same from deployment until termination of the server. This is comparable to a flat rate model in resource pricing. SPM is a realistic approach in contrast to an approach using a pre-defined PS provided by the Cloud Manager since the manager would give the highest PS to its Edge server. SPM, on the other hand, is affected by several factors that are measured by the Edge Manager and can differentiate the servers on an Edge node.

Table 2 lists four factors affecting the Static Priority Score (SPS) of an Edge server. The Cloud Manager of each workload is allowed to purchase the premium service to gain a higher PS. This factor is measured by the price a Cloud Manager has paid for premium service P_s . Such purchased priority has been adopted in designing scheduling algorithms in Cloud computing [31], which aims to guarantee that users who pay more can experience higher quality service. The Edge Manager maintains a record of the sequence in which requests were made for using Edge service; each Edge server is marked with an ordinal number ID_s when the Cloud Manager requests an

Table 2. Factors affecting static priority management

Factor	Notation	Description	Source
Premium Service	P_s	Price paid for purchasing priority on Edge server s	Cloud Manager
First Come First Serve	ID_s	The ID of Edge server s , set by the sequence of the request for launching s	Edge Manager
Ageing	Age_s	No. of times Edge server s has been rejected by an Edge node	
Loyalty	$Loyalty_s$	No. of times Edge server s has used services on an Edge node	

Edge service. This is in line with the first-come-first-serve policy widely adopted in Cloud resource provisioning [32]. However, using only P_s and ID_s may result in priorities becoming fixed and may deprive a workload of resources if it did not initially purchase a premium service. To avoid this, the Edge Manager also considers ageing Age_s , which is the number of times an Edge server has been denied resource access. Aging is implemented in many scheduling algorithms for Cloud computing as a technique to ensure that all jobs are eventually scheduled [33]. The last factor considered by the Edge Manager is loyalty $Loyalty_s$, which is the number of times the Edge service was used. Loyalty is used as an incentive for Cloud service users to adopt the Edge computing model [34]. SPS for an Edge server s is calculated as:

$$SPS_s = W_P * P_s + W_{ID} * \frac{1}{ID_s} + W_{Age} * Age_s + W_{Loyalty} * Loyalty_s \quad (3)$$

where W is the weight assigned to a factor. In this paper, SPS_s is used as a baseline combining the factors that are considered to be important in the literature as indicated above on resource management. This makes it a better baseline than using a single factor to decide the priority of an Edge server.

SPM is advantageous in that in SPM the priority score of an Edge application is calculated only once when it is deployed. This solution entails low overhead and is straightforward to implement. It is also proved that static factors are useful for designing effective resource management solutions in the context of Cloud computing [31, 32, 33, 34]. However, a disadvantage of SPM is that the PS of an Edge server may not remain the same in the real world. Dynamic factors affecting the

Table 3. Factors affecting dynamic priority management

Factor	Notation	Description	Source
	$Request_s$	No. of requests serviced by Edge server s	
Workload	$ U_s $	No. of users serviced by Edge server s	
	$Data_s$	Amount of data processed by Edge server s	Monitor
Reward	$Reward_s$	No. of times Edge server s has donated resources	
Penalty	$Scale_s$	No. of times Edge server s has been scaled on an Edge node	

performance of Edge servers, for example, whether a resource-intensive workload is executed needs to be considered. Hence, dynamic priorities are proposed in this paper, the aim is to further improve the effectiveness of priority management required for Edge computing.

3.2. Dynamic Priority Management

Table 3 lists factors affecting the Dynamic Priority Score (DPS) of an Edge server. Since vertical scaling occurs periodically (presented in Section 4), the values of these factors are updated in every round of vertical scaling. This ensures that varying priorities are taken into account for scaling decisions. In DPM, three approaches are introduced. The first is a workload-aware approach in which the workload of an Edge server in the previous round of scaling will affect its priority in the next round. The Monitor records three workload related metrics of each Edge server s , which includes the number of requests $Request_s$ and users $|U_s|$ serviced, and the amount of data transferred between the user and the Edge server $Data_s$.

The second DPM approach is community-aware in which a workload can donate resources to a shared resource pool. If an Edge server s donates, then it is recorded by the Monitor in $Reward_s$. In the subsequent scaling round, the server that donated will be rewarded a higher PS.

The third DPM approach is system-aware in which the adverse effect of continuously scaling on the Edge is mitigated. Frequent scaling may result in an unstable Edge node that is continuously aiming to meet the requirements of workloads, which cumulatively would result in large overheads. To avoid this the Monitor keeps a record of the number of times an Edge server s has been scaled $Scale_s$. This is used to penalise the workload when calculating its PS in the next scaling round.

3.2.1. Workload-aware Dynamic Priority Score (*wDPS*)

Workload-related factors presented in Table 3 differentiate individual workloads on the Edge node. For different pricing models, the workload-aware DPS needs to be different to ensure fairness among Edge service users. For example, in the PFP pricing model, the priority of an Edge server should be reduced if it will require more resources for a given time. However, in the PFR and Hybrid pricing models, it is assumed that additional use of resources is already paid for. Therefore, workload-related factors are not considered for issuing penalties in these models.

In the PFR and Hybrid pricing models, *wDPS* of an Edge server s is defined as:

$$wDPS_{s,PFR/Hybrid} = SPs_s + W_{Request} * Request_s + W_U * |U_s| + W_{Data} * Data_s \quad (4)$$

With the above Equation, *wDPS* may deprive small workloads by assigning them lower priority scores and assigning large workloads higher priority scores. This could result in the PFR and Hybrid pricing models being an unfair system for applications with varying workloads. Hence, to address fairness, Equation 5 defines *wDPS* in the PFP pricing model as:

$$wDPS_{s,PFP} = SPs_s + \frac{1}{W_{Request} * Request_s} + \frac{1}{W_U * |U_s|} + \frac{1}{W_{Data} * Data_s} \quad (5)$$

The above equation addresses fairness by assigning lower priority scores to large workloads. Through implementing *wDPS*, the Edge Manager captures the impact of different workloads in Edge servers and adjusts vertical scaling (considered in Section 4). For example, for a large workload in PFR and Hybrid pricing models, its PS will be set higher in the next round of scaling since compared to other workloads it requires more resources to maintain good performance. However, in the PFP pricing model, this workload will receive lower priority in the next scaling round because more resources may have to be added to the workload without being paid for.

A limitation of *wDPS* is that workload-related factors are objective and not controlled by the Cloud Manager. Therefore, a subjective factor to volunteer resources in exchange for rewards is further proposed to enrich DPM.

3.2.2. Community-aware Dynamic Priority Score (*cDPS*)

When an Edge server is scaled, it may be allocated the same resources as before, i.e. scaling was not required. This may occur when the Edge server has met its SLO but does not meet a specified threshold (considered in Section 4) to be scaled down. Under such circumstances, the Edge server

could still be scaled down as long as the risk of degraded performance is acknowledged. Reward credit is given to an Edge server if in a previous scaling round it agreed to donate to a shared resource pool given that this Edge server had met its SLO. This is derived from the community model in economics, which encourages voluntary contributions. In this paper, we implement resource donation as an incentive for Edge servers to share resources allocated to them on an Edge node. This is beneficial when no additional resources can be requested from the Edge node during busy hours since the resources may be adjusted locally within Edge services.

cDPS of an Edge server s is defined as:

$$cDPS_s = wDPS_s + W_{Reward} * Reward_s \quad (6)$$

With cDPS, Edge servers take partial control of their priority, which is in contrast to wDPS. Contrary to the premium service defined in SPM, the reward is free and could be added frequently in the life cycle of an Edge server.

The above DPS' do not consider the impact of frequent scaling. When many servers are running on an Edge node, continuous scaling could result in large overheads and a slower system response, which affects the performance of all Edge services. Hence, a penalty is imposed for frequently scaling in the next DPS approach.

3.2.3. System-aware Dynamic Priority Score (sDPS)

Frequent scaling may result in an increased overhead when consecutive scaling rounds occur between short time intervals. This is considered when updating the priority of all Edge servers. If an Edge server has to be scaled many times to not violate its SLO, then its PS will be set lower in the next scaling round as a penalty for slowing the system. One reason for frequent scaling may be due to unrealistic SLOs which are set by the Cloud Manager. Therefore, any adverse impact on the Edge node can be mitigated with a lower priority. sDPS of an Edge server s is defined as:

$$sDPS_s = cDPS_s + \frac{1}{W_{Scale} * Scale_s} \quad (7)$$

By implementing sDPS, the Edge Manager penalises servers that slow down the Edge node.

4. Dynamic Vertical Scaling

Dynamic vertical scaling is a mechanism to allocate or deallocate Edge resources for a workload at runtime. Efficient resource management is essential to better utilise Edge resources and ensure

that the overall QoS is not compromised. Most resource management mechanisms consider resource provisioning during workload deployment but ignore the need for post-deployment resource adjustment (after the workload has started execution). Without an efficient dynamic vertical scaling technique, an Edge node could be overloaded when executing bursty workloads. If additional resources are not allocated to the workload, then SLO violations are likely to occur. Therefore, a mechanism that is constantly aware of resources on an Edge node and makes scaling decisions for resource allocations periodically is necessary for an Edge environment.

Since it is disadvantageous to scale resources for multiple workloads concurrently (as presented in Section 3), a priority-based dynamic vertical scaling mechanism is required to reallocate resources for every Edge server. Workloads with the highest PS should be considered first and the one with the lowest PS at the end of a scaling round until there are no resources to support Edge servers with low priorities. Workloads that have insufficient resources will not be executed on the Edge node and will need to be executed elsewhere. In this research, workloads are deployed on the Edge node using LXC containers².

4.1. Priority-based dynamic vertical scaling

Procedure 1 presents a dynamic vertical scaling mechanism using priorities. The PS of a list of Edge servers executing on the Edge node is updated at the beginning of each round of vertical scaling (Line 1). The PS of an Edge server is calculated with the static and dynamic priority approaches defined in Section 3. The list of servers is then sorted by the updated PS (Line 2). The server with the highest priority is firstly considered by the Auto-scaler. The Auto-scaler checks if there is a need for the Edge server on the node (whether users are covered by this Edge node or whether the QoS of the workload can be improved on the Edge node; Line 4). Network latency is used to decide whether the Edge server can deliver the desired improvement or whether terminating the Edge server and servicing the users through the original Cloud server can be of more benefit (Line 20). In this paper, we assume that Edge nodes are not interconnected and therefore a workload cannot be migrated from one Edge node to another directly. Migration would need to occur via the Cloud.

The average latency is compared with the SLO provided by the Cloud Manager in its service request (Line 5). If the latency is higher than the SLO (i.e. the Edge server has not been performing

²<https://linuxcontainers.org/lxc/introduction/>

Procedure 1: Dynamic vertical scaling mechanism

Data: $S, L_s, uR, donation_s$

```

1 Update  $PS$  of all running Edge servers;
2 Sort  $S$  by  $PS$ ;
3 for  $\forall s \in S$  do
4   if  $s$  is active and network latency is acceptable then
5     if  $aL_s > L_s$  then
6       calculate  $VR_s$ ;
7        $scale(scaleup, VR_s, S)$ ;
8     else
9       if  $aL_s > dThr * L_s$  then
10        if  $donation_s == 1$  then
11           $scale(scaledown, uR, S)$ ;
12        else
13          no scaling for  $s$ ;
14        end
15      else
16         $scale(scaledown, uR, S)$ ;
17      end
18    end
19  else
20     $terminate(S, s, U_s)$ ;
21  end
22 end

```

as expected), then the container hosting the Edge server will be allocated more resources. The ratio of resources to be added is based on the SLO violation rate of the server (Line 6). This is to make sure that resources are adjusted based on a server’s performance. If the latency is lower than a pre-defined percentage $dThr$ of the SLO, for example, 80%, then resources are removed from the container hosting this server (Line 16). When latency is between the threshold and the objective, resources are not scaled down since performance may be negatively affected (Line 13). However, if the server is willing to donate resources in exchange for priority credits, it is scaled down and the number of times that the Edge server has donated resources is updated (Lines 10-11). If resources of a server are scaled, then the total number of times it has scaled is recorded. However, note that the scaling that occurs when resources are donated is not recorded, since it is not used for penalising a server. Sorting by PS in Line 2 of the above procedure has a computational complexity of $O(n \log n)$, where n is the number of running servers on an Edge node. The priority-based dynamic vertical scaling loops through the n running Edge servers, resulting in a complexity of $O(n)$. The total computation complexity of applying DYVERSE on an Edge node is then $O(n \log n)$ if there are 10 or more running Edge servers, or $O(n)$ if there are less than 10 running Edge servers; the actual complexity may be less since some of the low-priority Edge servers may be terminated before one round of the priority-based dynamic vertical scaling completes.

4.2. Multi-tenant Management

When a server is scaled, then the scaling mechanism as shown in Procedure 2 first checks the decision on whether to “scaleup” or “scaledown” (Lines 1 and 17). To scale up, i.e. to allocate more resources, the Monitor firstly calculates the amount of resources to add (Line 3) and then checks if there are additional available resources on the Edge node to support this (Line 4). If resources are available, then they are added to the container (Line 5). Re-configuring the resource limits of a container is realised through the *cgroup* command, i.e. control group, which is a feature of Linux kernel that limits, accounts for and isolates the resource usage (for example, CPU and memory) of a container³. If the available resource is not sufficient to support scaling up, then the container with the lowest priority in the container set will be terminated to release its resources so that there are more resources available (Line 8). During this process services for users connected to

³<https://linuxcontainers.org/lxc/manpages/man1/lxc-cgroup.1.html>

293 this Edge server are still available through the Cloud server. The containers with lower priorities
 294 will be terminated sequentially until the updated FR is sufficient to support the scale-up request or
 295 there are no more containers with lower priorities (Lines 10-13). To scale down, a unit of resource
 296 is removed from the server (Line 18). At the end of the scaling process, the Edge server is updated
 297 with a new quota of resources (Lines 5, 14 and 18).

Procedure 2: Scaling mechanism

Data: $decision, uR/VR_s, S$

```

1  if  $decision == scaleup$  then
2    Measure  $R_s, FR$ ;
3     $aR_s = R_s * VR_s$ ;
4    if  $FR \geq aR_s$  then
5      | Add  $aR_s$  to  $s$ ;
6    else
7      while  $FR < aR_s$  do
8        | terminate( $s_n, S, U_{s_n}$ );
9        | //  $s_n$  is the last server in  $S$ 
10       | measure  $FR$ ;
11       | if  $n == index\ of\ s$  then
12         | | break;
13       | end
14     | end
15     | Add  $FR$  to  $s$ ;
16   end
17 if  $decision == scaledown$  then
18   | Remove  $uR$  from  $s$ ;
19 end
  
```

298 Figure 2 illustrates the application of DYVERSE on an Edge node using an example. When
 299 applications (for example App a, b and s) from the Cloud are deployed into an Edge node, they have

Figure 2. Example application of DYVERSE on an Edge node.

an initial priority level, such as purchased priorities (for example, application a has a higher priority
 than b and b has a higher priority than s . The dynamic vertical scaling mechanism (marked in grey
 block) that was proposed in this section is carried out periodically, for example, every 5 minutes on
 the Edge node. If the Edge node is configured with the static priority management (SPM) defined in
 Section 3.1, the order of these applications being considered in the first round of the dynamic vertical
 scaling is calculated with Equation 3. For example, the order of applications in Figure 2 with SPM
 is a , s and b . Then in the subsequent rounds, the priority scores of these applications stay the same,
 because with SPM the priority scores of applications are only calculated during deployment. On the
 contrary, if the Edge node is configured with the dynamic priority management (DPM) defined in
 Section 3.2, the order of these applications being considered in every round of the dynamic vertical
 scaling is updated. For instance, the order of applications in Figure 2 with DPM in the first round is
 s , b and a . In the second round of dynamic vertical scaling, the order is updated to b , s , and a . The
 dynamic priority scores of these applications are calculated using either Equation 4, Equation 5,
 Equation 6 or Equation 7, depending on the selected DPM approach.

When Edge applications are initially deployed, an equal amount of resources are allocated to
 them, which is represented by a uniform size of the box. Before each round of dynamic vertical
 scaling (grey block), the order of Edge applications are either kept the same by SPM or updated
 by DPM, then the resources allocated to each of the applications are checked by Procedure 1.
 During the periodical dynamic vertical scaling, the resources allocated to Edge applications change

over time, which is represented in Figure 2 with changing sizes of the box. The decision to scale up/down or perform no scaling is based on monitoring the QoS of applications. In addition, the selected priority management approach also has an impact on the dynamic vertical scaling. For example, in the three rounds of dynamic vertical scaling illustrated in Figure 2, application *b* is completely removed from the Edge node when SPM is applied whereas it could have benefited from more Edge resources in the first two rounds when DPM is applied.

4.3. Discussion

In Edge computing, Edge servers when compared to Cloud servers will be used for shorter time intervals given the demand and limited availability of resources on Edge nodes. Therefore, in addition to dynamic vertical scaling, the Edge Manager will need to decide when an Edge server is removed from the Edge node. The Edge Manager terminates (i.e., the service does not run on the Edge, but continues to be hosted on the Cloud server) service in the following three cases, which are considered in dynamic vertical scaling. Firstly, there are no free resources to support the Edge service. Secondly, the Edge service is not required anymore (the Edge server has been idle for a while). Thirdly, the Edge service does not improve the QoS of the workload (the performance constraints cannot be satisfied by an Edge server).

When an Edge server is terminated, the associated data containing local updates are migrated to the Cloud (local data is appended to global data maintained by the Cloud Manager. This is realised through a key-value based data store, Redis⁴. Such a migration is use case dependent and only required by those applications that maintain data states across different geographic locations. For example, the Edge server of an application that processes video streams and does not require data to be stored could be terminated straightforwardly without requiring any data migration. In addition, for applications that require data migration, we assume the benefit of using the Edge outweighs the communication costs since only user data is moved from the Edge to the Cloud, rather than the entire server. Servers that do not rely on databases would have minimal data movement. User devices are redirected to the Cloud server during the termination process. For example, in the case of a location-based mobile application, location-specific data is maintained on the Edge node and would need to be migrated to the Cloud when the application no longer

⁴<https://redis.io/>

executes on the Edge. It is noted that DYVERSE can benefit applications with different workflows in improving their overall QoS - applications that maintain their data state on the Edge and those in which the data is stateless on the Edge.

5. Experimental Evaluation

In this section, the priority management approaches presented in Section 3 and the dynamic vertical scaling mechanism proposed in Section 4 are evaluated. The experimental setup, including the hardware platform and the distributed workloads employed in this research are firstly presented followed by the merits of DYVERSE against metrics including system overhead, SLO violation rate, and latency.

5.1. Setup

A Cloud-Edge platform is set up using Amazon Web Services Elastic Compute Cloud (AWS EC2) and an ODROID-XU board. On the Cloud, t2.micro Virtual Machine (VM) running 14.04 LTS from the Dublin data centre is used to host the Cloud server. Although we employ a low-cost and basic VM, we note that in this paper we make no comparisons about the compute capabilities between the Cloud server and the Edge server. Therefore, in no way would the results presented in this paper be affected had we chosen a larger VM from the same data centre.

The Edge node is located in the Computer Science Building of Queen’s University Belfast in Northern Ireland. The board has 2 GB of DRAM memory, and one ARM Big.LITTLE architecture Exynos 5 Octa processor running Ubuntu 14.04 LTS. Each server on the Edge node is hosted in an LXC container.

5.2. Workloads

Two workloads are chosen to evaluate DYVERSE: a location-based mobile game and a real-time face detection workload. Both workloads are server-based and a natural fit for Edge computing since they are latency critical – response time is affected by the distance between user devices and the server. Hence, a subset of the functionalities of the Cloud server can be brought closer to devices.

The above workloads are also representative of different workloads that can benefit from DYVERSE: the mobile game represents a multi-user application whose Edge server responds to incoming user requests; the face detection workload, in contrast, is representative of a data-intensive

streaming application, in which case the Edge server pre-processes incoming data and relays it to the Cloud.

(i) *Location-based mobile game*: The application is an open-source game similar to Pokémon Go, named iPokeMon⁵. iPokeMon comprises a client for the iOS platform, which can be used on mobile devices, and a server that is hosted on a public Cloud. User navigates through an environment in which virtual creatures named Pokémons are distributed. The iPokeMon game server was redesigned to be hosted on the Cloud and an Edge node. The Edge hosts a latency-sensitive component that updates the virtual environment as a user navigates; for example, the GPS coordinates of the player and the Pokémons. The local view on the Edge server is updated by frequent requests sent to the Edge server. If user requests are serviced from a distant data centre, then user experience is affected due to lags in refreshing. Hence, the Edge is beneficial to reduce latencies for this workload. Up to a maximum of 32 Edge servers are hosted, with each server randomly supporting between 1 and 100 users.

The server on the Edge node is tested using Apache JMeter⁶. One session of a connection (a user is playing the iPokeMon game) between the user device and the Edge server hosted in the LXC container is recorded for 20 minutes. During this time the number and type of requests and the parameters sent through the requests are recorded. Subsequently, JMeter load tests single and multiple Edge servers by creating virtual users and sending requests to the Edge servers from the virtual users in the experiments.

(ii) *Real-time face detection*: The original Cloud-based Face Detection (FD) workload captures video via a camera (such as CCTV footage) and transmits it to the Cloud where it is converted to greyscale (this is one-third the size of the color video) and then detects faces on it using OpenCV⁷. Pre-processing (converting to greyscale) is performed on the Edge and hosted via LXC containers. The converted stream is sent to the Cloud for reducing the bandwidth used instead of sending colour video to the Cloud. Up to 32 servers are load tested on the Edge node, with each server randomly pre-processing between 0.1 and 1 frame per second.

The experimental evaluation demonstrates: (i) the impact of varying priority factor weights on the priority score, (ii) the overheads in priority management and the overheads of dynamic vertical

⁵<https://github.com/Kjuly/iPokeMon>

⁶<http://jmeter.apache.org/>

⁷<https://opencv.org/>

scaling in a multi-tenant environment, (iii) the effect of DYVERSE on SLO violation rates, and (iv) the effect of DYVERSE on the latency of the above workloads.

5.3. Results

The experiments provide insight into the benefits of using the priority approaches and dynamic vertical scaling mechanism in a variety of scenarios, such as moderate bandwidth-consuming and bandwidth-hungry tasks, single and multi-tenant servers and varying user-defined SLOs.

5.3.1. Impact of Priority Factor Weights

In order to understand the impact of the static and dynamic factor weights introduced in Section 3, we firstly study the priority scores of a particular Edge server when the weights are ranged in $[0, 1]$. The nine individual factor weights are grouped into four based on the priority management approach they belong to, namely the factor weights for *SPS* in Section 3.1, the factor weights for *wDPS* in Section 3.2.1, the factor weights for *cDPS* in Section 3.2.2, and the factor weights for *sDPS* in Section 3.2.3.

Figure 3 presents the priority scores of an example Edge server when the priority weights have different values. The x axis range in $[0.1, 1]$ because the dynamic priority approaches are developed incrementally from the static priority as shown in Equation 3, Equation 4, Equation 5, Equation 6 and Equation 7. The comparisons between each pair of directly related priority approaches show that when the factor weights increase, the priority of the Edge server increases. By comparing the three sub-figures, it is inferred that the influence of the factor weights on the priority score, from the strongest to the weakest, is *SPS*, *wDPS*, *sDPS*, and *cDPS*.

There is an indirect impact of different priority scores resulted from the different configurations of factor weights on the overall performance of an Edge system. For example, in Figure 3a when the *wDPS* factor weight is set as 0, the priority score is the same as adopting the *SPS* approach. When the *wDPS* factor weight is set lower than the *SPS* factor weight, the priority score is mainly affected by the four factors defined in Table 2, whereas the workload factors defined in Table 3 have less impact on the priority score. This could lead to a situation when an idle Edge application has a similar priority score as a busy Edge application that is servicing a large number of user requests, which further causes unfairness.

DYVERSE aims to reduce the average SLO violate rate for all running applications on an Edge node, as defined in Equation 2.1. Therefore, it is necessary to consider both application-dependent

(a)

(b)

(c)

Figure 3. Priority scores of an Edge server when the priority factor weights vary.

(a) Priority management

(b) Dynamic vertical scaling

Figure 4. Overhead of priority management and dynamic vertical scaling per Edge server for iPokeMon and Face Detection (FD).

factors (e.g. workload factors) and system-related factors (e.g. reward and penalty. In the following experiments, the weights of Equation 3 to Equation 7 for calculating the priority scores are set to 1 to balance the impact of the different groups of factors.

5.3.2. System Overhead

Edge platforms are expected to have more transient system states than the Cloud, i.e. Edge workloads have shorter life cycles, so lower overheads in resource management are better. We measured the overheads of priority management (Figure 4a) and dynamic vertical scaling (Figure 4b) of each server on the Edge node for iPokeMon and FD. It should be noted that the servers are still servicing requests when priorities are updated. Similarly, the overhead of dynamic vertical scaling, which is the time taken to reallocate the resources of the server was measured. During this time the servers continue to service requests. Although three dynamic priorities (wDPS, cDPS, and sDPS) were proposed, only sDPS is considered for comparing the overheads between SPM and DPM; we noted that different dynamic priorities did not affect the overall delay in the experiments.

Figure 5. Violation rate when the SLO is 78ms for iPokeMon and 2.13s for FD. 32 Edge servers are deployed. Red blocks correspond to the overhead time when priority-based scaling occurs.

Figure 4a shows that SPM incurs small overheads, but it is more costly to employ DPM, especially for iPokeMon. For example, for 32 servers the overhead of using DPM with iPokeMon is 150 milliseconds (ms) when compared to 10ms for SPM. This is because static priorities assigned to servers do not change during execution and monitoring is needed for dynamic priorities. iPokeMon server records all requests and responses in a log file while supporting multiple users and the Monitor reads through this file in every scaling round to update the dynamic priority. The FD server maintains the values of dynamic factors while processing requests from the single streaming source and the Monitor uses these values to calculate the PS. One insight from the difference in DPM overheads is that monitoring workloads could be accounted for when designing Edge applications.

The overhead of dynamic vertical scaling in Figure 4b using DPM is on average higher than using SPM. This is because, during scaling, dynamic factors such as reward and penalty are monitored for each server. More time is required to scale iPokeMon servers because during termination user data needs to be migrated from the Edge to the Cloud. When an FD server is terminated no data is migrated. As described in Section 4.3, the termination procedure is use case dependent and only applies to iPokeMon in this paper. The application of DYVERSE on two use cases is to demonstrate the different types of workflows that can benefit from dynamic vertical scaling post-deployment of the application.

To understand the impact of the observed overheads on the QoS of iPokeMon and FD, the average SLO violation rate of 32 Edge servers every one minute is presented in Figure 5. The first five minutes window is the period when the Edge servers are launched with an equal amount of resources. During the 5th, 10th, and 15th minutes, priority-based dynamic vertical scaling

(Procedure 1) is executed. For DPM, the PS is first calculated for each server, after which scaling occurs. On the other hand for SPM, scaling based on the same initial PS is implemented. Updating priorities of 32 servers requires 30–40 seconds (s) and 15s for iPokeMon and FD, respectively. We infer from the figure that after the first scaling round, the SLO violation rates have decreased on average by 4% and 3% for iPokeMon and FD respectively. In subsequent scaling rounds, it is noted that the violation rates are further reduced. DPM on an average performs nearly 2% better than SPM. The figure highlights that the sub-minute delay caused due to priority management and scaling does not necessarily affect the QoS of our use-cases.

5.3.3. SLO Violation Rate

Meeting SLOs is key to achieving a high overall QoS. Higher SLO violation rates at the Edge indicate the possibility of losing Edge customers. Consequently, this results in the loss of revenue for the Edge service provider. Therefore, the SLO violation rate is chosen to highlight the difference that DYVERSE makes to the QoS of Edge services.

From the empirical analysis, we observed that the average time to service an individual request for iPokeMon and FD is 78ms and 2.13s, respectively. The aim is to study the performance of DYVERSE when different tolerance for the SLO violation is accepted. This is useful for understanding in what situations, for example when the SLO is more stringent or lenient, DYVERSE can perform better. Therefore, this study is pursued with SLOs that considers the average time taken to service a request and an upper tolerance of both 5% and 10% of the average time. For iPokeMon the SLOs are 78ms (0% tolerance), 82ms (up to 5% tolerance), 86ms (up to 10% tolerance) and for FD the SLOs are 2.13s (0% tolerance), 2.24s (up to 5% tolerance) and 2.34s (up to 10% tolerance).

The baseline we use for understanding the benefit of dynamic vertical scaling in static and dynamic priority management is by comparing them against a scenario in which there is no dynamic vertical scaling. This is because there are currently no alternative resource scaling solutions on a single Edge node. We understand the benefit of using dynamic priority management by comparing it only against static priority management so that the benefit of the former is not magnified.

Figure 6 and Figure 7 present the SLO violation rates of the two use cases when the SLO tolerates 0% (Figures 6a and 7a), up to 5% (Figures 6b and 7b) and up to 10% (Figures 6c and 7c) violations of the average application latency. The general trend is that for both workloads violation rates can be reduced by using the proposed scaling mechanism, either with SPM or DPM,

(a) SLO: 78 milliseconds

(b) SLO: 82 milliseconds

(c) SLO: 86 milliseconds

Figure 6. Violation rates of iPokeMon when applying varying SLOs.

(a) SLO: 2.13 seconds

(b) SLO: 2.24 seconds

(c) SLO: 2.34 seconds

Figure 7. Violation rates of the face detection workload when applying varying SLOs.

when compared to a scenario without dynamic vertical scaling. When there is no dynamic vertical scaling, the containers hosting the Edge servers continue to retain the same amount of resources they were allocated before they started execution. No additional resources are dynamically allocated to containers. For example in Figure 6a, when the SLO is 78ms for 32 servers and 149 KBs of data is processed per second, nearly 18% of the requests from users are violated. Scaling with SPM on an average reduces this violation by 4%. A reduction of 6% in SLO violation rate by applying SPM in FD is observed (Figure 7a) when the SLO is 2.13s for 32 Edge servers and 4 MBs of data is processed per second.

The scaling mechanism takes the server's performance into account. During dynamic vertical scaling, the measured latency is compared with the user-defined SLO. If there were violations, then the container hosting the workload would be allocated additional resources and vice versa. Allocation to one container is done by requesting free resources from the Edge node or deallocating resources from other containers that have lower priorities. This optimises the use of Edge resources and can reduce the overall violation rates. This is noted for all SLOs for both workloads. For example, in Figure 7c for an SLO of 2.34s scaling with SPM reduces the violation rate for FD by 3% over 32 servers (for iPokeMon with an SLO of 86ms in Figure 6c the violation rate is reduced by 4% over 32 servers).

However, the static priority of a server is set once before execution and does not change until it is terminated and needs to be executed again. This is a disadvantage since servers cannot change their priority during execution. Hence, we employ dynamic priorities that can be changed after deployment and during execution. Priorities can be implemented to take other factors, such as the number of users into account (in addition to factors listed in Table 2, DPM considers factors presented in Table 3). Resources can be scaled based on changing priorities. Dynamic priorities are better than static priorities since they further reduce the violation rates. When the violation threshold is stringent (SLO is 78ms for iPokeMon and 2.13s for FD), scaling with DPM reduces the violation by 1% (Figure 6a) and 4% (Figure 7a). Although we proposed three dynamic priorities (wDPS, cDPS, and sDPS) approaches, we have only considered sDPS for comparing the trends between no dynamic vertical scaling and scaling with SPM. This is because we observed that different approaches did not affect the overall violation rate. The effect of the three approaches is nonetheless considered in the context of latency.

When the violation threshold is lenient for iPokeMon (SLO is 86ms), a larger volume of requests

can be fully serviced by the servers resulting in no violations for up to 32 Edge servers (Figure 6c). In these cases, requests that would result in a violation when the threshold is more stringent can now complete execution within the specified SLO. Subsequently, the benefits of DYVERSE becomes less important when the Edge node is less stressed. It is also noted that a scale-up operation presented in Procedure 2 will not occur. However, containers may scale down during this time. With the same violation threshold, DPM for FD also results in a lower SLO violation rate, although the improvement is not as significant as in iPokeMon. This is because of larger data sizes - data transferred in FD is 30–150 times larger than data in iPokeMon. Images are frequently uploaded between the end device and the FD Edge servers, which results in longer processing times.

5.3.4. Latency of Workloads

The results shown above presented the violation rate when compared to the number of servers, but did not differentiate the three DPM approaches. This section aims to explore the impact of using different priority management approaches on latency. For iPokeMon, we define latency as the average round-trip time taken to service one request originating from the user by the Edge server. For FD, we define latency as the average single-trip time taken to detect faces by the Cloud server in one video frame originating from the camera. Figure 8 and Figure 9 show the distribution of application latency for the two use cases when the SLO tolerates 0% (Figures 8a and 9a), 5% (Figures 8b and 9b) and 10% (Figures 8c and 9c) of the average application latency. The distribution provides: (i) a time profile of requests violated, and (ii) the impact of DYVERSE on requests that are not violated.

It is inferred that when the SLO violation threshold is 5% and if no dynamic vertical scaling is used, then no requests are serviced in the lowest time band (from 0 to 80% of the SLOs). However, the number of requests that are serviced in the lowest time band increases when employing static priority (the normal in the distribution shifts to the left). For example, for an SLO of 82ms for iPokeMon (Figure 8b), 4% of requests are serviced within 66ms (the leftmost time band in the figure). This is due to the reallocation of resources from containers that meet SLOs to containers that violate SLOs in dynamic vertical scaling. Such a benefit is less obvious in FD, but a 1% improvement is noted by applying dynamical vertical scaling for an SLO of 2.24s (Figure 9b).

Dynamic priorities reduce the latency of both workloads when compared to static priority because they consider additional factors, such as the number of users connected to each server. In

(a) SLO: 78 milliseconds

(b) SLO: 82 milliseconds

(c) SLO: 86 milliseconds

Figure 8. Latency of iPokeMon when 32 servers are deployed on the Edge node.

(a) SLO: 2.13 seconds

(b) SLO: 2.24 seconds

(c) SLO: 2.34 seconds

Figure 9. Latency of face detection when 32 servers are deployed on the Edge node.

most cases, sDPS performs better than cDPS and wDPS. For example, given an SLO of 86ms for iPokeMon (Figure 8c), 20% of requests are serviced in the lowest time band using sDPS in contrast to 18% for cDPS and wDPS. Similar results are observed in FD. For instance, given an SLO of 2.13s (Figure 9a), 2% of video frames are processed by the face detector in the lowest time band using sDPS in contrast to 1% for cDPS and wDPS. This is attributed to the penalty that is imposed on containers that scale. As presented previously, scaling incurs an overhead and therefore the Edge Manager aims to minimise the impact of scaling by reducing priorities as shown in Equation 7.

5.3.5. Summary

The results obtained from the experimental evaluation considered above are summarised as follows:

1) *Overheads in using priorities and dynamic vertical scaling*: the overheads do not affect the execution of the servers and the servers continue to service user requests. In priority management, DPM takes a longer time if the monitoring of workload-related metrics is not incorporated in the Edge server (e.g. iPokeMon). Comparing with SPM, DPM has a 4%–8% higher overhead for dynamic vertical scaling.

2) *SLO violation rates for stringent thresholds (the average time to service one request)*: the SLO violation rate when there is no dynamic vertical scaling is observed to be nearly 18% and 23% for iPokeMon and FD respectively. Scaling with SPM improves the violation rate by 4% and 6% for iPokeMon and FD respectively when compared to no dynamic vertical scaling. Scaling with DPM improves the violation rate of iPokeMon by 1% when compared to SPM. This is confirmed by the 3% reduction of SLO violation rate of the face detection workload when comparing DPM and SPM.

3) *SLO violation rates for lenient thresholds (10% more of the average time taken to service one request)*: the SLO violation rate when there is no dynamic vertical scaling is observed to be nearly 12% and 18% for iPokeMon and FD respectively. Scaling with SPM improves the violation rate by 3% for FD and 4% for iPokeMon when compared to no dynamic vertical scaling. Scaling with DPM ensures no requests are violated for iPokeMon and improves the violation rate by 3% for FD when compared to SPM.

4) *Latency for stringent thresholds*: without dynamic vertical scaling, no requests are serviced within 66ms for iPokeMon and within 2.13s for FD. Using sDPS with iPokeMon, nearly 13%

of requests are serviced in 66ms–74ms, which is higher than cDPS and wDPS by 5% and 10% respectively. Using sDPS with FD offers a small advantage; 1% more requests are processed in 1.79–2.02s. DPM performs better than SPM. Using sDPS, 20% of requests of iPokeMon are serviced within 72ms, which is higher than SPS, cDPS, and wDPS by 4%, 2% and 2% respectively. Similarly for FD, 2% of video frames are processed in 1.7s–1.92s, which is higher than SPS, cDPS and wDPS by 2%, 1% and 1% respectively.

5) *Latency for lenient thresholds* : Fewer requests are serviced within 69ms for iPokeMon and 1.87s for FD with no dynamic vertical scaling compared to using it. Using sDPS with iPokeMon, 20% of requests are serviced in the lowest time band, which is higher than cDPS and wDPS by 2%. Using sDPS with FD, 14% of requests are serviced in the lowest time band, which is higher than cDPS and wDPS by 1% and 2% respectively. DPM performs better than SPM.

6. Related Work

Resource scaling is well studied in distributed systems and more recently on the Cloud. A classification of existing research is shown in Figure 10. In this article, six classifications are highlighted based on (i) the direction of scaling, (ii) the parameters that are optimised in the problem space, (iii) the virtualisation techniques that the resource scaling solutions target at, (iv) the number of applications the scaling algorithms consider, (v) the time when the resource scaling is performed, and (vi) the importance of workload being considered.

Based on the direction of resource scaling, computing resources for a workload can be allocated or deallocated from a cluster either horizontally [51] or vertically [17]. Horizontal scaling, also referred to as scale-out/scale-in, is an action to add/remove infrastructure capacity in pre-packaged blocks of resources. In addition to horizontal scaling, vertical scaling, also referred to as scale-up/scale-down, is an action to add/remove resources to an existing system [17].

The focus of this paper is vertical scaling on the Edge. The vertical scaling approaches used on the Cloud are not suited for the Edge due to three reasons. Firstly, monitoring techniques employed for vertical scaling are heavyweight (i.e., time consuming and require significant resources) [53], which cannot be directly applied to resource-limited Edge environments (a large cluster of resources may not be available at the edge of the network). Secondly, generating optimal scheduling plans for vertical scaling as employed in the Cloud [54] is prohibitive on the Edge. Integer linear programming or constraint programming that is usually used to generate optimal solutions are computationally

Figure 10. A classification of existing research on resource scaling.

intensive. Thirdly, vertical scaling on the Cloud is underpinned by workload prediction models [52], which are less suitable for use on the Edge. Edge services are likely to be short-running workloads in contrast to the Cloud and only limited data may be available for training machine learning models. Additionally, training models on the Edge may not be feasible. Therefore, in this paper, we propose a low overhead vertical scaling technique on an Edge node.

Based on the parameters that are optimised in the problem space, two important parameters are frequently considered, namely job deadline and monetary cost. Workload deadline has been employed as an important parameter to optimise during resource scaling [48, 49], as it directly represents the application performance. Moreover, cost budget also plays an important role [50, 49] because in many cases the more resources being utilised, the higher the monetary cost it entails.

Based on the virtualisation techniques being targeted, existing research has proposed resource scaling algorithms for both virtual machines [44] and containers [20, 46]. Given different types of VMs with variable CPU and memory resources, sometimes workloads are simply migrated onto another VM with resources more suitable to meet the optimising criteria (instead of scaling the resource on the original VM) [45].

Based on the number of applications being considered in the computing environment, resource scaling solutions have been designed for single tenancy [39] and more often multi-tenancy [40]. Multi-tenancy refers to the co-location of different workloads on the same computing resource. In a multi-tenant Cloud environment, migration is employed to resolve resource contention [40, 41]. Inter-tenant resource trading and intra-tenant weight adjustments are employed on the Cloud to ensure fairness when scaling [13].

Techniques employed on the Edge to deal with multi-tenancy need to be again lightweight and cannot incorporate complex models for workload predictions [24]. Complex schemes, such as convex optimisation, game theory or dynamic programming have been explored for Edge environments [55]. Heuristic-based mobile application partitioning and offloading algorithms have been designed to maximise the application's performance for multiple users [42] and the system utility [43]. These algorithms are scalable but still require $O(K^2)$, $O(K^3)$ time respectively, to service K mobile users. This is impractical on resource-constrained Edge nodes that use small processors since they cannot acquire hardware resources easily. Similarly, game theory based schemes have been designed for Mobile-Edge Cloud computing [56]. This approach is tested on a cloudlet, which naturally assumes more resources than the extreme edge of the network (cloudlets have dedicated

650 computing servers/clusters).

651 *Based on the time when resource scaling actions are performed*, the adjustment of resource
 652 allocation could happen either before or after the deployment of applications. Research on multi-
 653 tenancy at the Edge usually focuses on pre-deployment. Once a workload is deployed, the resources
 654 allocated to it do not change. Similarly, the QoS of applications is improved without considering
 655 multi-tenancy after deployment [38]. To the best of our knowledge, research on post-deployment
 656 resource scaling in Edge computing is sparse at best. However, multi-tenancy significantly impacts
 657 the SLO violation rate after deployment. In contrast to existing research that is reported in the
 658 literature on multi-tenancy in Edge computing, this research accounts for post-deployment condi-
 659 tions of workloads for scaling resources in a multi-tenant environment. This paper demonstrates
 660 that dynamic vertical scaling can improve the overall QoS of the application during its lifetime on
 661 the Edge after it is deployed.

662 *Based on the importance of workloads*, some resource scaling algorithms are priority-based
 663 whereas others treat the application equally important. Community models are used to develop a
 664 shared resource pool in Clouds [37]. The priority of Edge workloads is different because relation-
 665 ships between (i) the Cloud and the Edge, and (ii) the user and the Edge need to be accounted
 666 for. There is research that (i) adopts a threshold-based priority function to decide on partial or
 667 complete offloading of workloads [35], and (ii) adjusts the task execution order of workloads to sat-
 668 isfy different objectives in the pre-deployment phase [27, 36]. These show the impact of priorities
 669 on offloading, but do not demonstrate the influence of priorities on post-deployment resource man-
 670 agement. In addition, existing resource management solutions that consider workload priority only
 671 takes a single factor that impacts the priority, for example, the purchased priority [31], the sequence
 672 of incoming workloads [32] or the potential resource starvation [33]. To advance the state-of-the-art,
 673 the research reported in this paper considers a collection of static factors that are specific to the
 674 Edge, as well as dynamic factors that vary during the execution of workloads. These factors all
 675 together contribute to the workload priorities for efficiently managing resources in a multi-tenant
 676 environment.

677 7. Conclusions

678 Distributed applications will leverage the edge of the network to improve their overall QoS for
 679 which the challenge of multi-tenancy in resource-constrained environments need to be addressed.

Vertical scaling of resources is required to achieve multi-tenancy. However, existing mechanisms require significant monitoring and are computationally intensive since they were designed for the Cloud. They are not suitable for resource-limited Edge.

The research in this paper addresses the above problem by developing DYVERSE, a **DY**namic **VERT**ical **S**caling mechanism on **E**dge environments to facilitate multi-tenancy. The aim is to maximise the QoS of workloads executing in a multi-tenant Edge node without violating user-defined Service Level Objectives (SLOs). The proposed scaling mechanism generates post-deployment plans for workloads during execution so that SLOs are not violated. The mechanism is underpinned by one static and three dynamic priority management approaches. Three dynamic priorities that are workload-aware, community-aware and system-aware are proposed.

The feasibility of DYVERSE is validated using two workloads, namely a location-aware online game and a real-time face detection application in a Cloud-Edge environment. The priority-based dynamic vertical scaling has less than a sub-second overhead per server when 32 Edge servers are deployed on a single Edge node. Scaling with static priorities reduced SLO violation rates of user requests by up to 4% and 12% for the online game respectively and in both cases 6% for the face detection workload, when compared to no dynamic vertical scaling. For both workloads, system-aware dynamic vertical scaling effectively reduced the latency of non-violated requests, when compared to other methods.

7.1. Limitations and Future Work

Although system-related parameters that affect SLO violation rates and latency are considered, network-based parameters, such as bandwidth and how it affects the QoS were not considered. This is a part of our ongoing research and will be reported elsewhere.

An extensive study of the direct impact of different configurations of the factor weights on the system performance is not included in this paper. The current experimental results demonstrate a proof-of-concept of the feasibility of multi-tenant Edge environments. In the future, we aim to integrate the tuning of weights to optimise system performance.

The dynamic vertical scaling mechanism proposed in this paper considers only QoS metrics of an Edge server, namely workload latency and SLO violation rate, to invoke resource scaling. Utilisation metrics of containers, such as memory and swap space used for hosting Edge servers, could be further considered.

References

- [1] B. Varghese, R. Buyya, Next Generation Cloud Computing: New Trends and Research Directions, *Future Generation Computer Systems* 79 (3) (2018) 849–861.
- [2] B. Varghese, N. Wang, D. Nikolopoulos, R. Buyya, Feasibility of Fog Computing, in: R. Ranjan, K. Mitra, P. Prakash Jayaraman, L. Wang, A. Zomaya (Eds.), *Handbook of Integration of Cloud Computing, Cyber Physical Systems and Internet of Things*, Springer, 2019.
- [3] B. Varghese, N. Wang, S. Barbhuiya, P. Kilpatrick, D. S. Nikolopoulos, Challenges and Opportunities in Edge Computing, in: *IEEE Int. Conf. Smart Cloud*, 2016, pp. 20–26.
- [4] V. Cardellini, G. Mencagli, D. Talia, M. Torquati, New Landscapes of the Data Stream Processing in the Era of Fog Computing, *Future Generation Computer Systems* 99 (2019) 646–650.
- [5] B. Varghese, P. Leitner, S. Ray, K. Chard, A. Barker, Y. Elkhatib, H. Herry, C.-H. Hong, J. Singer, F. Tso, E. Yoneki, M.-F. Zhani, Cloud futurology, *Computer* 52 (9) (2019) 68–77. doi:10.1109/MC.2019.2895307.
- [6] T. Taleb, S. Dutta, A. Ksentini, M. Iqbal, H. Flinck, Mobile Edge Computing Potential in Making Cities Smarter, *IEEE Commun. Mag.* 55 (3) (2017) 38–43.
- [7] N. Wang, B. Varghese, M. Matthaiou, D. S. Nikolopoulos, ENORM: A Framework For Edge NNode Resource Management, *IEEE Transactions on Services Computing* (2017) 1–1doi:10.1109/TSC.2017.2753775.
- [8] T. Taleb, K. Samdanis, B. Mada, H. Flinck, S. Dutta, D. Sabella, On Multi-access Edge Computing: A Survey of the Emerging 5G Network Edge Cloud Architecture and Orchestration, *Commun. Surveys Tuts.* 19 (3) (2017) 1657–1681.
- [9] P. Liu, D. Willis, S. Banerjee, ParaDrop: Enabling Lightweight Multi-tenancy at the Network’s Extreme Edge, in: *IEEE/ACM Symp. Edge Computing*, 2016, pp. 1–13.
- [10] B. Varghese, M. Villari, O. Rana, P. James, T. Shah, M. Fazio, R. Ranjan, Realizing Edge Marketplaces: Challenges and Opportunities, *IEEE Cloud Computing* 5 (6) (2018) 9–20.
- [11] J. Zhu, D. Chan, M. M. S. Prabhu, P. Natarajan, H. Hu, Improving Web Sites Performance using Edge Servers in Fog Computing Architecture, *uS Patent App.* 13/904,327 (May 15 2014).

- [12] J. McChesney, N. Wang, A. Tanwer, E. de Lara, B. Varghese, Defog: Fog computing benchmarks, in: SEC 2019 : The Fourth ACM/IEEE Symposium on Edge Computing, IEEE, 2019.
- [13] H. Liu, B. He, Reciprocal Resource Fairness: Towards Cooperative Multiple-resource Fair Sharing in IaaS Clouds, in: Int. Conf. High Perf. Comp. Networking, Storage and Analysis, 2014, pp. 970–981.
- [14] J. B. Bernabe, J. M. M. Perez, J. M. A. Calero, F. J. G. Clemente, G. M. Perez, A. F. G. Skarmeta, Semantic-aware Multi-tenancy Authorization System for Cloud Architectures, *Future Generation Computer Systems* 32 (2014) 154–167.
- [15] A. Ahmed, E. Ahmed, A Survey on Mobile Edge Computing, in: Int. Conf. Intelligent Systems and Control, 2016, pp. 1–8.
- [16] M. Satyanarayanan, The Emergence of Edge Computing, *Computer* 50 (1) (2017) 30–39.
- [17] S. Sotiriadis, N. Bessis, C. Amza, R. Buyya, Vertical and Horizontal Elasticity for Dynamic Virtual Machine Reconfiguration, *IEEE Transactions on Services Computing* (99) (2016) 1–1.
- [18] J. Lango, Toward Software-defined SLAs, *Communications of the ACM* 57 (1) (2014) 54–60.
- [19] D. Liu, X. Sui, L. Li, Z. Jiang, H. Wang, Z. Zhang, Y. Zeng, A Cloud Service Adaptive Framework Based on Reliable Resource Allocation, *Future Generation Computer Systems* 89 (2018) 455–463.
- [20] A. Sangpetch, O. Sangpetch, N. Juangmarisakul, S. Warodom, Thoth: Automatic Resource Management with Machine Learning for Container-based Cloud Platform, in: Int. Conf. Cloud Computing and Services Science, 2017, pp. 103–111.
- [21] A. A. Rahmanian, M. Ghobaei-Arani, S. Tofighy, A Learning Automata-based Ensemble Resource Usage Prediction Algorithm for Cloud Computing Environment, *Future Generation Computer Systems* 79 (2018) 54–71.
- [22] R. Han, L. Guo, M. M. Ghanem, Y. Guo, Lightweight Resource Scaling for Cloud Applications, in: 12th IEEE/ACM Int. Symp. Cluster, Cloud and Grid Computing, 2012, pp. 644–651.
- [23] D. Çavdar, R. Birke, L. Y. Chen, F. Alagöz, A Simulation Framework for Priority Scheduling on Heterogeneous Clusters, *Future Generation Computer Systems* 52 (2015) 37–48.

- [24] C.-H. Hong, B. Varghese, Resource Management in Fog/Edge Computing: A Survey on Architectures, Infrastructure, and Algorithms, *ACM Computing Surveys* 52 (5). doi: 10.1145/3326066.
- [25] Y. Mao, J. Zhang, K. B. Letaief, Dynamic Computation Offloading for Mobile-Edge Computing with Energy Harvesting Devices, *IEEE Journal on Selected Areas in Communications* 34 (12) (2016) 3590–3605.
- [26] Y.-J. Ku, D.-Y. Lin, C.-F. Lee, P.-J. Hsieh, H.-Y. Wei, C.-T. Chou, A.-C. Pang, 5G Radio Access Network Design with the Fog Paradigm: Confluence of Communications and Computing, *IEEE Communications Magazine* 55 (4) (2017) 46–52.
- [27] Y. Yu, J. Zhang, K. B. Letaief, Joint Subcarrier and CPU Time Allocation for Mobile Edge Computing, in: *IEEE Int. Conf. Global Communications*, 2016, pp. 1–6.
- [28] Mobile Edge Computing (MEC); Framework and reference architecture, Standard, European Telecommunications Standards Institute (Mar. 2016).
URL http://www.etsi.org/deliver/etsi_gs/MEC/001_099/003/01.01.01_60/gs_MEC003v010101p.pdf
- [29] C. Pahl, B. Lee, Containers and Clusters for Edge Cloud Architectures – A Technology Review, in: *3rd Int. Conf. Future Internet of Things and Cloud*, 2015, pp. 379–386.
- [30] M. Al-Roomi, S. Al-Ebrahim, S. Buqrais, I. Ahmad, Cloud Computing Pricing Models: A Survey, *International Journal of Grid and Distributed Computing* 6 (5) (2013) 93–106.
- [31] H. Chen, F. Wang, N. Helian, G. Akanmu, User-priority Guided Min-Min Scheduling Algorithm for Load Balancing in Cloud Computing, in: *2013 National Conference on Parallel computing technologies (PARCOMPTECH)*, IEEE, 2013, pp. 1–8.
- [32] S. He, L. Guo, Y. Guo, C. Wu, M. Ghanem, R. Han, Elastic Application Container: A Lightweight Approach for Cloud Resource Provisioning, in: *IEEE Int. Conf. Advanced Information Networking and Applications*, IEEE, 2012, pp. 15–22.
- [33] M. Pastorelli, D. Carra, M. Dell’Amico, P. Michiardi, HFSP: Bringing Size-based Scheduling to Hadoop, *IEEE Transactions on Cloud Computing* 5 (1) (2015) 43–56.

- [34] L. Gao, T. H. Luan, S. Yu, W. Zhou, B. Liu, FogRoute: DTN-based Data Dissemination Model in Fog Computing, *IEEE Internet of Things Journal* 4 (1) (2016) 225–235.
- [35] C. You, K. Huang, H. Chae, B.-H. Kim, Energy-efficient Resource Allocation for Mobile-Edge Computation Offloading, *IEEE Trans. Wireless Commun.* 16 (3) (2017) 1397–1411.
- [36] M. Tortonesi, M. Govoni, A. Morelli, G. Riberto, C. Stefanelli, N. Suri, Taming the IoT Data Deluge: An Innovative Information-centric Service Model for Fog Computing Applications, *Future Generation Computer Systems* 93 (2019) 888–902.
- [37] R. Buyya, D. Abramson, J. Giddy, H. Stockinger, Economic Models for Resource Management and Scheduling in Grid Computing, *Concurrency and computation: practice and experience* 14 (13-15) (2002) 1507–1542.
- [38] A. Brogi, S. Forti, QoS-aware Deployment of IoT Applications Through the Fog, *IEEE Internet of Things Journal* 4 (5) (2017) 1185–1192.
- [39] T. C. Chieu, A. Mohindra, A. A. Karve, A. Segal, Dynamic Scaling of Web Applications in a Virtualized Cloud Computing Environment, in: *2009 IEEE International Conference on e-Business Engineering*, IEEE, 2009, pp. 281–286.
- [40] Z. Shen, S. Subbiah, X. Gu, J. Wilkes, CloudScale: Elastic Resource Scaling for Multi-tenant Cloud Systems, in: *ACM Symp. Cloud Computing*, 2011, p. 5.
- [41] G. J. L. Paulraj, S. A. J. Francis, J. D. Peter, I. J. Jebadurai, Resource-aware Virtual Machine Migration in IoT Cloud, *Future Generation Computer Systems* 85 (2018) 173–183.
- [42] L. Yang, J. Cao, H. Cheng, Y. Ji, Multi-user Computation Partitioning for Latency Sensitive Mobile Cloud Applications, *IEEE Transactions on Computers* 64 (8) (2015) 2253–2266.
- [43] X. Lyu, H. Tian, C. Sengul, P. Zhang, Multiuser Joint task Offloading and Resource Optimization in Proximate Clouds, *IEEE Trans. Veh. Technol.* 66 (4) (2017) 3435–3447.
- [44] L. Yazdanov, C. Fetzer, Lightweight Automatic Resource Scaling for Multi-tier Web Applications, in: *IEEE Int. Conf. Cloud Computing*, 2014, pp. 466–473.

- [45] M. Mao, M. Humphrey, Auto-scaling to Minimize Cost and Meet Application Deadlines in Cloud Workflows, in: *Int. Conf. High Perf. Comp., Networking, Storage and Analysis*, 2011, pp. 1–12.
- [46] F. Zhang, X. Tang, X. Li, S. U. Khan, Z. Li, Quantifying Cloud Elasticity with Container-based Autoscaling, *Future Generation Computer Systems* 98 (2019) 672–681.
- [47] A. Abdul Majeed, P. Kilpatrick, I. Spence, B. Varghese, Performance Estimation of Container-Based Cloud-to-Fog Offloading, in: *Proceedings of the 12th IEEE/ACM International Conference on Utility and Cloud Computing Companion*, 2019, pp. 151–156.
- [48] P. Marshall, K. Keahey, T. Freeman, Elastic Site: Using Clouds to Elastically Extend Site Resources, in: *IEEE/ACM Int. Conf. Cluster, Cloud and Grid Computing*, 2010, pp. 43–52.
- [49] J. Jiang, J. Lu, G. Zhang, G. Long, Optimal Cloud Resource Auto-scaling for Web Applications, in: *IEEE/ACM Int. Symp. Cluster, Cloud and Grid Computing*, 2013, pp. 58–65.
- [50] L. Thai, B. Varghese, A. Barker, Budget Constrained Execution of Multiple Bag-of-Tasks Applications on the Cloud, in: *IEEE Int. Conf. Cloud Computing*, 2015, pp. 975–980.
- [51] A. Ali-Eldin, J. Tordsson, E. Elmroth, An Adaptive Hybrid Elasticity Controller for Cloud Infrastructures, in: *IEEE Network Operations and Management Symp.*, 2012, pp. 204–212.
- [52] T. Li, J. Wang, W. Li, T. Xu, Q. Qi, Load Prediction-based Automatic Scaling Cloud Computing, in: *2016 International Conference on Networking and Network Applications (NaNA)*, IEEE, 2016, pp. 330–335.
- [53] S. Spinner, S. Kounev, X. Zhu, L. Lu, M. Uysal, A. Holler, R. Griffith, Runtime Vertical Scaling of Virtualized Applications via Online Model Estimation, in: *2014 IEEE Eighth International Conference on Self-Adaptive and Self-Organizing Systems*, IEEE, 2014, pp. 157–166.
- [54] S. Dutta, S. Gera, A. Verma, B. Viswanathan, Smartscale: Automatic Application Scaling in Enterprise Clouds, in: *2012 IEEE Fifth International Conference on Cloud Computing*, IEEE, 2012, pp. 221–228.
- [55] Y. Mao, C. You, J. Zhang, K. Huang, K. B. Letaief, A Survey on Mobile Edge Computing: The Communication Perspective, *Commun. Surveys Tuts.* 19 (4) (2017) 2322–2358.

- 843 [56] X. Chen, L. Jiao, W. Li, X. Fu, Efficient Multi-user Computation Offloading for Mobile-Edge
844 Cloud Computing, *IEEE/ACM Trans. Netw.* (5) (2016) 2795–2808.