**QUEEN'S
UNIVERSITY
BELFAST**

Modelling Fog Offloading Performance

Majeed, A. A., Kilpatrick, P., Spence, I., & Varghese, B. (2020). Modelling Fog Offloading Performance. In Y. Simmhan, & B. Varghese (Eds.), *Proceedings of the 4th IEEE International Conference on Fog and Edge Computing* (pp. 29-38). [9139185] (Proceedings - 4th IEEE International Conference on Fog and Edge Computing, ICFEC 2020). Institute of Electrical and Electronics Engineers Inc..
<https://doi.org/10.1109/ICFEC50348.2020.00011>

Published in:

Proceedings of the 4th IEEE International Conference on Fog and Edge Computing

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

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

Publisher rights

© 2020 IEEE.

This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

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.

Modelling Fog Offloading Performance

Ayesha Abdul Majeed, Peter Kilpatrick, Ivor Spence, and Blesson Varghese

School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, UK

E-mail: {aabdulmajeed01, p.kilpatrick, i.spence, b.varghese}@qub.ac.uk

Abstract—Fog computing has emerged as a computing paradigm aimed at addressing the issues of latency, bandwidth and privacy when mobile devices are communicating with remote cloud services. The concept is to offload compute services closer to the data. However many challenges exist in the realisation of this approach. During offloading, (part of) the application underpinned by the services may be unavailable, which the user will experience as down time. This paper describes work aimed at building models to allow prediction of such down time based on metrics (operational data) of the underlying and surrounding infrastructure. Such prediction would be invaluable in the context of automated Fog offloading and adaptive decision making in Fog orchestration. Models that cater for four container-based stateless and stateful offload techniques, namely Save and Load, Export and Import, Push and Pull and Live Migration, are built using four (linear and non-linear) regression techniques. Experimental results comprising over 42 million data points from multiple lab-based Fog infrastructure are presented. The results highlight that reasonably accurate predictions (measured by the coefficient of determination for regression models, mean absolute percentage error, and mean absolute error) may be obtained when considering 25 metrics relevant to the infrastructure.

Index Terms—Fog computing, offloading, edge computing, containers, performance estimation

I. INTRODUCTION

The limited resources of mobile devices, both power and computational, led to the Cloud-only concept of offloading services to a Cloud data centre. A desire to reduce the latency introduced by such offloading, together with an increase in the availability of resources (for example micro data centres or compute enabled routers) at the edge of the network, has resulted in the more flexible Fog computing approach. As illustrated in Figure 1, services can be offloaded: (i) from the Cloud to the Fog, to minimise the latency of communicating with the end-user device [1], (ii) from Fog to the Cloud, to exploit additional computational and storage resources [2], and (iii) from end-user devices to the Fog, to satisfy compute requirements unavailable on the devices [3]. We consider Cloud-to-Fog and Fog-to-Cloud offloading.

The benefits derived from offloading [4] can be negated if the very process of offloading results in a temporary loss of service (down time). We aim to predict the times that would be taken to offload a given service by a variety of known methods, which are estimates of the corresponding induced down times. Such predictions would be helpful in the following situations:

(i) *Automated Fog software development environments*: predictive models would enable a Fog software developer to assess the feasibility of offloading without establishing corresponding testbeds and conducting empirical investigations.

Fig. 1: Potential Fog offload scenarios, namely (i) Cloud-to-Fog, (ii) Fog-to-Cloud, and (iii) Device-to-Fog

(ii) *Adaptive decision-making in Fog orchestration platforms*: the transient nature of Fog environments means that the best selection of services, destinations and techniques for offloading may well be constantly changing. Predictive models would permit a continual reassessment of offloading decisions.

(iii) *Simulation platforms*: in the absence of corresponding computational and network infrastructure, researchers and practitioners can use simulation models to predict the performance of a particular configuration. Predictive models of offload times would make the models more complete in that down times as well as runtime performance could be modelled.

The paper aims to address the following *three research questions* - **Q1**: How can different offloading approaches be modelled? **Q2**: What runtime (system and network related) and offline (do not change during offload process) parameters influence offloading? **Q3**: Given the same empirical data of offloading from different experimental platforms how is the accuracy of estimation affected when using different estimation methods and machine learning algorithms?

In this investigation, we consider three ‘Stateless’ techniques (Save and Load, Export and Import, and Push and Pull) and one ‘Stateful’ technique (Checkpoint/Restore in Userspace), which are container-based techniques adopted in this paper for offloading from the Cloud to the Fog and vice-

versa. As a first step, the anatomy of each of the offloading approaches is considered by studying the individual components that make up the overall time required to offload (addresses Q1). The research establishes a method for estimating the time taken to offload a service which will also be applicable for different offloading techniques. 25 relevant runtime and offline metrics for the entire Cloud and Fog system have been identified to model the offloading time (addresses Q2). Four machine learning based predictive models (one linear and three non-linear) are applied to over 42 million data points obtained from two experimental platforms (to address Q3). The results indicate that the chosen runtime and offline metrics are relevant to predicting the offload time which can be estimated with reasonable accuracy.

This paper makes the following *three research contributions*: (i) Develops a method for modelling the offloading time for multiple container-based stateless and stateful techniques, (ii) Identifies 25 runtime and offline parameters that influence the model for offloading, and (iii) Presents an empirical investigation of the accuracy of multiple linear and non-linear machine learning algorithms for predicting the offloading time.

The remainder of this paper is organised as follows. Section II discusses stateless and stateful techniques for offloading. Section III proposes two methods for estimating the performance of offloading and highlights the models for estimating performance. Section IV presents the experimental studies pursued on two different platforms. Section V highlights research relevant to the discussion of this paper. Section VI concludes this paper by considering future work.

II. FOG OFFLOADING

This section presents three offloading scenarios highlighted in Figure 1 and four container-based offloading techniques.

A. Offloading Scenarios

Fog offloading can be considered in the following three scenarios, which are illustrated in Figure 1:

(i) *Cloud-to-Fog offloading* [5], [6] - this refers to transferring services of an application from a Cloud resource on to a Fog resource to meet latency and ingress bandwidth demands. Since the offloaded service is closer to the device that generates data, it reduces the communication latency and (pre)-processes data closer to the source, thereby reducing the volume of data that needs to be transferred to the cloud.

(ii) *Fog-to-Cloud offloading* [5], [7] - this refers to transferring the services that are resident on a Fog resource to the Cloud. This may be done in response to a change in the life-cycle of the application service - the service on the Fog resource needs to be terminated and resumed on the Cloud from where it was originally offloaded. This may be because the service requires additional resources (CPU cores, storage, memory) that are not available on the Fog resource, but can only be satisfied on the Cloud.

(iii) *Device-to-Fog offloading* [3], [8] - this refers to transferring services of an application from one or a collection of end user devices or sensors to a Fog resource. This is done in order

to preserve battery life of devices or to meet the computational demands of workloads that cannot be executed on user devices due to limited form factor and weak processing capabilities. The offloaded service executes on the Fog resource and the resulting output is provided back to the devices or sensors.

The focus of this paper is the Cloud-to-Fog and Fog-to-Cloud offloading scenarios.

B. Container-Based Offloading Approaches

There are two dominant approaches that facilitate offloading. The first is Virtual Machine (VM) migration. VM hand-off is used to move a service for supporting user mobility in the context of cloudlets [9]. A synthesis technique is adopted in which the VM is divided into two stacked overlays to optimise the downtime during VM hand-off.

A second approach to facilitate offloading is container migration. Containers are an alternative virtualisation approach that are lightweight and portable, thereby making offloading quicker than using VMs [10]. Therefore, containers are investigated within Fog computing, in which Fog resources have limited resources when compared to the Cloud [2], [5]. Popular container technologies, include LXC¹ and Docker² and they support migration, which is required for offloading.

Fog offloading, in this paper, is explored in the context of Docker containers. Docker packages an application as an image that consist of a file system with the required libraries, executables and configuration files. In practice, the image may comprise a series of layers that are stacked on top of a base image, for example, the Ubuntu operating system. When a container is executed Docker mounts all the layers of the image as 'read-only' using the Union File System (UnionFS) and the top layer as a writable layer as shown in Figure 2.

Docker supports two migration approaches that can be used for offloading: (i) *Stateless* - the state of the application is not transferred, instead a separate instance of the container is run elsewhere without its previous state (for example, Save and Load, Export and Import, and Push and Pull), and (ii) *Stateful* - the state of the running application is transferred with the container image. Stateful migration can be achieved by using the CRIU (Checkpoint/Restore in Userspace)³ approach.

1) *Stateless Techniques*: Three stateless techniques are considered in this paper for offloading. They are as follows:

(i) *Save and Load*: Figure 2 shows the Cloud-to-Fog offloading scenario using the Save and Load technique. The goal of this technique is to transfer the image of a running container and the underlying base layers. This offloading technique is a five step process as illustrated below:

Step 1 - Commit: A container is instantiated when it boots up from a series of base image layers when there is a request for offloading from the Cloud server. The commit operation stops the running container and saves the current state of the container, the accompanying stacked layers, and the modifications that were made within the container as a

¹<https://linuxcontainers.org/>

²<https://www.docker.com/>

³https://criu.org/Main_Page

new image, referred to as the committed image. This newly created image is stored in the local image registry of the Cloud server. The time taken for this step (store the configuration and run-time state of the container, create an image, and store the image in the local registry) is denoted as t_{commit} .

Step 2 - Save: The save operation, converts the committed image to a compressed `.tar` file and saves it on the hard disk of the Cloud server. The compressed file contains information on the contents of the container, including its parent layers and the size of each layer. The time taken to convert the committed image to the compressed file is denoted as t_{save} .

Step 3 - Transfer: The compressed image file is transferred to the Fog resource using a network transfer protocol, such as FTP, SCP or rsync; $t_{transfer}$ captures the time taken to transfer the compressed file to the Fog server.

Step 4 - Load: This operation initially decompresses the `.tar` image file and loads it from the hard disk on the Fog resource as an image into the local image registry. This time is captured as t_{load} .

Step 5 - Start: A container from the image in the registry is booted up on the Fog resource; t_{start} captures this time.

Based on the above five steps, the total time taken to offload a container-based service from the Cloud to the Fog or vice-versa is represented as:

$$t_{offload} = t_{commit} + t_{save} + t_{transfer} + t_{load} + t_{start} \quad (1)$$

In the Cloud-to-Fog offload scenario, t_{commit} and t_{save} are the operations on the Cloud server, and t_{load} and t_{start} are on the Fog resource. The transfer time will be for transferring the compressed file from the Cloud to the Fog. In the Fog-to-Cloud offload scenario, the commit and save times will be for the operation on the Fog resource, and the load and start for the operations on the Cloud server.

In this work it is assumed that Fog applications are designed and developed as micro-services (a collection of services can be orchestrated as a workflow and each service can be deployed as a container). The advantage of this offloading technique is that the base image layers are also transferred from the source (for example, the Cloud server in the Cloud-to-Fog offload scenario) to the destination Fog resource allowing for creation of further copies of the container. However, there will be a trade-off with the size of the image transferred.

A potential advantage is when multiple containers that rely on the same underlying libraries need to be offloaded. In this case, the base image layers will not be duplicated for multiple container services when it is offloaded. In addition, the history of the image is preserved (configuration settings, ports and entry points), which is an advantage, but incurs a large performance overhead during container deployment.

(ii) **Export and Import:** Figure 3 shows the Cloud-to-Fog offloading scenario using the Export and Import technique. The goal is to export the running state as a flattened (multiple image layers into one single layer) container. This offloading technique is a four step method as illustrated below:

Step 1 - Export: The export operation stops the running container and saves the current state of the container and the

Fig. 2: Five steps in offloading a container from the Cloud to the Fog using the Save and Load technique

Fig. 3: Four steps in offloading a container from the Cloud to the Fog using the Export and Import technique

modifications that were made within the container, referred to as the flattened container. The flattened container is stored on hard disk of the Cloud server as a `.tar` file. The time taken for this step (store the run-time state of the container, and store the container as a tar file) is denoted as t_{export} .

Step 2 - Transfer: The compressed container file is transferred to the Fog resource using SCP network transfer protocol. $t_{transfer}$ time captures the time taken to transfer the compressed file to the Fog server.

Step 3 - Import: This operation initially decompresses the `.tar` file and imports it from the hard disk on the Fog resource as an image into local image registry. This time is captured as t_{import} .

Step 4 - Start: A container is booted up on the Fog resource from the flattened container in the registry; this time is t_{start} .

The total time taken to offload is represented as:

$$t_{offload} = t_{export} + t_{transfer} + t_{import} + t_{start} \quad (2)$$

Flattening the container removes the history (configuration settings, ports and entry point settings) of the container which reduces the container size and improves the efficiency of the deployment process of a container. One disadvantage of this approach is that the flattened container that is exported cannot be used as a template to create a new container.

(iii) *Push and Pull Technique*: Figure 4 shows the Cloud-to-Fog offloading scenario using the Push and Pull technique - the image of a running container is pushed from the Cloud server to a central repository. Then the image is pulled and a container is instantiated on the Fog server.

This technique is a four step method as illustrated below:

Step 1 - Commit: The commit operation saves the current state of the container, and the modifications that were made within the container. The committed image is stored on the local registry of the Cloud server. The time taken for this step (store the configuration and run-time state of the container, and store the container in the local registry) is denoted as t_{commit} .

Step 2 - Push: The committed image is pushed into the Docker Hub using push command, t_{push} captures the time taken to push the committed image into the Docker hub.

Step 3 - Pull: This operation pulls the image from the Docker hub and loads the image into the local image registry of the Fog server. This time is captured as t_{pull} .

Step 4 - Start: A container from the image in the registry is booted up on the Fog resource; t_{start} captures this time.

The total time taken to offload a container-based service from the Cloud to the Fog or vice-versa is represented as:

$$t_{offload} = t_{commit} + t_{push} + t_{pull} + t_{start} \quad (3)$$

When the available network bandwidth is limited, pushing and pulling an image from the Docker Hub slows container deployment. Docker repository is a data-intensive application and the repository becomes a performance bottleneck, as the number of images and user requests increases, thereby affecting the offloading process.

2) *Stateful Techniques*: A container-based live migration technique is presented in this paper.

Checkpoint/Restore in Userspace (Live Migration): Figure 5 shows the Cloud-to-Fog offloading scenario using the live migration technique. Docker uses Checkpoint/Restore in Userspace (CRIU) to freeze a running container by checkpointing it and then restoring it. The process is as follows:

Step 1 - Checkpoint: This procedure freezes the container process, collects and saves the complete state of the container and then stops the container ($t_{checkpoint}$ captures this time). The container state file (dump files) is shared (or transferred) to the Fog server through shared Network File System (NFS).

Step 2 - Restore: a container is restored from the dump files on the mounted directory, process execution is resumed (this time is $t_{restore}$). Container is initialised after it is restored on the Fog server. The offload time is represented as:

$$t_{offload} = t_{checkpoint} + t_{restore} \quad (4)$$

Stateful techniques are beneficial for applications in which a client is serviced by an application that must retain the internal

Fig. 4: Four steps in offloading a container from the Cloud to the Fog using the Push and Pull technique

Fig. 5: Two steps in offloading a container from the Cloud to the Fog using the Stateful offloading technique

state of the application. One disadvantage is that services cannot be further replicated on a different Fog node since data relevant to individual users is stored in the application.

III. ESTIMATION METHODS AND MODELS

In this section, the parameters that influence the stateless and stateful offloading techniques, the methods used for estimation, and the machine learning algorithms used for the estimation model that predicts $t_{offload}$ are presented.

A. Methods for Estimation

Table I describes a catalogue of parameters that are considered in this work by the estimation models while offloading a container. Two types of parameters are considered, namely runtime and offline parameters. The *runtime* parameters are collected when the steps of the migration techniques are executed to offload a container. These parameters relate to both the offloading process and the entire Cloud and/or Fog system as highlighted in the table. The network properties between

TABLE I: Parameters that impact overall offloading time; Bps - Bytes per second, bps - bits per second, BW - bandwidth

Parameter	Description	System/ Process	Cloud/ Fog
Runtime			
P_1, P_9	CPU utilisation (%)	System	$P_1 - P_3$ (Cloud) $P_9 - P_{11}$ (Fog)
P_2, P_{10}	Memory utilisation (%)		
P_3, P_{11}	Disk utilisation (%)		
P_4, P_{12}	CPU utilisation (%)	Offloading Process	$P_4 - P_8$ (Cloud) $P_{12} - P_{16}$ (Fog)
P_5, P_{13}	Memory utilisation (%)		
P_6, P_{14}	Disk throughput (Bps)		
P_7, P_{15}	Bytes sent (KB/sec)		
P_8, P_{16}	Bytes received (KB/sec)		
Offline			
P_{17}	Image size (MB)	Offloaded container	
P_{18}, P_{21}	No of Cores	$P_{18} - P_{20}$ (Cloud) $P_{21} - P_{23}$ (Fog)	
P_{19}, P_{22}	Memory Size (GB)		
P_{20}, P_{23}	Hard disk Size (GB)		
P_{24}	Network BW (bps)	Between Cloud and Fog	
P_{25}	Network Latency (ms)		

TABLE II: Parameters that impact the offloading sequence of the Save and Load Technique

Time	Parameters that impact
t_{commit}	$X_{commit} = \{ P_1, \dots, P_6, P_{17}, \dots, P_{20} \}$
t_{save}	$X_{save} = \{ P_1, \dots, P_6, P_{17}, \dots, P_{20} \}$
$t_{transfer}$	$X_{transfer} = \{ P_7, P_8, P_{15}, P_{16}, P_{17}, P_{24}, P_{25} \}$
t_{load}	$X_{load} = \{ P_9, \dots, P_{14}, P_{17}, P_{21}, \dots, P_{23} \}$
t_{start}	$X_{start} = \{ P_9, \dots, P_{14}, P_{17}, P_{21}, \dots, P_{23} \}$

the Cloud and the Fog are considered to capture the state of the network when the offload occurs. The *offline* parameters are statically determined and do not change during the offload process; for example number of cores, network bandwidth and container image size.

Two methods are used to estimate the offload time - the first uses a single estimation model and the second uses multiple estimation models for the individual time components shown in Equation 1-4.

The first method is *using a collective model*, referred to as (CM), which is a reference to the use of a single model that estimates the offload time. The collective model uses all the parameters listed in Table I as input. Consider a collective model, $M_{collective}$ that estimates the offload time and let $X_{offload} = \{P_1, \dots, P_{25}\}$ be the input to the model, then we represent $t_{offload} = M_{collective}(X_{offload})$.

The second method is *using individual models*, which refers to the use of separate models for estimating the individual times of Equations 1-4.

Table II shows the parameters that affect the individual times of the Save and Load technique. Let M_{commit} be an individual model to estimate t_{commit} using the input X_{commit} shown in Table II. Then the estimation of $t_{commit} = M_{commit}(X_{commit})$. Similarly, $t_{save} = M_{save}(X_{save})$, $t_{transfer} = M_{transfer}(X_{transfer})$, $t_{load} = M_{load}(X_{load})$, and $t_{start} = M_{start}(X_{start})$. The offload time can be estimated as

TABLE III: Parameters that impact the offloading sequence of the Export and Import Technique

Time	Parameters that impact
t_{export}	$X_{export} = \{ P_1, \dots, P_6, P_{17}, \dots, P_{20} \}$
$t_{transfer}$	$X_{transfer} = \{ P_7, P_8, P_{15}, P_{16}, P_{17}, P_{24}, P_{25} \}$
t_{import}	$X_{import} = \{ P_9, \dots, P_{14}, P_{17}, P_{21}, \dots, P_{23} \}$
t_{start}	$X_{start} = \{ P_9, \dots, P_{14}, P_{17}, P_{21}, \dots, P_{23} \}$

TABLE IV: Parameters that impact the offloading sequence of the Push and Pull Technique

Time	Parameters that impact
t_{commit}	$X_{commit} = \{ P_1, \dots, P_6, P_{17}, \dots, P_{20} \}$
t_{push}	$X_{pull} = \{ P_1, \dots, P_8, P_{17}, \dots, P_{20}, P_{24}, P_{25} \}$
t_{pull}	$X_{push} = \{ P_9, \dots, P_{14}, P_{17}, P_{21}, \dots, P_{23} \}$
t_{start}	$X_{start} = \{ P_9, \dots, P_{14}, P_{17} \}$

shown in Equation 5.

$$t_{offload} = M_{commit}(X_{commit}) + M_{save}(X_{save}) + M_{transfer}(X_{transfer}) + M_{load}(X_{load}) + M_{start}(X_{start}) \quad (5)$$

Table III shows the parameters that affect the individual times of the Export and Import technique. Let M_{export} be an individual model to estimate t_{export} using the input X_{export} shown in Table III. Then the estimation of $t_{export} = M_{export}(X_{export})$. Similarly, $t_{transfer} = M_{transfer}(X_{transfer})$, $t_{import} = M_{import}(X_{import})$, and $t_{start} = M_{start}(X_{start})$. The offload time can be estimated as shown in Equation 6.

$$t_{offload} = M_{export}(X_{export}) + M_{transfer}(X_{transfer}) + M_{import}(X_{import}) + M_{start}(X_{start}) \quad (6)$$

Table IV shows the parameters that affect the individual times of the Push and Pull technique. Let M_{commit} be an individual model to estimate t_{commit} using the input X_{commit} shown in Table IV. Then the estimation of $t_{commit} = M_{commit}(X_{commit})$. Similarly, $t_{pull} = M_{pull}(X_{pull})$, $t_{push} = M_{push}(X_{push})$, and $t_{start} = M_{start}(X_{start})$. The offload time can be estimated as shown in Equation 7.

$$t_{offload} = M_{commit}(X_{commit}) + M_{pull}(X_{pull}) + M_{push}(X_{push}) + M_{start}(X_{start}) \quad (7)$$

Table V shows the parameters that affect the individual times of the stateful technique. Let $M_{checkpoint}$ be an individual model to estimate $t_{checkpoint}$ using the input $X_{checkpoint}$ shown in Table V. Then the estimation of $t_{checkpoint} = M_{checkpoint}(X_{checkpoint})$. Similarly, $t_{restore} = M_{restore}(X_{restore})$. The offload time can be estimated as shown in Equation 8.

$$t_{offload} = M_{checkpoint}(X_{checkpoint}) + M_{restore}(X_{restore}) \quad (8)$$

B. Models for Estimation

Four machine learning algorithms (one linear and three non-linear) were explored for predicting the offload time.

TABLE V: Parameters that impact the offloading sequence of the Stateful Technique

Time	Parameters that impact
$t_{checkpoint}$	$X_{checkpoint} = \{ P_1, \dots, P_8, P_{17}, \dots, P_{20} \}$
$t_{restore}$	$X_{restore} = \{ P_9, \dots, P_{16}, P_{21}, \dots, P_{23} \}$

The approach used for estimation is based on historical data that is collected from the experimental platform (presented in Section IV) to predict $t_{offload}$. The algorithms used are:

(i) Multivariate Linear Regression (MLR): The model developed using this algorithm captures the relationship between multiple input variables $X = \{P_1, P_2, \dots, P_n\}$ and the dependent output variable $t_{offload}$ by a straight line equation [11].

(ii) Polynomial Multivariate Regression (PMR): The regression model developed captures the relationship between the input variables $X = \{P_1, P_2, \dots, P_n\}$ and the dependent output variable $t_{offload}$ as an n^{th} degree polynomial in X .

(iii) Random Forest Regression (RFR): An ensemble model generates k different training subsets from the original data set, and then k different decision trees are built based on the generated training subsets. Each sample of the testing data set is predicted by all decision trees, and the final result is obtained by averaging a score specific to each decision tree [11].

(iv) Ridge regression (RR): A non-linear approach that adds an penalty ($L2$), which equals the sum of the squared value of the coefficients $Error_{L2} = Error + \sum_{i=0}^N \lambda \cdot W_i^2$ [12].

IV. EXPERIMENTAL STUDIES

This section presents the experimental setup and the results obtained from running experiments for the Stateful and Stateless container based offloading approaches.

A. Experimental setup

The proposed methods, namely the Collective Model (CM) and the Individual Models (IM) for estimating the performance of offloading using four estimation models, namely MLR, PMR, RFR and RR are evaluated on two different lab-based platforms (each platform is a combination of a Cloud and Fog VM that executes the container based offloading techniques). Docker 18.09-ce is installed on all VMs. Both experimental platforms use 64-bit x86 architectures for the Cloud and Fog environment. A number of recent Fog-enabled nodes, such as the Dell Edge Gateway 5000, use 64-bit x86 processors [13].

The first platform is the combination of a Cloud VM running Ubuntu 18.10 with 6 virtual CPUs, 30GB hard disk and 6GB RAM and Fog VM running Ubuntu 18.10 with 2 virtual CPUs, 20GB hard disk and 2GB RAM.

The network bandwidth between the Cloud and Fog VMs are emulated using Linux Traffic Control (tc)⁴. The bandwidth is varied as 25Mbps, 50Mbps, 100Mbps, 1000Mbps with a latency of 10ms and 30ms (values are based on the literature [14]).

The second platform is another combination of a Cloud VM and Fog VM running on OpenStack. Both VMs run Ubuntu

18.10; the Cloud VM has 4 virtual CPUs, 80GB hard disk space and 8GB virtual RAM whereas the Fog VM has 2 virtual CPUs, a 40GB hard disk, and 4GB virtual RAM. The default network connection is approximately 3.2Mbps.

1) *Stateless Techniques*: The following is considered for the set up of the stateless techniques. For the Cloud VM in the *Save and Load* technique, the container state is saved using the Docker commit process and the save process saves the image as a .tar file. The tar file is transferred to the Fog server using the scp protocol. When the tar file is received on the Fog VM, *inotify* loads the image into the local registry. A container is started from the loaded image. During the offloading process, the values of the run-time parameters are collected at one-second intervals. System CPU utilisation is captured by P_1 and P_9 , which are obtained by monitoring `/proc/stat`. CPU utilisation of the offloading process captured by P_4 and P_{12} are obtained by monitoring `/proc/PID/`. RAM utilisation of the system, denoted by P_2 and P_{10} are obtained by the Linux utility tool `ps`. RAM utilisation of the offloading process is monitored using `/proc/PID/smmaps` file for the parameters P_5 and P_{13} . Disk utilisation of the system, denoted by P_3 and P_{11} are obtained using `iostat` utility. Disk throughput of the offloading process is obtained by monitoring `/proc/PID/io` to record the number of bytes written to and read from disk for parameters P_6 and P_{12} . Network utilisation of the offloading process is obtained using the *nethogs* tool for parameters $P_7 - P_8$ and $P_{15} - P_{16}$. The values for offline parameters (i) P_{17} is the size of the offloaded container, (ii) $P_{18} - P_{23}$, are obtained from the settings defined for the Cloud and Fog VMs, and (iii) P_{24} and P_{25} are acquired during network configurations using `tc`.

To simulate the varying availability of CPU, memory and hard disk resources in the experimental environment, the CPU, memory and I/O stress were gradually increased for different experimental runs using `stress-ng`⁵; CPU stress was increased by 10%, memory stress on the Cloud VM by units of 1GB until 75% of capacity and for the Fog VM by units of 512MB until 75% of capacity, and disk stress on the Cloud VM by units of 4GB until 75% of capacity and for the Fog VM by units of 2GB until 75% of capacity.

The data values using the estimation methods (of Section III) are used to build the model for estimating $t_{offload}$.

In the *Export and Import* technique, the container state is saved using the Docker export process as a .tar file on the Cloud VM. The tar file is transferred to the Fog server using the scp protocol. When the tar file is received on the Fog VM, *inotify* loads the image into the local registry from which a container is booted up. During the offloading process, the values of the runtime and offline parameters are collected similar to the Save and Load technique presented above.

In the *Push and Pull* technique, the container state is pushed to a central repository on the Cloud VM using the Docker push process. On the Fog VM, the container state is pulled using the Docker pull process and a container is started from the

⁴<https://linux.die.net/man/8/tc>

⁵<https://manpages.ubuntu.com/manpages/artful/man1/stress-ng.1.html>

loaded image. During offloading, the values of the runtime and offline parameters are collected same as above.

Stateful technique: For *CRIU-based live migration*, the container state is stopped and dumped on the shared directory of the Cloud VM. *inotify* triggers the process of restoring the container state on the Fog VM. During the offloading process, the values of the runtime and offline parameters are collected similar to the discussion above.

The dataset that is generated from all the parameters for the prediction models of the four offloading techniques consists of 5,700 instances across the two experimental platforms for the Cloud-to-Fog and Fog-to-Cloud based offload scenarios for varying combinations of the offline parameters. Since the runtime parameters are collected at one-second intervals each Cloud-to-Fog and Fog-to-Cloud offload scenario for a given set of offload parameters generates a large number of intermediate instances (on an average 60 intermediate instances). The runtime values are averaged to obtain the aggregate instances. A total of $(25 \times 5,700 \times 60 = 8,550,000)$ data points are used. The experiments were repeated five times and a total of 42,750,000 data points were collected.

B. Results

The experimental results for the Cloud-to-Fog and Fog-to-Cloud offloading scenarios are presented in this section. The goal is to demonstrate the feasibility of estimating the time to offload in both scenarios. The results are presented for three metrics, namely (i) R^2 , also known as the coefficient of determination, highlights how much of the observed variations by the regression model can be explained by the model's input. In other words, a higher value indicates a better fit of the regression model to the inputs, (ii) Mean Absolute Percentage Error (MAPE) is the average of percentage errors. A lower value indicates that the model estimates the offload time with a higher accuracy, and (iii) Mean Absolute Error (MAE) is the average error of the model. A lower value indicates that the estimation error of the model is low.

Two validation approaches are employed in each of the offloading scenarios. The first is a train-test validation approach in which 70% of the data is used for training and the remaining 30% of the data is used for testing. This approach is data-driven and requires that operational data from the infrastructure is gathered at a relatively high frequency. The second validation approach is k-fold cross validation, where $k = 10$. The results are presented for each validation approach in both offload scenarios considered in this paper.

1) *Cloud-to-Fog offload:* Results for the train-test validation and k-fold validation approaches are shown in Figure 6 and Figure 7, respectively.

Figure 6a (using the train-test validation approach) and Figure 7a (using the k-fold validation approach) show the degree of influence (R^2) of the chosen runtime and offline parameters (as shown in Table I) on $t_{offload}$ using the four machine learning approaches when employing individual and collective models (IM and CM). Across all machine learning approaches and for both IM and CM, RFR provides the

Fig. 6: Train-test validation approach for Cloud-to-Fog offload

highest R^2 . It is observed that the Push and Pull technique has the lowest R^2 . In this technique, the central repository is accessed for uploading and downloading the container image over the Internet, and the model does not capture the associated uncertainty. Based on R^2 of the estimation models, it can be concluded that the chosen parameters influence $t_{offload}$.

Figure 6b and Figure 7b show the Mean Absolute Error (MAE) in the estimations using multiple validation approaches; lower value indicates a more efficient prediction in terms of the standard deviation. The RFR has the lowest MAE with the exception of the Push and Pull technique across both CM and IM. This is expected since RFR is known for more efficient feature selection and fitting.

Figure 6c and Figure 7c show the Mean Absolute Percentage Error (MAPE) in the estimations; lower value indicates a more efficient model and RFR has the lowest MAPE.

Across all metrics of evaluation, RFR is a superior estimation model. The results for MAE and MAPE are surprising given that although RR applies a penalty for reducing the errors, it is less accurate in its prediction.

2) *Fog-to-Cloud offload:* The aim of presenting results for offloading from the Fog-to-Cloud is to examine whether any

(a) R^2 of the prediction models

(b) Mean Absolute Error (MAE) of the prediction models

(c) Mean Absolute Percentage Error (MAPE) of the prediction models

Fig. 7: k-fold validation approach for Cloud-to-Fog offload

additional factors influence the cost of offloading, such as the cost of data preparation on the Fog due to its limited resources when compared to the cloud. The results for R^2 obtained for the Fog-to-Cloud offload scenario are similar to the Cloud-to-Fog scenario (Figure 8a and Figure 9a for the two validation approaches). The observations are similar to those noted for the Cloud-to-Fog scenario. RFR emerges as a superior model in relation to MAE (except for Push and Pull; Figure 8b and Figure 9b) and MAPE (Figure 8c and Figure 9c). It is noted that the offloading results obtained from Cloud-to-Fog offloading and vice versa in our experimental setup are similar with small variations that do not affect the offloading model.

C. Summary

The following summarises the experimental results:

(i) RFR has superior performance compared to the MLR, PMR and RR. These approaches with the runtime and offline parameters were considered to address Q1.

(ii) With respect to R^2 , it is noted that: (a) for the Cloud-to-Fog scenario between 72%–99% of the chosen input parameters (runtime and offline) influence $t_{offload}$, and (b) for the Fog-to-Cloud scenario between 71%–95% of the parameters

(a) R^2 of the prediction models

(b) Mean Absolute Error (MAE) of the prediction models

(c) Mean Absolute Percentage Error (MAPE) of the prediction models

Fig. 8: Train-test validation approach for Fog-to-Cloud offload

affect the offload time. This addresses Q2 posed initially by determining the influence of input parameters on offloading.

(iii) MAE as low as 1.4 seconds is obtained for Cloud-to-Fog and 1.1 seconds for Fog-to-Cloud. The lowest MAPE for Cloud-to-Fog is 4.21% and for Fog-to-Cloud is 3.19%. The MAE and MAPE measures provide insight into the estimation accuracy of different models to address Q3.

(iv) It is known that the CM approach is easier to use in practice than the IM approach (since a collection of models are required). However, it is demonstrated that accuracy obtained using the CM method is comparable with the IM method.

V. RELATED WORK

Fog offloading is necessary to meet the overall Quality-of-Service requirements of an application. This section considers three relevant areas of Fog offloading, namely the scenarios, approaches, and the decisions to be made for offloading.

1) *Offloading Scenarios*: Offloading offers the benefits of meeting the computational requirements of an application and latency demands, load balancing, and managing energy consumed [2], [15]. Section II highlighted three offloading

(a) R^2 of the prediction models

(b) Mean Absolute Error (MAE) of the prediction models

(c) Mean Absolute Percentage Error (MAPE) of the prediction models

Fig. 9: k-fold validation approach for Fog-to-Cloud offload

scenarios, namely Cloud-to-Fog, Fog-to-Cloud, and Device-to-Fog with the offloading approaches.

Cloud-to-Fog: Current techniques to partition a monolithic application (i.e. not designed as a micro-service based application) are manual [5]. Multiple approaches are adopted to make decisions on Fog placement. Examples include a heuristic-based scheduling algorithm to balance makespan and the monetary cost of Cloud resources has been proposed [16].

Fog-to-Cloud: Fog resources are anticipated to be hardware constrained and geographically dispersed when compared to a Cloud data center [2]. Therefore, services may not be able to easily scale across the Fog if they have significant compute or storage requirements. Thus, a Fog service may need to be offloaded back to the Cloud. One strategy is to employ an analytical queuing model that is based on the LRU filter [7].

Device-to-Fog: Osmotic computing provides an architecture to deal with Device-to-Fog offloading [3]. Major concerns that need to be addressed while offloading from user devices or sensors onto the Fog include minimising delays [8], data transferred to the Cloud [17], and communication overheads [18].

2) **Offloading Approaches:** Virtual Machine (VM) and containers are the two main approaches considered in the literature

for offloading services. VMs when compared to containers have larger overheads and are generally used in resource abundant environments, such as the Cloud. In the context of limited compute and storage resources as seen in the Fog, containers may be more appropriate for offloading services. Nonetheless, VM placement in the Fog by taking user mobility into account using integer linear programming has been proposed [19].

The majority of research in Fog computing that focuses on offloading takes containers into account. Service hand-off for offloading services between Fog resources is proposed [14] and live container migration using CRIU is considered [20].

3) **Offloading Decision:** Three aspects relevant to the decision-making process for Fog offloading need to be considered: (i) *Which services need to be offloaded* - an application may be a composition of multiple services and not all services would equally benefit from being offloaded. Therefore, the subset of services that can benefit from offloading needs to be identified. (ii) *Where to offload services* - multiple edge resources may be available within the same geographic region (for example, a street or city) to offload services. The most appropriate resource that can execute the service to meet all service objectives needs to be determined. (iii) *How to offload services* - there are multiple techniques for offloading services. For example, when using containers the most appropriate approach (for example with least overhead) needs to be chosen from multiple stateless and stateful techniques (Section II).

Which services to offload: Service selection has been formulated using Mixed Integer Non Linear Programming and solved using linearisation and genetic algorithms [21]. Search-based algorithms have also been considered to determine eligible services of an IoT application by taking into account the hardware capacity of Fog resources [22]. A QoS-aware model is proposed for deploying IoT services to Fog infrastructure. The model considers the available infrastructure (latency and bandwidth), interactions among software components, and business policies during the selection of services [23].

Where to Offload: Fog nodes are selected using multiple criteria mixed-integer linear programming [24] and Markov-Decision-Process [25] approaches. The FOLO framework performs dynamic task allocation in vehicular Fog computing that takes the resource constraints of Fog resources [26].

How to Offload: Low end-to-end latency for mobile users is achieved by a service hand-off mechanism using live container migration that is achieved using CRIU [14]. Cloud4IoT is a platform that performs vertical and horizontal migration of IoT services using a stateless approach, namely Push and Pull [27]. A Cloud to Fog based model to predict offloading time for stateless container technology has been proposed using machine learning techniques [28]. Live container migration techniques, namely Pre Copy, Post Copy, and Hybrid have been considered in the context of Fog computing environments [29]. Stateful and stateless migration with low down time has been investigated [14], [29], [30].

The research above presents techniques for optimising offloading, service placement in the Fog and selecting the offloading approach. The research in this paper aims to charac-

terise offloading by estimating the time taken to offload when multiple stateless and stateful techniques are employed.

VI. CONCLUSIONS

This paper proposes the design and implementation of methods for estimating the offload time using containers in Fog computing. Four container-based offloading technique (three stateless techniques, namely Save and Load, Export and Import and Push and Pull, and one stateful technique, namely CRIU-based Live Migration) are investigated. Four estimation models based on Multivariate Linear Regression, Polynomial Multivariate Regression, Random Forest Regression and Ridge Regression are considered to estimate the offload time. A catalogue of 25 parameters collected at the system and process levels during runtime and offline are used as input to the models. Two estimation methods, namely using a collective model and individual models are proposed. Experimental studies are pursued on two Cloud-Fog lab-based platforms from which 42 million data points are obtained. The results are presented in relation to: (i) the influence of the input parameters on the estimation models measured by the coefficient of determination for regression, and (ii) the accuracy of estimation measured by the mean absolute error and mean absolute percentage error. It is noted that the Random Forest Regression (RFR) has superior performance compared to the other approaches. No specific trend was observed in relation to the collective or individual models. The input parameters are appropriate for the estimation models since up to a 99% influence is observed on the offload time. Moreover, it is noted that the RFR can estimate the offload time with less than a 10% mean absolute percentage error.

In the future, an integrated decision-making approach that considers: (i) which services of Directed Acyclic Graph based applications should be offloaded, and (ii) given multiple Fog resources where should they be offloaded will be developed.

ACKNOWLEDGMENT

The first author is funded by a Schlumberger Scholarship.

REFERENCES

- [1] B. Varghese and R. Buyya, "Next Generation Cloud Computing: New Trends and Research Directions," *Future Generation Computer Systems*, vol. 79, pp. 849–861, 2018.
- [2] C.-H. Hong and B. Varghese, "Resource Management in Fog/Edge Computing: A Survey on Architectures, Infrastructure, and Algorithms," *ACM Computing Survey*, vol. 52, pp. 97:1–97:37, 2019.
- [3] M. Villari, M. Fazio, S. Dustdar, O. Rana, and R. Ranjan, "Osmotic Computing: A New Paradigm for Edge/Cloud Integration," *IEEE Cloud Computing*, vol. 3, pp. 76–83, 2016.
- [4] M. Satyanarayanan, "The Emergence of Edge Computing," *Computer*, vol. 50, no. 1, pp. 30–39, 2017.
- [5] N. Wang, B. Varghese, M. Matthaiou, and D. S. Nikolopoulos, "ENORM: A Framework For Edge NNode Resource Management," *IEEE Transactions on Services Computing*, pp. 1–14, 2017.
- [6] J. McChesney, N. Wang, A. Tanwer, E. de Lara, and B. Varghese, "DeFog: Fog Computing Benchmarks," in *ACM/IEEE Symposium on Edge Computing*, 2019.
- [7] M. Enguehard, G. Carofiglio, and D. Rossi, "A Popularity-Based Approach for Effective Cloud Offload in Fog Deployments," in *30th International Teletraffic Congress*, vol. 01, 2018, pp. 55–63.
- [8] A. Yousefpour, G. Ishigaki, R. Gour, and J. P. Jue, "On Reducing IoT Service Delay via Fog Offloading," *IEEE Internet of Things Journal*, pp. 998–1010, 2018.
- [9] K. Ha, Y. Abe, Z. Chen, W. Hu, B. Amos, P. Pillai, and M. Satyanarayanan, "Adaptive VM Handoff Across Cloudlets," *Technical Report CMU-CS-15-113*, 2015.
- [10] S. Soltész, H. Pötzl, M. E. Fiuczynski, A. Bavier, and L. Peterson, "Container-Based Operating System Virtualization: A Scalable, High-Performance Alternative to Hypervisors," in *ACM SIGOPS Operating Systems Review*, vol. 41, 2007, pp. 275–287.
- [11] T. P. Pham, J. J. Durillo, and T. Fahringer, "Predicting Workflow Task Execution Time in the Cloud using A Two-Stage Machine Learning Approach," *IEEE Transactions on Cloud Computing*, pp. 1–13, 2017.
- [12] V. Nikolaenko, U. Weinsberg, S. Ioannidis, M. Joye, D. Boneh, and N. Taft, "Privacy-Preserving Ridge Regression on Hundreds of Millions of Records," in *IEEE Symposium on Security and Privacy*, 2013, pp. 334–348.
- [13] C. Puliafito, E. Mingozzi, F. Longo, A. Puliafito, and O. Rana, "Fog Computing for the Internet of Things: A Survey," *ACM Transaction on Internet Technology*, vol. 19, pp. 18:1–18:41, 2019.
- [14] L. Ma, S. Yi, and Q. Li, "Efficient Service Handoff Across Edge Servers via Docker Container Migration," in *Proceedings of the 2nd ACM/IEEE Symposium on Edge Computing*, 2017, pp. 11:1–11:13.
- [15] B. Varghese, N. Wang, S. Barbhuiya, P. Kilpatrick, and D. S. Nikolopoulos, "Challenges and opportunities in edge computing," in *IEEE International Conference on Smart Cloud*, 2016, pp. 20–26.
- [16] X.-Q. Pham and E.-N. Huh, "Towards Task Scheduling in a Cloud-Fog Computing System," in *18th Asia-Pacific Network Operations and Management Symposium*, 2016, pp. 1–4.
- [17] T. Yu, X. Wang, and A. Shami, "A Novel Fog Computing Enabled Temporal Data Reduction Scheme in IoT Systems," in *IEEE Global Communications Conference*, 2017, pp. 1–5.
- [18] X. Hou, Z. Ren, W. Cheng, C. Chen, and H. Zhang, "Fog Based Computation Offloading for Swarm of Drones," in *IEEE International Conference on Communications*, 2019, pp. 1–7.
- [19] D. Gonalves, K. Velasquez, M. Curado, L. Bittencourt, and E. Madeira, "Proactive Virtual Machine Migration in Fog Environments," in *IEEE Symposium on Computers and Communications*, 2018, pp. 742–745.
- [20] S. Nadgowda, S. Suneja, N. Bila, and C. Isci, "Voyager: Complete Container State Migration," in *37th International Conference on Distributed Computing Systems*, 2017, pp. 2137–2142.
- [21] A. M. Maia, Y. Ghamri-Doudane, D. Vieira, and M. F. de Castro, "Optimized Placement of Scalable IoT Services in Edge Computing," in *Symposium on Integrated Network and Service Management*, 2019, pp. 189–197.
- [22] R. Mahmud, K. Ramamohanarao, and R. Buyya, "Latency-Aware Application Module Management for Fog Computing Environments," *ACM Transaction on Internet Technology*, vol. 19, pp. 9:1–9:21, 2018.
- [23] A. Brogi and S. Forti, "QoS-Aware Deployment of IoT Applications Through the Fog," *IEEE Internet of Things Journal*, vol. 4, pp. 1185–1192, 2017.
- [24] R. A. da Silva and N. L. da Fonseca, "On the Location of Fog Nodes in Fog-Cloud Infrastructures," *Sensors*, vol. 19, pp. 2445–2465, 2019.
- [25] T. Taleb, A. Ksentini, and P. A. Frangoudis, "Follow-Me Cloud: When Cloud Services Follow Mobile Users," *IEEE Transactions on Cloud Computing*, vol. 7, pp. 369–382, 2019.
- [26] C. Zhu, G. Pastor, Y. Xiao, Y. Li, and A. Ylae-Jaeeski, "Fog Following Me: Latency and Quality Balanced Task Allocation in Vehicular Fog Computing," in *15th Annual IEEE International Conference on Sensing, Communication, and Networking*, 2018, pp. 1–9.
- [27] C. Dupont, R. Giaffreda, and L. Capra, "Edge Computing in IoT Context: Horizontal and Vertical Linux Container Migration," in *Global Internet of Things Summit*, 2017, pp. 1–4.
- [28] A. Majeed, P. Kilpatrick, I. Spence, and 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.
- [29] C. Puliafito, C. Vallati, E. Mingozzi, G. Merlino, F. Longo, and A. Puliafito, "Container Migration in the Fog: A Performance Evaluation," *Sensors*, vol. 19, pp. 1488–1510, 2019.
- [30] A. Machen, S. Wang, K. K. Leung, B. J. Ko, and T. Salonidis, "Live Service Migration in Mobile Edge Clouds," *IEEE Wireless Communications*, vol. 25, pp. 140–147, 2018.